Title
Generative Neural Activity Patterns in the Hippocampus

Author
Comrie, Alison Emelie

Publication Date
2024

Peer reviewed|Thesis/dissertation

Generative Neural Activity Patterns in the Hippocampus

by
Alison Emelie Comrie

DISSERTATION
Submitted in partial satisfaction of the requirements for degree of
DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Approved:

DocuSigned by:
3A5C1850C2F1412...
Michael Brainard
Chair

DocuSigned by:
Loren Frank
Loren Frank

DocuSigned by:
Joshua Berke
Joshua Berke

DocuSigned by:
12024A1A01564EB...
Alexandra Nelson

Committee Members

For my teachers

ACKNOWLEDGEMENTS

I am grateful to those who have contributed to this work, supported me in this work, and otherwise made this work a gratifying experience. Among them, I am especially indebted to:

Loren Frank, for your trust, patience, and imagination, which I continue to learn from, and aspire to extend to others. You bring a combination of quantitative dexterity, technical skill, and well-founded narrative building to neuroscience that I find extraordinary, making me lucky to have been trained by you. I also know that I am fortunate to have the freedom and responsibility that you have given me, and challenged me to make the most of, in my graduate training.

Nathaniel Daw, for your dependable and incisive input, no matter the time zone. It has been a privilege to enjoy your computational and theoretical perspectives at every step of the way from our experimental design through our analysis and writing.

Joshua Berke, for your critical questions and honesty, and for teaching me that one is not merely motivated to act but that acting is motivating. Thank you also for actively making our collaboration truly collaborative.

Michael Brainard, for your insistence and willingness to puzzle over what models we are testing, and for your openness.

Alexandra Nelson, for your rigorous thinking and practical insights, for listening, and for leading by example.

Emily Monroe, for your questions, your keen eyes, your deep appreciation for animals, and for teaching me to be a better mentor, supervisor, and scientist.

Ari Kahn and Tim Krausz, for your generous teamwork.

Abhilasha Joshi, Anna Gillespie, Mari Sosa, Eric Denovellis, Hannah Joo, and Kenny Kay, for being brilliant role models in science and beyond. Your friendship, moral strength, kindness, integrity, and willingness to share your time and expertise, have shaped me and my aspirations.

Anya Kiseleva, for your graceful management of the laboratory.

Past and current members of the Frank lab, for creating an intellectually playful environment, always pushing me to revisit analyses, and never running out of questions and new ideas.

Collaborators at SpikeGadgets, LBNL, LLNL, and the Flatiron Institute, including Mattias Karlsson, Maxim Borius, Ryan Ly, Allison Yorita, Razi Haque, and Jeremy Magland, for working together to develop new tools and methods that made our experiments, data management, and data analyses possible.

Dirk Kleinhesselink, for keeping our data safe and our computers running smoothly.

Viktor Kharazia for sharing your expertise in anatomy and histology.

The LARC staff and especially Ellie Karlsson, for helping us keep our animals healthy.

The staff who serviced Sandler, and Frank Plut, for keeping our spaces clean and functional, and for always explaining to me how the building works as I built and denoised experiments.

Kevin Bender and Mazen Kheirbek, for enthusiastically training me in my rotations, and nurturing my interest in neurophysiology.

The Neuroscience Graduate Program administration, including Anatol Kreitzer, Pat Veitch, and Lucita Nacionales, for creating an empowering training environment.

Those in the Office of Career and Professional Development, Department of Physiology, and Office of Sponsored Research, who have supported me professionally.

My classmates at UCSF, the Woods Hole Methods in Computational Neuroscience community, the Fog at Bay producers, the Brain Camp team, and especially Don Woodson, for creating a better future for neuroscience.

Carol Barnes, for welcoming me into your laboratory as a high school student and ensuring that I have never really left the lab, giving me experience in a wide range of experimental approaches that shaped my perspectives on neuroscience and cognition, and developing my appreciation for studying the healthy brain and its natural proclivity for change.

Peter Somogyi, for sitting me down at a microscope with Golgi-stained slides and waiting for me to figure out what I was looking at myself, for teaching me about the power of long-range inhibition, for always ensuring that I wrote up my observations, and for appreciating nature’s diversity from cells in the brain to flora and fauna outside.

Leslie Tolbert, for your public lecture on neuroplasticity that sparked my interest in neuroscience before I knew what research was, for letting me wander into your office to stare in awe at microscope images on your wall as I began to find out, and, once I became your student, for encouraging me never to settle.

My university professors, especially Alan Nighorn, Lynne Oland, John Hildebrand, Gail Burd, and Lynn Nadel, for being sources of constant encouragement to follow my curiosity both in and outside of the classroom and lab.

The leadership of the KEYS, UBRP, and BRAVO Programs, including Carol Bender, Jennifer Cubeta, and Marti Lindsey, for making it possible for me to actively participate in research as a student, and Emily Kotay, for training me to apply for grants and support my own work.

My grade school educators, especially Ms. Robinson, Mr. Bindschadler, Ms. Shane-Boyd, Mr. Erbe, and Ms. Keri, for valuing communication, fostering creativity, encouraging grit, teaching me to teach myself, and showing me that we can be part of something bigger than ourselves.

The rats who participated in my training and experiments, for providing company and comfort through the loneliest days and sometimes nights of the pandemic, and whose creativity, curiosity, and persistent hypothesis testing should serve as an example for scientists at large. Thank you for contributing your lives to this work.

My grandparents, and the generations of our family before, for giving me the privileges I have in my life and work today, and especially Harold and Emelie Loewhenheim, for instilling in me the values that really matter.

My sister, for your companionship and your example of determination.

My parents, for raising me to be fascinated by nature, and for making everything possible.

Daniel Silversmith, for our love, and for embracing me fully along with my work. No other discovery could possibly compare.

And those whose names I have not included, for whom I am also deeply grateful.

This work was supported by NSF GRFP 1650113 and NIH F31MH124366, and thereby the taxpaying American public, as well as the UCSF Jonas Cohler Discovery Fellowship.

CONTRIBUTIONS

All chapters were written by Alison E. Comrie (A.E.C.) and edited by Loren M. Frank (L.M.F.).

Chapter 1 is reproduced in adapted form from:

> Comrie, A. E., Frank, L. M., Kay, K. (2022). Imagination as a fundamental function of the hippocampus. *Philosophical Transactions of the Royal Society B, 377*(1866), 20210336. https://doi.org/10.1098/rstb.2021.0336.

> Conceptualization, A.E.C., L.M.F., K.K. Writing – Original Draft, A.E.C., Writing – Review and Editing, A.E.C., L.M.F., K.K. Funding Acquisition, A.E.C., L.M.F, K.K.

Chapter 2 contains research performed and written by Alison Comrie, and is reproduced in adapted form from an archived preprint:

> Comrie, A.E., Monroe, E.J., Kahn, A.E., Denovellis, E.L., Joshi, A., Guidera, J.A., Krausz, T.A., Berke, J.D., Daw, N.D., Frank, L.M. (2024) Hippocampal representations of alternative possibilities are flexibly generated to meet cognitive demands. *bioRxiv*. 2024.09.23.613567.

> Conceptualization, A.E.C., L.M.F.; Methodology, A.E.C., A.E.K., T.A.K., J.D.B., N.D.D., L.M.F.; Software, A.E.C., E.J.M., A.E.K., E.L.D., N.D.D., L.M.F.; Formal Analysis, A.E.C., E.J.M., A.E.K.; Investigation, A.E.C., E.J.M., A.J., J.A.G.;

Resources, L.M.F.; Data Curation, A.E.C., E.J.M., E.L.D., L.M.F.; Writing – Original Draft, A.E.C., L.M.F.; Writing – Review & Editing, A.E.C., E.J.M., A.E.K., E.L.D., A.J., J.A.G., T.A.K., J.D.B., N.D.D., L.M.F.; Visualization, A.E.C., A.E.K.; Supervision, L.M.F.

GENERATIVE NEURAL ACTIVITY PATTERNS IN THE HIPPOCAMPUS

ALISON EMELIE COMRIE

ABSTRACT

Animals can leverage prior experience to guide adaptive decision making. To decide where to forage, for instance, an animal may recall previous locations and internally simulate paths to take next. These functions are thought to rely on the hippocampus, a brain structure long implicated in learning, memory, and navigation. Accordingly, hippocampal neural activity can represent an animal's current position, as well as generate representations of alternative possible locations. These representations of alternative "non-local" possibilities are hypothesized to enable internal simulation of previous experiences, alternative pasts, and potential futures to support cognition and, in turn, experience-guided decision making. However, it remains unclear whether or how internally generated hippocampal non-local representations are regulated during active behavior depending on changing cognitive needs for learning about and deciding among alternatives. In this work, I first synthesize evidence describing hippocampal non-local representations that suggests that they are well-suited to serve a wider range of cognitive abilities than previously thought. This work advances the idea that hippocampal function is well characterized not only by its representation of actual experience, but also by its regular representation of alternatives to actual experience. I then present experimental findings that show that the hippocampus generates representations of a wide range of spatial possibilities during active navigation, and that

representations of these distinct possibilities are distinctly modulated with learning and decision making in a complex and dynamic foraging environment. These findings indicate that the brain regulates the generation of alternatives in the hippocampus to meet momentary cognitive demands for adaptive behavior.

TABLE OF CONTENTS

TABLE OF FIGURES

CHAPTER 1: IMAGINATION AS A FUNDAMENTAL FUNCTION OF THE HIPPOCAMPUS

Summary

Imagination is a biological function that is vital to human experience and advanced cognition. Despite this importance, it remains unknown how imagination is realized in the brain. Substantial research focusing on the hippocampus, a brain structure traditionally linked to memory, indicates that firing patterns in spatially tuned neurons can represent previous and upcoming paths in space. This work has generally been interpreted under standard views that the hippocampus implements cognitive abilities primarily related to actual experience, whether in the past (e.g., recollection, consolidation), present (e.g., spatial mapping), or future (e.g., planning). However, relatively recent findings in rodents identify robust patterns of hippocampal firing corresponding to a variety of alternatives to actual experience, in many cases without overt reference to the past, present, or future. Given these findings, and others on hippocampal contributions to human imagination, we suggest that a fundamental function of the hippocampus is to generate a wealth of hypothetical experiences and thoughts. Under this view, traditional accounts of hippocampal function in episodic memory and spatial navigation can be understood as particular applications of a more general system for imagination. This view also suggests that the hippocampus contributes to a wider range of cognitive abilities than previously thought.

Introduction

The ability to imagine is essential to human experience. At a broad level, imagination has a major role in human creativity, agency, and everyday thoughts and actions. More specifically, humans have and express many types of imagined experiences. These include recollections, predictions, simulations, counterfactuals, fantasies, suppositions, and mind-wandering – and, in pathological cases, hallucinations and confabulations. These wide-ranging forms of imagination are relevant, if not essential, to a similarly wide range of cognitive domains, such as memory, planning, learning, and inference. Despite this fundamental importance, our understanding of how imagination is realized as a biological process in the brain remains nascent. Indeed, the sheer diversity of imagined experiences makes it challenging to begin to envision a possible biological approach.

As starting point, we identify a unifying characteristic of imagined experiences: they do not refer to actual present experience, or directly reflect ongoing circumstances in the external world. Rather, imagined experiences refer to non-actualities, and arise from a source internal to the subject. Awake healthy subjects can, in other words, "mentally" self-generate thoughts and experiences and distinguish them from thoughts and experiences driven by ongoing stimuli in the actual present. We refer to this fundamental ability to generate possibilities that do not correspond to the actual present as generativity. By this definition, generativity is a basic function that underlies imaginative abilities broadly, regardless of more specific properties, such as references in time (e.g., remembering the past or simulating futures). As further clarification, we also note that our present use of "generativity" differs from its senses in linguistics and

in statistical models (notwithstanding potential connections between these uses[1–3]). Defining generativity enables us to focus on a single characteristic ability that may ultimately facilitate our understanding of the diverse types and components of imagination.

Crucially, generativity can be understood at the level of the brain. Mirroring the subject-level ability to distinguish actual from imagined experience[4], specific neural processes in the healthy brain must "parse" internal representations as ongoing experience (actual) versus internally generated alternative experience (imagined). Importantly, this substrate-level generativity does not presuppose features such as mental imagery, mental time travel, or conscious awareness. Indeed, defining generativity enables us to refer to the brain's capacity to internally generate experiences that are distinguished from externally driven present experience, without invoking these features that are associated with subjective human imagination. As an example, a soccer player approaching a moving ball can rapidly assess numerous dynamic ongoing events and stimuli, consider multiple possible responses, and decide on a play, all in a split second and without overt awareness of each internally represented possibility. In animals, ethologically relevant scenarios such as predation and escape make similar demands on cognition[5]. Thus, direct investigation of the brain may be essential to understand generativity.

In this review, our overall aim is to describe and advance our understanding of how generativity – an ability underlying imagination – is realized in the brain. Our review is guided by five questions: (1) where generativity might be implemented in the brain, (2) how generative neural activity can be identified, (3) what candidate generative

neural activity patterns and representational correlates have been previously described, and (4) how the brain can organize actual versus generative activity patterns. This discussion establishes that the hippocampus, a brain structure in the medial temporal lobe, is a candidate biological substrate of generativity, and that patterns of hippocampal neural firing reflect generative processes by representing a diverse range of alternatives to ongoing experience. Finally, we consider (5) what these observations suggest about the biological basis of generativity and its role in cognition. More specifically, in light of recent findings at the level of neuronal firing patterns in rodents, in addition to brain research related to imagination in humans, we suggest that the hippocampus – often understood as a system that characteristically represents actual experience, whether in the past, present, or anticipated future – may be better understood as a system that also represents imagined alternatives to actual experience.

The Hippocampus as a Locus of Generativity in the Brain

What structure within the brain might implement generativity? One approach to this question is to determine whether damage to specific parts of the brain causes deficits in imaginative abilities relying on generativity, including recollecting the past, envisioning the future, or constructing fictional scenarios. Notably, the earliest case studies linking imagination of the future to specific brain areas are in individuals with previously established deficits in memory of the past[6–10]. In one classic case, patient H.M. suffered severe amnesia after his hippocampus and adjacent medial temporal areas were surgically removed, which established the hippocampus as an important site for memory, particularly episodic memory[11,12]. Notably, while episodic memory impairments are most traditionally reported, H.M. and many other patients with hippocampal damage

have since been examined and found to have severe impairments in future-oriented thinking and constructing fictional events more generally[9,13–18]. These findings raise the possibility that recollection of the past, anticipation of the future, and imaginative abilities more broadly may share common underlying functions as well as dependence on the hippocampus[17,18].

Complementing lesion studies, functional brain imaging has revealed activation of the hippocampus during a variety of self-reported imagined experiences that overtly differ from subjects' actual circumstances[19–22]. In such studies, subjects are typically asked to imagine experiences that differ from present experience through changes in time, space, and/or personal perspective. The hippocampus, in addition to a group of cortical areas known as the default mode network, is consistently activated during, for instance, recalling autobiographical experiences, imagining anticipated future episodes, imagining counterfactuals, mentally simulating common activities (e.g., brushing teeth), constructing fictional scenes, imagining non-actual events and stories, taking on others' perspectives, and unprompted mind-wandering[19,20,23–27]. These results highlight that the hippocampus, along with other brain regions in the default mode network, is important for the capacity to generate mental displacements from actual present circumstances, whether in time, space, personal perspective, and possibly other domains[13,18,19,28,29]. Thus, although the cognitive role of the hippocampus is often conceptualized in relation to prior experience (i.e., episodic recollection, recall) or explicitly anticipated experience (i.e., planning, prospection)[30–32], the hippocampus appears to play a more general role in imaginary experience[28].

In efforts to clarify this role, studies have often probed the availability and character of mental imagery. Several further studies help refine the role of the hippocampus beyond the observation mentioned above that hippocampal damage is associated with deficits in vividly visualizing fictional scenes. First, patients with partial hippocampal lesions show activation of residual hippocampal tissue when tasked with imagining complex scenes[33,34]. Second, one patient with longstanding hippocampal damage found it effortful but possible to visualize single imaginary objects and simple scenes, yet could not readily imagine complex scenes in one automatic and coherent picture – instead, he built up the scenes "bit by bit"[34]. Residual hippocampal tissue in this patient was not activated during these tasks as it was in control participants[34]. These findings suggest that the hippocampus is not strictly required for mental imagery, and therefore that the role of the hippocampus in imagination may be only indirectly related to mental imagery. The requirement of the hippocampus for readily constructing complex scenes in particular suggests a different basis or principle by which the hippocampus contributes to imagination [34]; we revisit this issue in the section "Generativity as a function of the hippocampus."

The above lesion and functional imaging work implicates the hippocampus as a candidate substrate for generative thinking, typically by relying on conscious verbal or behavioral reports. This approach is, however, limited in addressing how generative processes are implemented at a neuronal level. For example, the timing of underlying processes relative to eventual behavioral reports remains unclear. Generative processes may also unfold at timescales considerably faster than behavior, which suggests the need for complementary approaches with finer temporal resolution. Here

animal models provide an important advantage by enabling greater access to neural firing. This potential approach in turn raises the question of whether animals also exhibit behaviors indicating generative thought, and if so, whether the hippocampus is also implicated, as in humans.

From work dating at least a century, it is clear that animals behave based on memory of prior experience and conceptual insight rather than solely trial and error, instinct, and presently sensed information[35–37]**.** This implies a corresponding ability to construct and use internal representations and suggests the existence of generative neural processes in animals. In the case of rats, a common model for hippocampal studies, a seminal example of behavior based on internal representations is spatial navigation. When navigating, rats can take novel paths (for instance, short-cuts to goal locations), implying an internal model enabling the ability to generate such novel courses of action[38,39]. Rat behavior can also appear deliberative and regretful, suggestive of internally generating representations of possibilities, including counterfactual pasts[40–42]. In service of these and other behaviors, the hippocampus is thought to be essential for using an abstract internal model, or "cognitive map," that relates items, events, and features of experience[40,43,44]. Indeed, hippocampal damage impairs various behaviors thought to rely on abstract internal representations such as rats' abilities to infer relationships between stimuli[45]. Further, hippocampal lesions impair rats' abilities to make choices dependent on an internal model and predictions or plans made by that model[46]. These findings suggest that the hippocampus is an important locus in the rodent brain for constructing abstract mental models, which in turn could be

used to generate representations of prior, new, and otherwise not presently experienced possibilities, enabling insightful behaviors.

With the hippocampus as a starting point for investigating generativity in both humans and animals, we now aim to clarify what neural firing patterns have been observed in the hippocampus and what internal representations they suggest. To do so, it is necessary to address our second question: how can generative neural activity patterns be identified?

Identifying Neural Firing Patterns That are Generative

Identifying neural firing patterns that may represent imagined experiences requires us first to identify neural firing that corresponds to actual experience. Here, we focus on studies of neural firing in the rodent hippocampus. To investigate internal representations at the level of neurons, neurobiologists have leveraged the well-established relationship between spatial location and hippocampal firing in freely moving rats[47]. Over fifty years of work have established that principal neurons in the rodent hippocampus exhibit increased firing rates when the animal is in distinct physical locations[47,48] (Figure 1.1). As the rat moves through an environment, each of these "place cells" consistently increases its firing rate when the animal is in the neuron's "place field" location(s)[47,48]. Importantly, place cell firing also varies based on numerous factors besides location[49]; for example, in linear environments, a large proportion of place cells fire more when the animal is traveling in a particular direction[50]. Therefore, at a broader level, it is important to note that a place field describes average firing over many individual runs through a location, even though there is often substantial variability in a place cell's firing across individual runs through the same place (Figure 1.1).

The basic notion of a place field, along with the ubiquity of place cells in the rat hippocampus, provides a possible approach to identifying actual and generative activity at a neural level. If we take a place cell's activity to represent its place field location, then each instance of firing by that neuron can be provisionally understood as representing that location. By this interpretation, a place cell will reliably fire when the animal is in the cell's place field, thereby representing the animal's actual present location.

Importantly, in certain moments, a place cell can also fire when the animal is not actually in the cell's time-averaged place field location[51–53] (Figure 1.1). Accordingly, these moments can be provisionally understood as times in which a representation of the place field location is internally generated, even though the animal actually occupies a different location at that moment.

Strikingly, place cells have been found to fire outside of their place fields in coordination with each other[54,55] (Figure 1.1). During these events, the collective activity of place cells can be understood to express a representation that corresponds to locations different from the animal's current location[52,53]. In other words, this neural firing is consistent with a generative representation; while it appears displaced from the animal's actual state and present stimuli, it is internally coordinated across cells (Figure 1.1).

A variety of analysis methods have been used to investigate these generative firing events and internal spatial representations in the hippocampus[56–58]. Briefly, one approach is to model the firing of many individual place cells as their time-averaged place field locations, and then invert that model to produce an estimate of the neurally-

represented location at each moment in time[59–61]. Doing so enables us to infer, or decode, the animal's moment-to-moment "mental location" based on hippocampal firing patterns. Thus, by identifying periods when the decoded representation of location (or direction) differs from the animal's actual state, we can examine periods when hippocampal activity is collectively inconsistent with a representation of actual experience and may instead be generative. This enables us to address our third question: what kinds of generative representations have been observed in the hippocampus?

Generative Representations in Hippocampal Neural Firing

Single-cell and population decoding approaches have revealed a striking variety of putative generative representations in the rat hippocampus over the past several decades[62–65]. Traditionally, these representations have been accounted for as specific episodes and abstracted experiences that are based on the past, or that anticipate experiences in the future[66,67]. Recent results, however, imply that the hippocampus also regularly represents alternatives to actual experience, whether in the past, present or anticipated future[68–70]. Together, these findings suggest that the hippocampus may generate a substantially wider range of internally constructed alternatives to the animal's actual experience than traditionally understood.

Representations Consistent With Past Experiences

The first reports of hippocampal activity patterns related to past experiences focused on sleep[51,54]. Firing sequences of place cells that were active during running on a maze

were found to reactivate in similar sequential order during subsequent sleep, as if briefly "replaying" past spatial experience[71–73]. These replays occur on the order of tens to hundreds of milliseconds, far faster than the seconds-long timescale over which the actual behavioral traversal of those locations unfolds[71] (Figure 1.1). Importantly, replay events were subsequently found also to occur during waking periods in which rats are behaviorally immobile, such as sitting still or eating[74,75] (Figure 1.1). During wake and sleep, replay typically occurs during a burst-like hippocampal network-level activity pattern, the sharp wave-ripple (SWR), that is itself internally generated (rather than externally driven), consistent with the notion of generativity[76].

As suggested by its name, replay has been interpreted as recapitulating specific episodes of prior experience. An early observation was that after an animal ran towards and then came to rest at a reward location, a path was replayed starting at the animal location and proceeding in reverse, as if retracing the path that led to the reward[74,77,78]. Replay representations not only initiate at a stationary animal's location[75], but can also correspond to paths that start farther away from the animal within the current maze, as well as on a different maze experienced beforehand[79,80] (Figure 1.1). These examples are evocative of the hippocampus' long hypothesized role in cognitive functions that rely on experiences from the past, such as memory consolidation and episodic recall[62,81].

Additional findings on replay suggest a more complex picture. Unlike a rigidly recapitulative process that uniformly represents recent experiences, replay can be enriched for previously taken paths associated with reward, paths associated with aversive outcomes, nearby locations, and paths that have not recently been taken[61,82–84]. Further, these and several additional findings[82,84–88] suggest that replay events are

collectively well described as reflecting an abstract internal spatial model of the encountered environment, or a spatial “cognitive map”[43,53,64]. For instance, replays can be biased toward paths that are less behaviorally traversed, and replays can be consistent with random trajectories through a familiar space[87,88]; replays like these may sample locations that are not the most behaviorally salient or the most physically occupied to support the maintenance of a flexible model of the environment, and this function could help explain why replays are inconsistent with a rigid recapitulation that passively records recent experience[82,87,88]. These reports suggest that replay, instead of directly reinstating specific episodes, may abstractly reflect past experience via an internal spatial map.

While there is little doubt that replays can be derived from prior experience, both in the case of a rigid recapitulation or abstract model based on the past, what remains unclear is whether neural processes within or beyond the hippocampus interpret replay events as temporally situated in the past. For example, a replay of recently traversed locations behind the animal, that are not subsequently traversed, is better correlated with past than future behavior, but this does not rule out the possibility that this replay represented a potential future traversal of those locations, or a spatial sequence without a projection in time. Despite this ambiguity, replay can indeed be related to prior behavioral experiences. And, moreover, these findings on replay exemplify how generative activity in the hippocampus can represent various possibilities that differ from the actual present – here, in the form of spatial paths in known environments.

In parallel to replay during rest, neural firing in the hippocampus during movement has also been suggested to be recapitulative. During movement, an

internally generated network-level activity pattern, the 8 Hz theta rhythm, is observed throughout the rodent hippocampus[89–92]. Place cells are known to fire systematically in relation to the theta rhythm, such that neurons with place fields behind, at, and ahead of the animal fire at early, intermediate, and later phases of theta cycles, respectively[55,93,94]. Accordingly, collective place cell firing during a single cycle can represent a series of locations consistent with sweeping from the immediate past and present ahead to anticipated future locations[63] (rightmost example in Figure 1.1). Although firing in early phases of the theta rhythm can recapitulate locations just traversed by the animal, this firing appears to be consistent with the immediate actual past (for instance, as opposed to alternative past (counterfactual) locations)[63,95]. This suggests that early theta phase representations may also be best understood as reflecting actual experience, and not possible experience. That said, hippocampal firing during movement can correspond to locations behind the animal and is often thought to reflect the recent past[52,96,97].

Representations Consistent With Anticipated Futures

Place cell firing can also correspond to upcoming spatial paths, suggesting that generative representations may anticipate future experience. As introduced above, place cells firing in late phases of theta cycles tend to have place fields in locations ahead of the animal[52,55]. The extent to which this activity projects ahead of the animal can correlate with the distance the animal subsequently traverses, consistent with the possibility of future anticipation or prediction[98]. When multiple paths are available (such as a path bifurcating into two), hippocampal firing has been found to proceed ahead

along only one path at a time[69,99]. Furthermore, place cell firing corresponding to the left or right path ahead can occur on interleaved theta cycles, consistent with serially representing alternatives[69] (Figure 1.2). These internally generated representations are consistent with generatively representing anticipated possibilities, and are reminiscent of deliberation[99]. However, while theta-associated neural firing can predict the animal's subsequently taken path[99–101], firing patterns associated with alternation between paths fail to reliably predict the animal's subsequent choice[69,99].

Apart from generative activity associated with theta, replays suggestive of anticipated future experience have also been reported. In early work, replay was found to correspond to sequences of locations starting near and projecting ahead of the animal, just prior to running along that same path in the linear maze, consistent with anticipation of upcoming experience[75,79,80]. Since then, several studies have reported that replay in environments with more options (an open arena or multi-arm maze) is biased toward goal locations that the animal subsequently visits[102,103]. While replay can indeed correspond to subsequently taken paths, recent work from our group shows that replay fails to predict upcoming choices[84].

Seeking to relate generative firing to behavioral episodes in subjects' past or future (e.g., the choice of maze arm in the previous or next trial) has been a common approach in investigating the contributions of hippocampal activity to cognitive functions, especially past-oriented functions such as episodic recall and future-oriented functions such as planning. Task paradigms that disambiguate prior from upcoming experience are well-suited for this approach[84]. However, relating generative neural activity to particular locations behaviorally occupied in the past and future does not necessarily

indicate that such activity is an internal representation that refers temporally to the past or future. For example, neural firing corresponding to one of two paths ahead of the subject is consistent with a possible future, yet may also reflect recall of a prior traversal of that location, or simply not have any reference in time. In this sense, it remains an open question whether generative firing patterns observed in the hippocampus can refer to experiences projected into the future. Apart from this, it remains the case that some instances of generative firing during theta and replay can correspond to potential future locations, and may thereby contribute to explicitly anticipatory functions such as planning.

Representations Consistent With Alternative Possibilities

Firing patterns corresponding to locations different from a subject's actual location indicate that the hippocampus can generate representations of alternatives to actual ongoing experience. As discussed above, it has been hypothesized that these firing patterns reflect internal representations referring to episodes of experience in the past or anticipated future. Critically, recent findings indicate that generative firing patterns exhibit properties that may be more consistent with an underlying process that generates representations of non-actual hypotheticals and possibilities more broadly, rather than a process characterized primarily by projection of actual experience in time[68–70].

In recent work focusing on periods of movement, we found that neural firing in the rat hippocampus can regularly represent various alternatives with striking speed and regularity[69] (Figure 1.2). In initial observations, we found that alternative locations ahead

of moving animals could be represented not only as quickly as the frequency of the theta rhythm (~125 millisecond cycles), but also sustained across many consecutive theta cycles[69] (Figure 1.2). As in previous work showing that place cell activity can serially alternate between upcoming paths[99], or correspond to paths subsequently taken[98,102], one possibility is to interpret this pattern of neural firing as reflecting an essentially anticipatory function, such as planning or deliberating over future behavior.

However, we also found that place cell firing corresponding to opposite directions of travel exhibited the same pattern of serial alternation: sustained 8 Hz cycling between the animal's actual direction and an alternative, or non-actual, direction[69] (Figure 1.2). Toward clarifying what this pattern of alternating activity might reflect about the underlying process in the hippocampus, we highlight three points of consideration.

First, this generative firing pattern has no overt or intuitive temporal reference. Unlike the case of alternative locations ahead of the animal, alternative direction is neither more consistent with upcoming experience, nor more consistent with previous experience. This was especially the case given the experimental setting, in which rats routinely traveled in either direction through maze as part of navigating in an alternation task[69] (similar to Figure 1.2). Thus, neural firing signaling the non-actual direction was just as plausibly a recollected past as an anticipated future. Importantly, this ambiguity regarding time extends further: it is also just as plausible that the firing pattern reflected a representation of a counterfactual past, an alternative present, or an experience with no specific reference in time. This last possibility is reminiscent of imaginative thoughts in humans which do not explicitly project experience into the past or future, but nonetheless differ from a subject's present circumstances. Without further knowledge, it

may be relatively parsimonious not to attribute temporal reference to the observed hippocampal firing pattern – rather, a simpler interpretation is that this neural activity corresponded to non-actual experience.

Second, the speed of alternations between actuality and location or direction may be at odds with conscious human thought processes that are, at least subjectively, slower than ~125 ms theta cycles. For this reason, we speculate that a neural process at this speed is unlikely to be directly coupled to conscious awareness, such as during a human subject's internal deliberation over two choices, or mental imagery of a remembered episode of navigating a path. Rather, these generative neural firing patterns suggest a function that, like generativity, is marked by moment-to-moment variability and productivity.

The third point is that this generative hippocampal activity, which alternated between possibilities not actually being presently experienced, was largely independent from behavior[69]. This was the case both for cycling of non-actual locations and direction. Specifically, generative alternating firing patterns occurred commonly across classes of locomotor behaviors (e.g., running, crawling, turning, head scanning, brief pauses in running). Additionally, the number of theta cycles corresponding to alternatives varied widely between instances of otherwise similar trajectories through the maze. Further, activity that cycled between to two paths ahead at a bifurcation did not reliably predict animals' upcoming turn behavior on individual run trajectories[69]. These observations suggest an underlying process that can be uncoupled from behavior at three levels: classes of behavioral state, behavioral trajectories, and upcoming behavioral choices.

These three points are complemented by an additional line of investigation. Examination of the timing of neural firing within theta cycles revealed a surprising commonality between firing corresponding to alternative locations and direction: for either representational correlate, firing corresponding to the non-actual circumstance occurred specifically in late phases of theta cycles[69] (Figure 1.2). This observation indicates that late theta phases, previously understood to contain firing related to upcoming paths[52,67], are not exclusive to locations ahead of the animal (Figure 1.3). Rather, late phases can also contain firing related to alternative direction, which raises the possibility that firing corresponding to alternative possibilities in any other domain encoded by the hippocampus may also be generated in late theta phases. Toward this, further observations from our group suggest two additional examples of generative firing during late phases of theta; firing may correspond to locations behind the animal on a path that was not just taken, suggesting an alternative past representation[69], and may also correspond to locations relatively far from the animal, not only locations immediately ahead[70]. Together, these findings illustrate that late-phase firing can correspond to multiple kinds of alternatives to actual ongoing experience (direction and various locations). This is surprising because it is not consistent with the canonical understanding of theta cycles as organizing a sequential representation of locations along a single path, sweeping from past to future locations, or behind to ahead of the animal in space and time[52,55,67,94] (Figure 1.3). Rather, these findings suggest revising the established view that hippocampal firing during late theta phases corresponds to locations immediately ahead of the animal. A more inclusive view is that late theta

phases may be enriched for firing related to a diversity of alternative possibilities and hypotheticals, including, but not limited to, anticipated experiences (Figure 1.3).

In the preceding discussion, we largely focused on recent findings regarding hippocampal neural activity associated with movement. Several parallel results indicate that replay events, occurring during periods of immobility, can also represent possible or hypothetical experiences that are not clearly recapitulative nor anticipatory. Replays can represent trajectories that link physically connected spatial paths that the animal has not traversed behaviorally, as if simulating short-cut paths[82,104]. Such synthesized trajectories were not directly experienced by the animal, and therefore are inconsistent with strict recapitulation of the past. Furthermore, in some cases subjects never took the short-cut paths, suggesting that these replays may not have been anticipatory. Recently, another study found that replay is biased to an unchosen path even when that path would not fulfill the animal's motivational state (i.e. biased to water when hungry, and food when thirsty). This finding is inconsistent with both replay of the recent past and of the immediate future[68]. In sum, generative neural firing in the hippocampus during both movement (theta) and rest (replay) may reflect a process that represents a diversity of possibilities that are alternatives to actual present experience (Figure 1.3, Figure 1.4).

Organization and Origin of Generative Activity in the Brain

Having reviewed multiple types of generative neural activity in the hippocampus, we turn to our next question of how generative representations may be organized and "parsed" from representations of actual, ongoing experience. One would expect that neural processes are in place to separate actual and generative activity to avoid their confusion, reminiscent of the subject-level ability to internally distinguish actual from

imagined experience[4]. Multiple organizational schemes are possible; different sets of neurons could participate in actual versus generative representations, these representations could occur at different relative times, or some combination of these schemes could take place.

Findings in the rodent hippocampus indicate that neural firing corresponding to actual and generative representations occur at different relative times that are internally determined[105]. Generative representations tend to occur not only with temporal separation from representations of actuality, but also in alignment with underlying network-level activity patterns in the hippocampus that are internally generated: sharp wave-ripples (SWRs) and the theta rhythm[69,106] (Figure 1.3). This results in a serial alternation of neural firing corresponding to actuality and generativity, or a temporal "multiplexing" of actual and generative representations in the brain.

This serial alternation is present across behavioral states. During immobility, neural firing corresponding to the animal's actual present location is maintained for prolonged periods, transiently suppressed during sharp-wave ripple events that typically contain generative replays (10s-100s of milliseconds), and then subsequently restored[106,107] (Figure 1.3).

Similarly, during movement and exploratory behaviors, neural firing corresponding to actual present and non-actual alternative experience, or actual and generative representations, occurs serially and in alignment with characteristic phases of the theta rhythm[1,69]. More specifically, early phases characteristically contain representations of the animal's actual past and present experience, while late phases may contain firing corresponding to a variety of hypothetical experiences, resulting in

alternating actual and generative representations[69] (examples in Figure 1.2, schematic in Figure 1.3). Furthermore, there are multiple levels of alternation between actual and generative activity during movement – representations not only alternate within ~125 millisecond theta cycles (e.g., actual and upcoming position), but also across consecutive theta cycles[69] (e.g., alternation of two possible paths ahead) (Figure 1.2). Additional findings are also consistent with the idea that multiple representations can be accommodated in the hippocampus via serial alternation at a sub-second timescale. For instance, studies in the rat hippocampus have reported theta-modulated "flickering" between representations of two environmental contexts, as well as dynamic switching between two spatial reference frames, and separate reverse and forward-ordered location sequences within theta cycles[108–110].

The organization of actual and generative neural firing in the hippocampus also extends to other brain areas, consistent with the engagement of a distributed network in these representations[20,111,112]. Network-level neural activity patterns underlying generative representations can be coherent across the hippocampus and prefrontal cortex during replays and along the theta rhythm, with some reports of concurrent expression of actual versus alternative location representations across both regions[107,113–117]. Additionally, some generative firing events in the hippocampus are not only coordinated with but also predicted by the activity of cells in the medial prefrontal cortex[70]. Numerous other cortical and subcortical areas also share coordinated firing patterns with the hippocampus, during both replay events and the theta rhythm [67,118–125]. Recruitment of a large network of brain areas during activity related to actual and generative experience appears to reflect brain-wide organization, and the question of

how firing patterns in other regions across the brain specifically contribute and respond to generative representations in the hippocampus remains an active area of research[116,124].

How might organized generative neural firing patterns in the hippocampus come about through hippocampal and extrahippocampal processes? This remains largely unknown, but some initial points can be made. First, one would expect generative firing patterns, which do not correspond to immediately ongoing circumstances, to arise primarily from internally driven activity patterns, as opposed neural activity driven directly by external stimuli. Consistent with this, generative events are observed during SWRs and in association with the theta rhythm – and both of these activity patterns are generated internally in the brain (spontaneously) rather than elicited by external stimuli[76,126]. More specifically, SWRs spontaneously occur during sleep in the absence of dynamic sensory stimuli and can be intrinsically generated in isolated hippocampal slices in vitro[76]. Hippocampal theta oscillations arise in vivo in coordination with a rhythm generator region, the medial septum, and can also be generated in isolated rodent hippocampus in vitro[127,128]. Furthermore, late phases of theta, during which generative representations tend to occur, are associated with increased recurrent network activity from within the hippocampus, and relatively weaker influence from cortical areas that are thought to provide multimodal information to the hippocampus[63,67,129,130].

While SWR and theta oscillations are understood to be internally generated and are associated with the occurrence of generative neural firing patterns in the hippocampus, the question of how specific groups of neurons (such as place cells with

overlapping place fields) are recruited during generative events remains open[131]. In addition to mechanisms that support SWR and theta generation, it is likely the case that input from brain regions beyond the hippocampus have a role in this process[67]. One possibility is that the activation of particular sets of spatially tuned neurons during generative events is guided by extrahippocampal areas, such as the prefrontal cortex, that are also implicated in the default mode network[20]. This possibility is consistent with evidence that cortical activity can predict generative spiking during theta oscillations several cycles in advance, as well as SWR activity during sleep, and would argue against the idea that hippocampal ensembles are activated by exclusively unstructured input [70,123]. Studies focusing on the internal correlates of generative activity within the brain, over external behavioral correlates, may be especially important to understand what determines the generative neural firing patterns observed in the hippocampus.

The segregation of generative and actual representations in the hippocampus also raises the question of whether the hippocampus further differentiates subtypes of generative representations. For example, are events that reflect veridical experience from the past somehow distinguished from those that reflect constructed alternatives, or those that are predictive of future choices? At the level of neural firing, it remains unclear whether or how the hippocampus might separate these possible representations. However, two points of reference in the human literature offer clues that the relevant neural substrates may be outside the hippocampus. First, patients with hippocampal amnesia can entertain thoughts that distinguish the past or the future, despite impairments in episodic memory[132,133]. Additionally, hippocampal activation during mental simulations without temporal placement versus those specifically set in

the future result in similar activation levels in the medial temporal lobe and default mode network[133,134]. These results are consistent with the idea that temporally differentiating representations related to the past or the future may not be hippocampally dependent. Second, healthy human subjects can subjectively discriminate internally and externally derived information, an ability known as reality monitoring[135]. Based on functional imaging studies in both healthy subjects and patients with schizophrenia who experience hallucinations, reality monitoring is thought to rely primarily on prefrontal cortical networks[112]. In contrast, another study reports that hippocampal activation was similar across cases of true and false recognition memory[136], further suggesting that this ability does not strictly rely on the hippocampus. Although probing reality monitoring in rodents is not straightforward, it would be notable if, for example, frontal cortical firing patterns systematically differed based on the representation of possibilities in the hippocampus that reflected veridical experience versus constructed alternatives. Such a result would be consistent with the idea that the hippocampus alone may not distinguish subcategories of generative events, but that the brain may do so via the engagement of prefrontal circuits.

Looking beyond rodents, it remains an open question as to which patterns of generative activity in the hippocampus are shared across species[137]. On the one hand, SWRs have been observed in a range of vertebrates, as have neural reactivation patterns suggestive of replay[138–144]. In humans, replay and replay-like patterns have also been reported, including activity patterns consistent with reactivating prior experience, as well as inferred sequential activity that is not simply recapitulative[145–149]. In contrast to the ubiquity of SWRs across vertebrates, the theta rhythm appears to be

more prominent and continuous in the rodent hippocampus than in various other species[137]. A notable example is the bat hippocampus, which shows network-level activity fluctuations that are not generally rhythmic yet still organize place cell firing according to phase[140,150–153]. This may suggest that actual and generative representations can be organized via temporal multiplexing even in the absence of strong rhythmicity. In nonhuman primates and humans, the hippocampal theta rhythm appears to occur in intermittent bouts and at a lower frequency[140,150–153]. Recently, theta phase coding has also been shown in single cells in human subjects[154,155]. In all, these results indicate some conservation across species of the organization of neural firing with respect to network-level hippocampal activity. More generally, they leave open the possibility that the brains of many species temporally multiplex actual versus generative internal representations.

Generativity as a Function of the Hippocampus

Recent findings described above suggest that the hippocampus regularly generates a wider range of representations than previously thought. What functional implications does this suggest?

The cognitive roles commonly ascribed to the hippocampus offer a starting point. Existing theories of hippocampal function are often based on the established role of the hippocampus in human episodic memory[30–32]. Under an essentially episodic view, hippocampal activity necessarily represents or refers to both to time and space, in accordance with the definition of an episodic experience[156] (Figure 1.3). Along these lines, hypotheses based on this view propose that hippocampal neural firing corresponds to cognitive processes such as past-oriented memory retrieval and

consolidation, or future-oriented planning and prospection[1,62]. The view that hippocampal neural firing can support memory of past episodes has been suggested by findings that causally link sharp wave-ripple events (which generally co-occur with replay) to performance on tasks requiring memory of a choice made on a previous trial[157,158], although the diversity of generative neural firing patterns reviewed above during replays suggests a more complex picture[53,62,64]. Both replay as well as theta-associated generative representations have also been posited as anticipatory processes, such as planning, in support of decision-making[159]. These ideas have been reinforced by observations reviewed above that replay and theta-associated activity can relate not only to past but also potentially upcoming behavior. To account for findings that neither of these firing patterns appears to encode animals' upcoming choice with high reliability[53,69,84,99] (but see[98,100,101]), another version of the anticipatory planning hypothesis posits that the hippocampus generates a "menu" of relevant options evaluated by other brain areas prior to the decision[42,99].

Each of these functional interpretations is plausible and conceptually important, yet only consistent with a subset of the instances and properties of generative neural firing patterns across studies reviewed above. Specifically, it is unclear how retrieval of past experience and planning for the future account for the prevalence and variety of generative representations observed, particularly those that are ambiguously related or unrelated to behavior in the past or future, and those that run counter to immediate experiences and choices[68,69,82].

As an alternative view, generative firing patterns may be understood as characteristically expressing alternatives to actual circumstances, irrespective of

whether those circumstances are in the past, present, or future. Indeed, we suggest that temporal referencing to the past, present, or future for a given firing pattern in the hippocampus may not be intrinsic or essential (Figure 1.3), a view which has also been posited to account for recent findings in human subjects[133].

By this interpretation, past- and future-oriented cognitive functions that require or involve the hippocampus would be particular applications of a broader and more essential underlying role in generativity, or representing non-actual experiences including possibilities and hypotheticals. This view of the hippocampus' role at the level of cognition is closer to imagination (Figure 1.4). A neural system implementing generativity may often construct a variety of potentially useful representations that do not necessarily relate to known circumstances, or predict immediate behavior, yet remain relevant for behavior in an indefinite horizon of time. This advantage is akin to that of insightful thoughts generated in the course of seemingly undirected mental activity, such as free reflection or mind-wandering.

If generative neural activity in the hippocampus does not intrinsically refer to actual experiences, what type of relevance might it have to experience, thought, and behavior? One unifying idea is that the hippocampus is a system for "relational" memory: a system for inferring abstract relationships between observable events (such as sensory stimuli, actions, and internal states). Integrating information in relational memory is beneficial in that it enables inference and generalization to novel circumstances, such as those where elements of previous experience are reconfigured[40,160,161]. This can be advantageous regardless of whether those novel circumstances can be anticipated at the time of generating the relational information,

consistent with the idea that generative activity can but does not always relate to immediate behavior.

Generativity could support relational thinking by bringing together elements from experience that are not actually experienced together. Specifically, the role of generative activity could be to combine otherwise separate elements of experience so that their relationship can be inferred. To illustrate, one classic example is inference of a transitive relationship; if a subject learns that A should be chosen over B and B over C from real-world experiences, then the subject can infer that A should be chosen over C, despite never having experienced A and C together. This sort of inferential ability is dependent on the hippocampus in rodents and is furthermore associated with hippocampal activation in humans[45,162–164]. A second example is inferring spatial relationships, which often relies on the ability to link physically discontiguous prior experiences. Rodent behavior has long indicated the ability to infer novel routes, including more efficient shortcuts, through a spatial environment[38,39]. At the level of neural firing, replay events can stitch together into one coherent representation two track segments that the rat has never traversed in a single run, and can represent novel paths to goals that have not been taken before[82,104]. Further, the mouse hippocampus has been shown to coactivate neurons related to representation of distinct events during SWRs in an inferential reasoning task[164]. Generative representations such as these appear to combine information across separate prior episodes into internally constructed possibilities that may have the potential, but are not required, to inform behavior. Consistent with these findings, not only is human hippocampal activation associated with correct inferential choices[164], human subjects also exhibit internally generated sequences of hippocampal

activity that reorder elements of experience into novel, inferred sequences that do not simply recapitulate previously experienced sequences[147]. Taken together, these results suggest that generative hippocampal activity may be well suited to contribute to relational thinking, and ultimately the internal generation of new knowledge that goes beyond actual experience.

Importantly, inferences in such a relational memory system can operate not only across various modalities, such as sensory, motor, and internal states, but also generate all kinds of relations[160,165–167]. Under this view of the hippocampus, temporal and spatial relations are instances of relations which are rich and prevalent – and experimentally accessible – yet not fully comprehensive. This is evidenced by the requirement of the hippocampus for animals to infer relations that are neither temporal nor spatial[45,163,164,168].

Understanding the hippocampus as intrinsically representing alternatives to actuality suggests that the hippocampus may have a broader role in cognition than is often described. If the role of the hippocampus in cognition is not restricted to particular types of relations such as in time and space, generative neural activity might involve the construction and application of any number of relations across additional domains. In rodents for instance, neural firing in the hippocampus is often studied in relation to space yet can also encode a wide variety of variables from experience such as odors and sounds[100,169–174]. Accordingly, we would expect the hippocampus to exhibit generative activity corresponding to alternatives to actual experience in terms of such variables. This could include linking aspects of experience across modalities into internally constructed representations of possibilities or hypotheticals that have not

actually been experienced. In the case of humans, it is particularly notable that the hippocampus has long been linked not only to the acquisition of episodic memory, but of declarative memory more generally, which entails acquisition of semantic memory. By this token, it may be plausible that generative neural activity in humans (in addition to animals) can represent alternatives to actuality by engaging in semantic relations – for instance, in language comprehension or production, and in creativity understood more broadly[2,34,133,175–177].

This broader view of generativity in the hippocampus may have additional implications at a higher level than representation. An advantage of internal models is that they enable internally directed exploration, or generative simulations and hypothesis formation intended to yield maximum information gain[159]. Interestingly, though such exploration is recognized to be ultimately adaptive, it might have little or no immediate utility, and, further, neural activity implementing this process could be uncorrelated with immediately upcoming behavior. Exploration can also be driven by curiosity, an intrinsic motivation that has notably been linked to the construction of rich internal models[178]. These are several points of contact between information-based exploration and generative activity patterns. Yet even beyond information-based exploration, it is increasingly recognized that humans and a range of animals can harbor intrinsic motivations expressed as self-determined and self-guided goals, manifesting in behavior as “play”[179]. Critically, like exploration, play has (practically by definition) little or no immediate utility to subjects, though its relevance or role in advanced cognition is potentially crucial[179]. It is therefore worth speculating that analogous play-like adoption of seemingly arbitrary internal aims and constraints is relevant to understanding

generative activity patterns in the hippocampus, both in animals and humans. Thus, generative hippocampal representations and the hippocampus at large may represent alternatives to actual experience not only to navigate immediately relevant environments and objectives, but also to pursue any number of internally directed and invented goals. Ultimately, doing so may be crucial not only to evolve a greater understanding of past and immediately relevant experience, but also to deal adaptively and flexibly with unexpected scenarios and unknown circumstances in the future.

Conclusion

Imagination requires the ability to generate experience, thoughts, or representations that do not refer to the actual present. This essential ability, termed “generativity,” can be understood at the level of the brain and need not entail conscious awareness or mental imagery. Human studies have linked imagination and the construction of hypotheticals to the hippocampus and complementary studies in the rodent hippocampus have identified neural firing patterns corresponding to experiences that do not reflect the actual present. Traditional accounts of hippocampal function often interpret these generative firing patterns, such as those observed during SWR replays and late phases of the theta rhythm, in relation to actual experience in the past and future. This important view may be limited in accounting for the diversity of generative hippocampal firing patterns suggested by recent findings. Rather, we propose that representing alternatives to actual present experience is itself essential to the hippocampus. These representations may span a wide range of self-generated possibilities, hypotheticals, and non-actualities of all kinds: from past episodes and

anticipated futures, to counterfactuals, alternative presents, novel combinations of experiences, and to creative or even playful simulations, including those without spatial or temporal reference. We further suggest that diverse generative hippocampal activity patterns may be used to learn, infer, and consider various abstract relations. This in turn would suggest that functions of the hippocampus that refer to time or space (such as episodic memory and mental time travel) may be particular applications of a broader system of imagination[7,15,18,180] (Figure 1.4). Notably, this view advocates that the function of generative neural activity in the hippocampus may not be characterized by the strength of its correlation to immediate behavioral choices, but rather by its relationships to internal processes[69,116,124,181]. The contribution of generativity to behavior may not be immediate, and in fact might have an indefinite horizon in the lifetime of subject. This system may be elaborated in humans, supporting frequent and at times seemingly undirected flights of creativity and imagination.

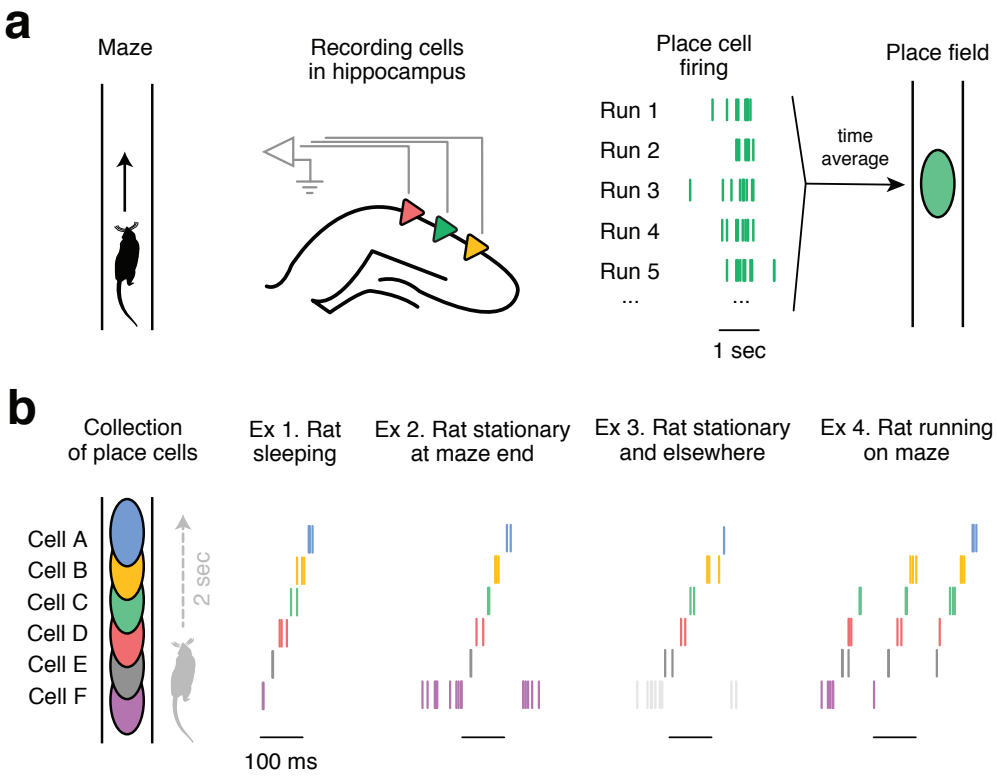

Figure 1.1 Hippocampal neural firing patterns
(a) Place cell firing. Left: a rat runs on a linear maze. Middle: the firing activity of multiple neurons is recorded simultaneously from the hippocampus while the rat runs. Right: one example recorded place cell fires in the same spatial location on the track over many runs, but with notable variability across individual runs (firing denoted by raster lines). The time averaged firing of the example place cell over many runs forms the cell's place field (oval). **(b)** Generative place cell firing examples. Left: place field locations (ovals) of multiple place cells (A-F) that fire as the rat runs along the linear maze (one run takes ~2 sec.). Right: examples (Ex 1-4) of generative firing patterns that occur when the rat is not actually in the cells' place field locations. Place cell firing is denoted by cell-colored raster lines. Ex 1: replay occurs while sleeping in an environment separate from the maze. Ex 2: replay occurs while rat is stationary at the lower maze end. Ex 3: replay occurs while rat is stationary and in a different maze environment (light grey lines indicate firing of a cell active in the different maze, but not active on the maze shown at Left). Ex 4: occurs while the rat is running on the maze, during the theta rhythm. In Ex 1-4, place cell firing corresponds to a series of locations not presently occupied by the rat.

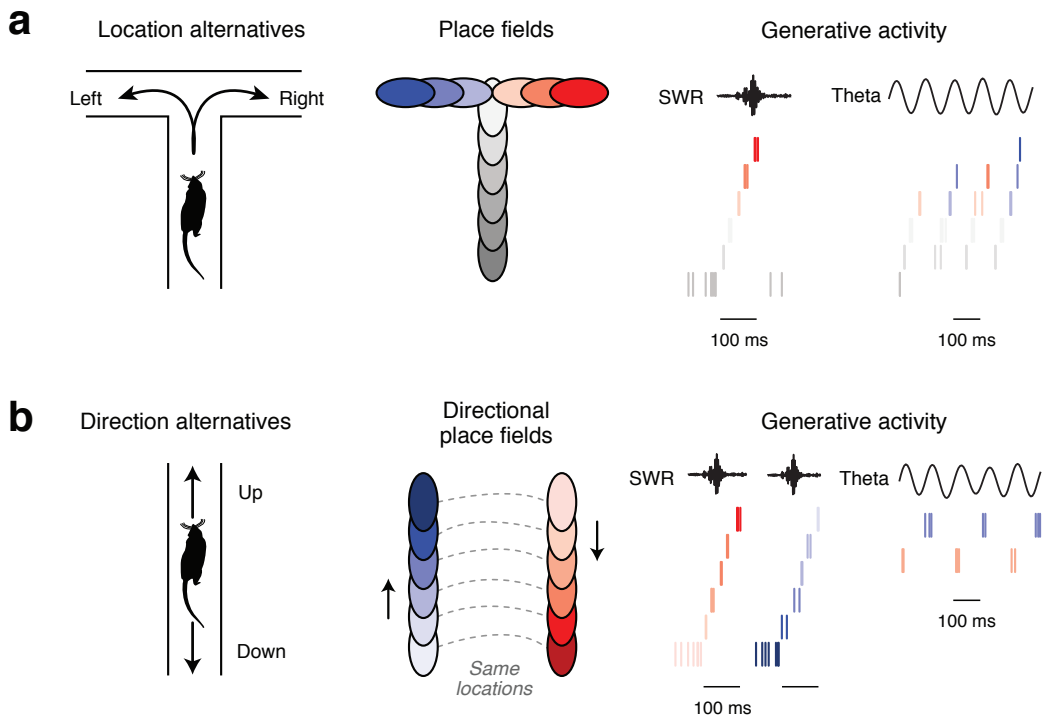

Figure 1.2 Generative activity corresponding to alternative possibilities
(a) Example of neural firing corresponding to alternative locations. Left: a rat running on the central arm of a bifurcating maze can either turn left or right. Middle: individual place cells fire when the animal is in particular place field locations on the maze (ovals colored by cell). Right: generative activity. While the rat is resting at the end of the central maze arm, generative neural firing during a sharp wave-ripple (SWR) can replay other locations (raster lines colored by cell). At other times while the rat is running up the central maze arm towards the bifurcation, neural firing during the theta rhythm alternates between current and upcoming locations within each cycle, and between the left and right trajectories ahead across the second halves of theta cycles. **(b)** Example of neural firing corresponding to alternative directions. Left: a rat on a linear maze can run in up or down directions. Middle: different sets of direction-selective place cells fire in their place field locations (ovals) when the rat runs up or down (ovals colored by cell). Right: generative activity. Neural firing during SWRs replays locations that do not correspond to the rat's actual location at the end of the maze. Additionally, while the rat is running in one direction, cells corresponding to the rat's actual and alternative directions fire in alternation along the theta rhythm.

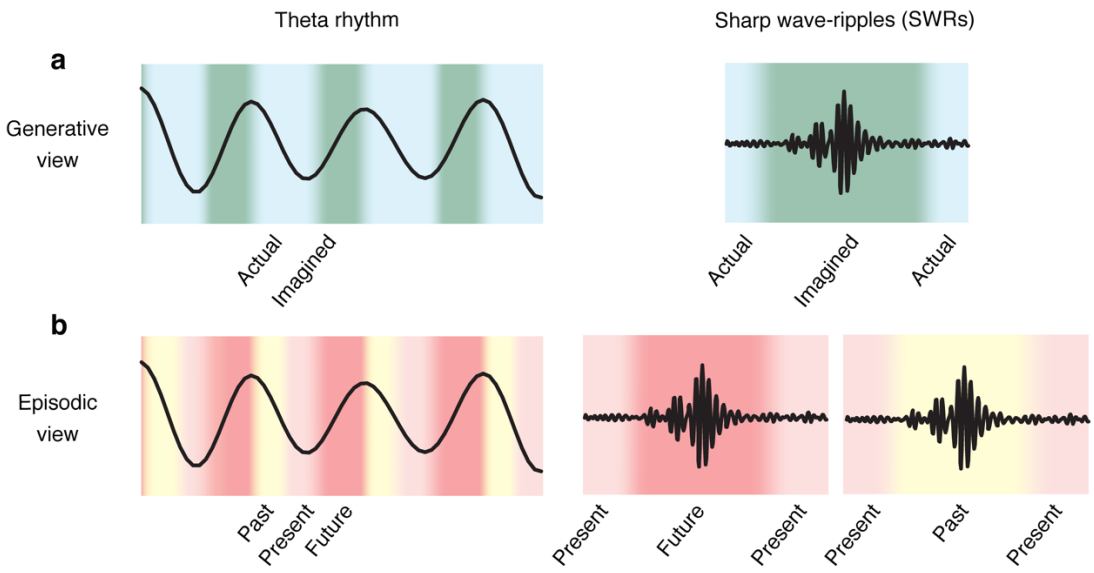

Figure 1.3 Two views of hippocampal firing during network-level activity patterns **(a)** Generative view (proposed). Neural firing during early theta phases represents actual circumstances, while representations during late theta phases and sharp wave-ripple (SWR) replays are imagined (generative). **(b)** Episodic view (traditional). Neural firing during each theta cycle corresponds to sequential locations in the past, present, and future, and neural firing during SWR replays correspond to past or future locations. Note that the Episodic view can be understood as a particular application of the proposed Generative view.

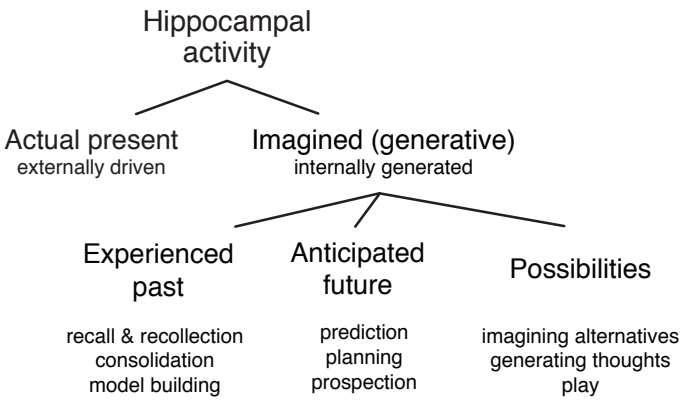

Figure 1.4 Schema for interpreting hippocampal activity
We suggest that generative activity arising from the hippocampus may not only correspond to the experienced past or anticipated future, but also to a wide range of possibilities. This view may also organize or suggest various cognitive functions.

CHAPTER 2: HIPPOCAMPAL REPRESENTATIONS OF ALTERNATIVE POSSIBILITIES ARE FLEXIBLY GENERATED TO MEET COGNITIVE DEMANDS

Summary

The cognitive ability to go beyond the present to consider alternative possibilities, including potential futures and counterfactual pasts, can support adaptive decision making. Complex and changing real-world environments, however, have many possible alternatives. Whether and how the brain can select among them to represent alternatives that meet current cognitive needs remains unknown. We therefore examined neural representations of alternative spatial locations in the rat hippocampus during navigation in a complex patch foraging environment with changing reward probabilities. We found representations of multiple alternatives along paths ahead and behind the animal, including in distant alternative patches. Critically, these representations were modulated in distinct patterns across successive trials: alternative paths were represented proportionate to their evolving relative value and predicted subsequent decisions, whereas distant alternatives were prevalent during value updating. These results demonstrate that the brain modulates the generation of alternative possibilities in patterns that meet changing cognitive needs for adaptive behavior.

Introduction

Animals are continually faced with decisions about what to do and where to go next. In the context of behavioral tasks, a long tradition of animal experiments and behavioral models suggest that the brain can make adaptive decisions by comparing the expected values of available options, which are learned through experience. As an example, in spatial settings, an animal may compare the expected values of different rewarded locations and associated paths. After making a choice, a rewarding outcome would lead to an update that increases the stored value of the rewarded location and the path taken to get there, enabling subsequent adaptive decisions even as outcomes change[53,65,182,183].

While the idea of retrieving and updating expected values of possible options seems relatively simple, in real-world situations like navigation, choices often lead to outcomes that are distant in space or time. This poses a challenge: the brain must go beyond current experience to decide among or learn about alternative “non-local” possibilities. For instance, when making a choice among nearby routes, the brain may retrieve values related to their ultimate destinations. Further, in complex structured scenarios, rewards received in one location can imply information about the availability of rewards in other places, as in zero-sum situations or multi-step planning tasks. In such cases, the brain may take advantage of learned structure to make inferences across space, and update not only the value of the current location but also the values of alternative locations and paths[46,74,184,185]. Additionally, unlike many laboratory tasks, naturalistic scenarios can have many alternatives available at once[159,186–190]. This indicates a need to prioritize[191]; when deciding among or updating different possibilities,

some judicious mechanism is required to consider the most relevant non-local alternatives to meet current demands.

The neural mechanisms that enable such prioritized computations about relevant non-local possibilities during complex behavior are not understood. Existing data indicate the hippocampus and its representations of space as a starting point. First, the hippocampus is critical for rapid learning and performance of spatial tasks where animals must learn the locations of and routes among rewarded locations[192–199]. Second, the hippocampus is well known for spatially tuned “place cells” whose activity typically signals the actual location of the animal[47]. These cells are often described as the substrate for a cognitive map, or an internal model of the world, that encodes the relationships among both locations and experiences more broadly[43,49]. Critically, while place cells are best known for coding an animal’s actual position, they are also capable of expressing non-local representations of alternative locations at a sub-second timescale[54,55,74,80,200–203]. Finally, natural behavior often involves experience-guided decision making during active navigation, and hippocampal non-local representations can be regularly expressed during movement[69,204]. Thus, during active behavior, generating a non-local representation corresponding to a particular alternative location could serve to retrieve or update information associated with that location, including its value.

Non-local representations have been linked to cognitive processes for both decision making and learning during navigation. These population-level representations are most often associated with the sequential firing of place cells corresponding to a trajectory through locations behind, at, and ahead of the animal’s actual

location[55,93,200,95,181]. In the context of decision making, as animals approach a choice point, these representations can sweep along future paths ahead, which is evocative of a role in retrieving at least immediately upcoming options[69,98–101,117,159,205,206]. These sequences also engage place cell activity on timescales consistent with synaptic plasticity, suggesting a role in learning[55,200,207–210], and potentially updating internal representations based on experience.

Yet, whether or how the brain generates representations of different alternatives as cognitive demands change throughout experience-guided decision making and learning remains unclear. We therefore combined approaches typically used to separately study decision making and reinforcement learning, or experience-guided navigation. We developed a dynamic patch-foraging task where changing reward probabilities across six locations challenged animals to continually update their internal models to make experience-guided choices about where to go for reward. By leveraging a computational model, we estimated internal cognitive variables from animal behavior related to both value-guided decision making and value updating. As these cognitive variables evolved across successive trials of experience, we monitored hippocampal neural activity to identify non-local representations expressed during active navigation. We observed a range of representations of alternatives, corresponding to potential paths not only ahead of, but also behind the animal, including along counterfactuals, and to distant locations in remote foraging patches. We further found that these representations were modulated across successive trials in two distinct patterns, one for representational content and the other for spatial extent, each of which was related to distinctly evolving cognitive variables. These findings demonstrate that mechanisms

exist that regulate the expression of distinct non-local possibilities in conjunction with cognitive needs.

Results

Rats Make Experience-Guided Decisions in the Spatial Bandit Task

We developed a dynamic foraging task where performance could benefit from representing alternative possibilities, both for deliberating among alternatives and for updating information about alternatives. This “spatial bandit” task combined features of spatial memory and decision-making paradigms (Figure 2.1). First, as in classic spatial memory behaviors, rats (n=5) navigated a maze based on prior experience to reach reward locations. The track was made up of three Y shaped “foraging patches” radiating from the center of the maze, and each patch contained two reward ports at the ends of the linear segments (Figure 2.1). The multiple bifurcations and reward locations provided many opportunities for deliberation among alternative options and reward-based updating. Second, as in classic decision-making tasks, we introduced uncertainty by dispensing rewards probabilistically. Each port was assigned a nominal probability of reward, p(R), of 0.2, 0.5, or 0.8 (Figure 2.1), and one of the three patches had a greater average p(R) than the other two. On any visit to a port the rats either did or did not receive reward, determined by the p(R), and consecutive visits to the same port were never rewarded. Thus, animals could only accurately infer each port’s hidden reward probability state through experiences across multiple visits and multiple port locations.

During run sessions, reward probabilities around the track covertly changed in blocks (as in Figure 2.1) every 60 (n=4 animals) or 80 (n=1 animal) trials. Each day of

task experience consisted of up to 8 run sessions, separated by rest sessions in a rest box, and each run session consisted of 180 (n=4 animals) or 160 (n=1 animal) trials. A trial was defined as the period between a departure from one port to a departure at another port, and animals individually completed 5760, 8970, 10191, 7917, and 5220 trials, providing a large dataset compatible with behavioral and neural analyses. Importantly, the reward contingency block transitions were uncued and changed which patch was associated with the highest overall nominal probability of reward. These reward probability transitions thereby encouraged experience-based, adaptive decision making.

Animals exhibited choice behavior akin to patch foraging[211] (Figure 2.1). They began by serially exploring the ports in the three patches, often making repeated choices to alternate between ports within a patch (Stay trials) and occasionally making flexible choices to navigate to an alternative patch (Switch trials) (Figure 2.1, Figure 2.6). Animals adapted their choices based on their dynamic reward experiences across reward ports, trials, and contingencies (as in Figure 2.1). As a result, animals generally learned and remained within the patch with the highest nominal reward probabilities by the end of each block (Figure 2.1).

The patterns of these Stay and Switch choices indicated a decision strategy that used reward history information, including reward history related to both the current patch and to alternative patches. To understand how reward information related to animals' choices at a single-trial level, we fit a behavioral learning model to each animal's sequence of port choices and reward outcomes (see Methods). We estimated weights related to the expected values of the animal's two potential options on each

trial: Staying within the current patch or Switching between patches. Logistic regressions predicting Stay or Switch choices revealed significant effects of not only the Stay value (measured as the value of the upcoming port in the current patch) but also the Switch value (measured as the value of the more valuable unoccupied patch) in each animal (Figure 2.1). The coefficients were negative for Stay values, indicating that animals were less likely to make Switch choices as the value of Staying increased (Figure 2.1). In contrast, the coefficients were positive for Switch values, indicating that animals were more likely to choose to Switch as the value of Switching increased (Figure 2.1). Thus, animals' behavior depended on previous reward experiences both from nearby Stay locations and more spatially distant Switch locations across the maze (Figure 2.1, Figure 2.10).

Since Switch choices were value-guided and punctuated often longer, stable bouts of Stay choices (Figure 2.1, Figure 2.6), we anchored further analyses around these self-paced patch changes. Here, a central decision variable is the relative value between these two options on each trial: Switch value - Stay value (Figure 2.1). This relative value was low during trials far from a Switch and on average increased over the course of the Stay trials leading up to a Switch choice. On the Switch trial, relative value was highest, and then decreased again across trials after the Switch (Figure 2.1). Note that these relative values do not incorporate the bias to Stay (or cost to Switch) as illustrated by the negative constants in Figure 2.1. We also note the relative value fluctuations on alternating trials before a Switch, which are consistent with animals tending to Switch after visiting the higher value port within a patch (Figure 2.1). Importantly, the increasing and decreasing relative value pattern around Switches

reflects a gradual updating of values over successive trial outcomes. This suggests that animals may similarly access internal estimates of values associated with Switch and Stay options when deciding whether to leave the current patch and, after Switching, whether to remain in the new patch or again navigate elsewhere.

Taken together, these findings provide evidence that animals learn from their changing reward experiences and leverage this experience to make value-guided decisions among alternative paths. Furthermore, they suggest that seemingly isolated and behaviorally overt Switch choices occur in the context of gradually changing covert reward expectancies that place evolving cognitive demands on the animal across successive trials.

Representations of Alternative Paths Ahead of the Animal

Given these behavioral results, we then asked whether the content of hippocampal representations of alternative paths was also modulated around Switch choices. As previous work has identified non-local representations consistent with possible future locations[69,98,99,101,117], we began by examining non-local representations extending ahead of the animal's actual position that occurred while the animal was in the first segment of each trial. In this period, the animal approached the first choice point, which is associated with the choice to Stay or Switch. We used an established state space decoding algorithm[212,213] (2 cm spatial bins, 2 ms temporal bins, see Methods) to assess the instantaneous hippocampal representation of space[59,60,214] during navigation (animal speed > 10 cm/s) across all five rats, each implanted with tetrode microdrives targeting the CA1 region of the hippocampus (Figure 2.6). We limited our analyses of

non-local representations to periods with high confidence that the decoded hippocampal representation was in a non-local track segment, defined as a segment distinct from the one corresponding to the rat's actual location (Figure 2.2, see Methods).

We observed non-local representations that reflected locations along either of the two paths ahead of the choice point (Figure 2.2). We also verified the expected organization of these non-local representations of paths ahead based on the phase of the hippocampal theta rhythm[55,89,91,93,200] (Figure 2.7). On each trial, we then classified non-local representations as reflecting paths that, if physically traversed, would either lead the animal to Stay in the patch or to Switch between patches (Figure 2.2). We began by examining the Stay trials leading up to and following Switch choices. Critically, this enabled us to assess the content of non-local representations as the relative value of Stay and Switch paths changed (Figure 2.1), while the behavioral choice to Stay remained constant.

Strikingly, the relative representation of Switch and Stay paths mirrored their evolving relative value. Across trials preceding a patch Switch, non-local representations of Switch paths ahead became increasingly prevalent relative to Stay paths (Figure 2.2). All animals exhibited this pattern, showing an approximately 1.3-1.5 fold increase in the likelihood of non-local representations consistent with the unchosen Switch versus the chosen Stay paths over the 20 trials before a Switch choice (Figure 2.2, Figure 2.7). This culminated in a ratio approaching equal (0.5) representation of each path ahead on the trial before a Switch trial (Figure 2.2, Figure 2.7). Additionally, on Stay trials after the animal arrived in a different patch following a Switch, non-local representations were again enriched for the unchosen Switch path. Then, the longer the

animal chose to Stay, the less the non-local representations reflected the Switch path (Figure 2.2, Figure 2.7). These increases and decreases in the proportion of Switch path representations were driven largely by changes in the amount of time spent representing the Switch path, while the level of Stay path representation remained, on average, relatively stable across successive Stay trials (Figure 2.8). We also confirmed that the roughly symmetrical increasing and decreasing pattern around a Switch was not driven by periods when the animal and decoded non-local spatial representation were very close by, as is possible near choice points, by requiring representations to extend at least 10 cm from the animal (Figure 2.8). Additionally, this pattern was not driven only by short Stay bouts, but was also seen when analyses were restricted to long bouts (Figure 2.8).

Notably, the modulation of non-local representations of paths ahead of the animal before and after Switch trials resembles the changes in relative value between the Switch and Stay options (Figure 2.1). That is, trials with a higher relative value of Switching were associated with greater relative representation of Switch paths. This same relationship could also explain differences in non-local representations across Stay and Switch trials. Stay trials were on average associated with approximately 20-30% non-local representation of paths consistent with Switching, whereas the complementary bias was seen on Switch trials, in which 70-80% of the alternative representations were of the Switch path (Figure 2.2).

These findings indicate that the hippocampus continually tunes the retrieval of spatial alternatives in a pattern related to their evolving relative value. In the context of experience-guided decision making, our results are consistent with an across-trial

process where internal sampling of alternatives is biased by relative expected value, such that as these estimated values become more similar—and the cognitive demand for distinguishing between them potentially greater—the options are sampled more equally. Recruiting representations of relevant paths based on their viability for making the best choice, in turn, could enable more accurate comparisons of the values[215].

The possibility that non-local representations are engaged across trials in an internal sampling process led us to ask whether the proportion of non-local representation of the Switch path was predictive of the future choice to Stay or Switch. We reasoned that proportionally more representation of the Switch option on the Stay trial an entire trial in advance of the Switch could provide samples of the relatively high value Switch option (Figure 2.1) that in turn could influence the decision on the next trial. To investigate this at the level of single trials, we used a cross-validated logistic regression to predict whether an animal would choose to Stay or Switch on each trial based on the proportion of non-local representation corresponding to Switch paths as the animal approached the first choice point of each trial (using the same metric as in Figure 2.2). Training data were balanced such that chance level was 0.5.

We found evidence that non-local representations could be used in an across-trial decision process. Non-local representations occurring on the previous trial predicted the subsequent trial's Stay or Switch choice better than chance. This indicates that the increased alternative representation on the Stay trial before a Switch trial was predictive of a future Switch choice an entire trial in advance, even when the immediately upcoming choice on the current trial was still to Stay (Figure 2.2). This choice prediction was even more accurate when considering only the non-local

representations that occurred during the approach of the choice point on the current trial (Figure 2.2), which is expected from Figure 2.2. Strikingly, including non-local representations on both the previous trial and the current trial led to an even better prediction of the choice the animal would make on the current trial (Figure 2.2). While Figure 2.2 indicate a ramping process on average across bouts of Stay trials before patch Switches, this regression result extends the observation of an across-trial modulation to the level of individual choices within bouts.

Representations of Alternative Paths Behind the Animal

The idea that representations of alternatives may be flexibly generated to meet decision-making needs across trials, rather than only for the immediate future choice, led us to ask next whether non-local representations of alternatives could be expressed at times when there was no immediately upcoming choice point. Here we focused on periods when the animal was traversing the final track segment of a trial and approaching a reward port. During this period, we found that non-local representations not only corresponded to the path the animal recently traversed, but also the unchosen Stay or Switch path that the animal could have come from or gone to but did not. This is consistent with the representation of counterfactuals (Figure 2.3).

These non-local representations of Stay or Switch paths behind the animal (Figure 2.3) showed a very similar pattern of modulation around Switch choices as did non-local representations of paths ahead of the animal (Figure 2.2). Even though the representations behind were expressed after the animal had behaviorally indicated a choice on the current trial (Figure 2.3), the relative representation of the counterfactual

Switch path ramped up across Stay trials before a Switch and ramped down across the subsequent Stay trials after arriving in a different patch (Figure 2.3, Figure 2.7). Again, this increasing and decreasing pattern was roughly symmetric on average around the Switch trial. Furthermore, just as for representations ahead (Figure 2.2), representations behind the animal were biased on average to represent locations along the actual path taken on the current Stay or Switch trial (Figure 2.3). Thus, the non-local representations behind the animal were systematically modulated with the evolving relative values of the options (Figure 2.1).

As with non-local representations extending ahead of the animal, the dynamic generation of representations of alternatives behind the animal was primarily driven by changing levels of Switch path representation (Figure 2.8). This is consistent with an increasing relative representation of the unchosen alternative path leading up to and following a choice to Switch. These results were also consistent when analyses were limited to representations extending at least 10 cm behind the animal (Figure 2.8), as well as in bouts within a patch lasting at least 10 Stay trials (Figure 2.8). Further, these non-local representations behind were concentrated in early phases of the theta rhythm, as expected from previous work[55,91,93,200] (Figure 2.7).

Representations behind the animal could also predict future behavior, consistent with an across-trial decision process. The relative representation of non-local paths behind the animal on the previous trial were predictive of whether the animal chose to Stay or Switch on the entirely subsequent trial (Figure 2.3). These predictions were, as expected (Figure 2.3), not as accurate as those from non-local representations occurring as the animal traversed the final segments of trials, after the animal had

already behaviorally expressed a choice to Stay or Switch (Figure 2.3). Nonetheless, previous trial non-local representations predicted subsequent choice well above chance. Thus, the proportion of non-local representation of the Switch path behind the animal on a given trial could predict what the animal would do several seconds later following traversal down the track segment to the reward port and back up the track segment to the choice point.

Together with the results from representations ahead of the animal (Figure 2.2), these findings indicate that animals can progressively engage non-local representations associated with a progressively more valuable unchosen alternative (here Switching) both before and after the choice point leading to that option on each trial. These findings provide further support that the hippocampus dynamically samples alternative options, and can do so across successive trials with varying cognitive demands throughout flexible decision making.

Representations of Remote Alternatives

Beyond retrieving information from an internal model related to an ongoing decision-making process, non-local information could be useful for additional purposes, including relating experiences or making inferences across space. In this task, where reward ports are distributed far across the maze, representing distant locations[70,95,98,108,109], including those in alternative patches, could support relating reward experiences across space or updating values of distant locations based on local reward experiences. We therefore asked whether there was evidence for representations of distant locations across the maze, and whether these representations were specifically generated in relation to patch Switching, as animals learned from dynamic reward experiences.

In addition to non-local representations corresponding to track segments neighboring the animal's actual location (Figure 2.2, Figure 2.3, Figure 2.4), we also observed non-local representations corresponding to alternative patches, closer to other reward ports (Figure 2.4). Alternative patches were represented on 10-20% of trials with any non-local representation (Figure 2.9). To quantify the extent of non-local representations, we then measured the maximal distance in centimeters between the animal's actual position and the most likely decoded position represented by the hippocampus on each trial, during running at least 10 cm/s (Figure 2.4). Given our previous observation of representations of paths both ahead (Figure 2.2) and behind (Figure 2.3), we investigated non-local distance both as animals approached the first choice point and after the animal passed the final choice point in each trial (Figure 2.4).

We found that non-local distance ahead was strongly modulated (Figure 2.4) in a pattern surprisingly distinct from the modulation of non-local content (Figure 2.2). The maximum non-local distance gradually ramped up across Stay trials and the first segment of Switch trials (Figure 2.4). The extent was then greatly elevated once the animal arrived in a new patch, and then decreased rapidly across the subsequent few Stay trials (Figure 2.4). This is in marked contrast to the modulation of non-local content, which increased and decreased in a roughly symmetrical manner. The maximum non-local distances could reach an average of approximately 90 cm from the animal (Figure 2.4, Figure 2.9), and, across animals, increased by approximately 40 cm from baseline (Figure 2.4) on the Stay trial following a Switch. This asymmetric pattern of non-local extent around Switch trials suggests a second kind of modulation of the generation of

non-local representations across successive trials by the hippocampus during experience-guided decision making.

We then examined how these distances changed throughout the Switch trials themselves (Figure 2.4). Here we considered maximally distant representations both ahead and behind the animal during traversal of each of the four track segments of Switch trials. We found that while on average non-local representations extended as far as 40-60 cm from the animal on the first segment, this distance dropped to approximately 20 cm on both the second and third segment. Then, interestingly, the maximum distances rose again to approximately 40-80 cm as the animal traversed the final track segment of the trial and approached the reward port in the new patch.

The extent of non-local representations behind the animal also showed an asymmetric pattern of across-trial modulation around Switches (Figure 2.4, Figure 2.9). That is, the maximum distance of non-local representations behind the animal on the final segment of trials was also elevated upon Switching patches, and then decreased back to baseline over subsequent Stay trials (Figure 2.4). Thus, the extent both ahead and behind showed a modulation pattern around Switch trials that was distinct from the modulation of Switch versus Stay path content (Figure 2.2, 2.3).

We next asked whether non-local representations that corresponded to alternative, unoccupied patches were predictive of past or future patch choices. We found no evidence for such a relationship. Representations of specific alternative patches were neither biased to represent the patch the animal just arrived from nor predictive of the patch the animal would visit next (Figure 2.9). And, while in some animals these alternative patch representations did tend to overrepresent the track

segments containing reward ports, rather than the central three track segments, this effect was not consistently observed across animals (Figure 2.9). These results indicate that representations of specific distant spatial alternatives in other patches are not directly reflective of specific prior or upcoming Switch choices.

Together, these patterns indicate that non-local extent is modulated above and beyond distance to a goal[98] or choice point[95]. These findings also raise the possibility that more distant representations are involved in computations particularly required upon patch switching, including the initial approach of the first reward port in the new patch. This is a period in which the animal has the opportunity to learn about the reward values in the new patch, and potentially how they relate to values across the maze. Importantly, the reduced maximum distances observed in the third segment of Switch trials (Figure 2.4) indicates that the relative novelty or time since last visit for each track segment cannot be the main driver of the elevated distances observed after a patch Switch (Figure 2.4). This is because the third segment has a very similar level of relative novelty to the final segment, but the expression of non-local representations was very different.

Non-local Distance is Enhanced During Learning Opportunities

The observation that the extent of non-local representations was greatly elevated on the final segment of a Switch trial and for a few trials thereafter, after entering the new patch, suggested that representing distant locations might be particularly useful at those times. Specifically, we hypothesized that these representations might support computations that enable inference and updating related to alternative options across

the environment. We then returned to our computational model of choice behavior, which suggested why this might be the case.

This model is an abstract statistical learning model that captures what a rational observer would expect about the probability of reward at each port given the animal's actual history of reward experiences, which themselves occurred across different ports. Additionally, we fit the parameters of a choice model, similar to logistic regression, to estimate how this value information was used in relation to choice behavior. These analyses provided the opportunity to ask when information about the values of distant ports around the maze are most likely to be updated in relation to behavior at the level of trials.

Our model choice was guided by the blocked reward contingency structure in this task, in which high reward probabilities tend to be mutually exclusive across patches at any given time. For example, the presence of a highly rewarding patch indicates that the other patches are less rewarding, and vice versa. An animal could account for this structure by estimating value jointly across ports rather than separately. Accordingly, the model estimates the expected value of each port on each trial by inferring the underlying contingency state across the entire maze. Specifically, the learning model is a hidden Markov model (HMM) whose hidden states are the sets of reward probabilities across all six ports (Figure 2.5, left and upper right, see Methods). This "global" model enables non-local value updating, where a reward outcome at one port can impact the expected values of other ports across the maze by providing evidence for specific contingency states.

We first asked whether this non-local updating was important to model animal behavior in this task. If so, then a model that included this feature should fit the data better than an otherwise closely matched model that does not allow for non-local updating. We therefore implemented a second “local” model that inferred the value state of each port independently, as separate HMMs (Figure 2.5, lower right) that were otherwise identical to the initial “global” model (see Methods).

The global model that allowed for non-local inference across ports significantly improved the fit to behavior in all animals (Figure 2.5). This suggests that animals did not learn port values independently, and instead adopted a non-local learning strategy in which outcomes at one port informed reward estimates at other ports, across the maze.

We then asked whether the dynamics of variables related to value updating in the model corresponded to the dynamics of non-local representation distance in the neural data (Figure 2.5, Figure 2.4). Importantly, in this model, the degree of updating is not required to be uniform for every trial’s experienced reward outcome. Instead, as is typical in Bayesian statistical learning, an outcome drives greater learning when its reward probability is more uncertain. We therefore quantified both outcome uncertainty as well as global value updating using the model (Figure 2.5). We first assessed outcome uncertainty, defined as the entropy over the value state of the upcoming port.

As with non-local extent, this outcome uncertainty was, on average, relatively stable in trials before animals Switched, increased substantially on Switch trials, and then decayed rapidly back to baseline (Figure 2.5, Figure 2.10). Higher values upon Switching patches reflect the lack of recent information about likely outcomes at the

locations in the new patch, which, in turn, could drive stronger value updating. The asymmetric pattern of outcome uncertainty around Switch trials is similar to that of the hippocampal extent results (Figure 2.5), potentially reflecting preferential updating or propagation of value information to internally represented distant non-local places during these uncertain periods.

We next aimed to specifically capture this global value updating process using the model. To do so, we quantified the total change in the estimated values across all ports, including both the local (occupied) and non-local (unoccupied) patches in the maze from trial to trial (Figure 2.5, Figure 2.10). This revealed that value updating at locations across the maze was also asymmetric around patch Switch trials, with higher degrees of updating on and after the Switch, and a rapid decay thereafter (Figure 2.5), again akin to the pattern of hippocampal extent (Figure 2.5). Taken together, these results are consistent with the idea that the hippocampus preferentially samples more distant locations across the environment when expected values are uncertain, and new reward outcome information must be integrated into the animal's internal model across the environment.

Discussion

Our findings demonstrate the distinct modulation of the content and spatial extent of non-local representations of alternatives in the hippocampus, in patterns that are well-matched to the animals' evolving cognitive needs throughout flexible decision making. The task structure encouraged animals to continually update their estimates of the values of different options based on experience and to use those values to decide

whether to Stay in the current patch or Switch to a different patch. Moment-by-moment decoding of hippocampal spatial representations during navigation, in combination with behavioral modeling, revealed two distinct patterns of modulation: (1) non-local representations of alternatives were expressed in relation to their relative expected values, with the higher value alternative making up the greater fraction, and (2) non-local representations extended further from the animal at times where our model suggested that non-local learning was most prevalent. These patterns of modulation spanned successive trials and engaged non-local representations of locations not only ahead of, but also behind the animal. Further, the representations included locations that were not only close to, but also very distant from the animal. These findings reveal the broad representational capacity of the hippocampus during locomotion and indicate that this representational capacity is specifically engaged in patterns appropriate to subserve cognitive demands for both making decisions among and updating internal representations of alternatives.

Our behavioral model provided a framework for understanding these contrasting patterns between non-local content corresponding to alternatives (Figure 2.2, 2.3, which are approximately symmetric around Switch trials) and non-local extent (Figure 2.4, which is asymmetric). In the trials leading up to a patch Switch, the hippocampus tended to reflect locations at distances within the neighboring track segments (Figure 2.4), along either the Stay or Switch path (Figure 2.2, 2.3), rather than at the remote port locations where reward was previously obtained. This is consistent with a key idea of many reinforcement learning models—that choices are guided by a learned representation of the long-term reward consequences of some local action. In this

framework, an animal would evaluate Stay or Switch options leading up to a Switch choice by retrieving the up-to-date values associated with the neighboring paths nearby rather than their ultimate remote destinations. The observation that non-local representations extend further during greater non-local value updating, particularly after Switching, is further consistent with the idea that non-local updating of locations across the maze could help to propagate value information from those distal reward port targets to proximal choice points. Thus, whereas the elevated extent of non-local representations after Switches may reflect a non-local updating process (Figure 2.5), the representation of alternative Stay and Switch paths may reflect judicious retrieval of nearby segments when they are relevant to an ongoing choice between alternatives (Figure 2.1).

Modulation of Non-Local Representations With Relative Choice Value

Previous work on non-local representations and cognitive functions during locomotion has often focused on representations ahead of the animal and identified a variety of inconsistent relationships between the alternatives represented and current trial behavior. These results included an association between alternative representation content and head-scanning behavior[205], although non-local representations are also observed in the absence of this behavior[69,216]. These results also include reports that representations of specific alternatives either did[98,100,101] or did not[69,99,117] predictably relate to the upcoming choice on the current trial.

Our results provide a potentially unifying framework to explain these previous observations. We identified a modulation of non-local representational content with

relative value across successive trials. Specifically, when the difference between the values of the Stay and Switch options was large, the higher value alternative dominated non-local representations. When the differences were smaller, the relative proportion of non-local representation of each was more similar. These findings suggest that the expected values of options or related decision variables influence the extent to which those options are expressed in the context of non-local activity. This provides a potential prioritization of options to be considered, as the hippocampus would be most likely to represent the best, most viable options. Similar prioritization, in which the most promising options are internally sampled most extensively, also arises in rational computational analyses of choice under uncertainty[215].

Thus, when the preferred choice is clear, the hippocampus primarily represents that path, but nevertheless samples occasionally from the alternative, potentially to retrieve information about the path not taken. In conditions where the preferred choice becomes less clear, the hippocampus generates progressively more balanced representations of the two options. Critically, therefore, hippocampal non-local representations are neither constantly predicting the planned choice, nor strictly providing a menu that evenly represents all potential options—rather, representations can be more or less predictive of behavior depending on the current decision-making demands.

At the same time, relative value-related modulation of non-local representations suggests that the mechanisms selecting non-local representations already have an estimated value of the alternatives. If so, then why would it be useful to generate a non-local representation? We suggest that non-local representations activate place and

associated value representations on the actual path toward possible goals. This sampling of place and value could progressively refine and provide a more accurate estimate of the value, and thus serve to better inform decision processes[159,206]. This is consistent with our observation that the proportion of representation of the Switch path ahead was predictive of choice on even the next trial, and that similar patterns of non-local content extending behind the animal also related to upcoming choice, well in advance of physically approaching the next choice point.

Modulation of Non-Local Extent During Periods of Updating

Our results also suggest that non-local representations of distant locations may particularly be recruited to support updating an internal model. Behavior in the Spatial Bandit task was better fit by a model that included non-local value updates wherein experiences at one reward port could alter the values at other ports. The pattern of model-estimated updates around Switch trials was similar to the measured extent of non-local representations on these trials, with the largest updates and the most distant representations preferentially expressed after the Switch decision was made and on subsequent trials in the new patch.

One way to understand why value updating is associated with the representation of more distant locations is that it may enable more proximal value retrieval at the time of later choices. In the context of a choice to Switch patches, we found that animal behavior was driven in part by the reward value experienced at remote destination ports. Yet, the non-local representations in trials immediately preceding a patch Switch largely sampled paths along the neighboring track segments, rather than sampling distal

reward port locations. This is consistent with the widespread idea that the brain learns a value function, mapping locations near each choice point to the rewards to which the paths ultimately lead[185]. Thus, when outcomes are experienced at reward ports, the hippocampus may support non-locally updating information at distant locations across the maze in part to simplify later retrieval when that up-to-date information is needed for evaluating nearby choices.

While the potential mechanisms by which non-local representations during navigation could support updating merit further investigation, some initial points can be made. First, the extent of non-local representations was low as animals traversed the central segments on a Switch trial but increased substantially in the final segment; that is, non-local extent increased immediately before rather than only after the first outcome in the new patch was revealed. This suggests that the role of these remote representations is not dependent on having just experienced an outcome at a reward port. Instead, the hippocampus may facilitate post-outcome updating by activating representations of distant alternatives that should be subsequently updated after new outcomes are observed in the chosen patch. One possibility is that activation during running could promote reactivation during subsequent replay events at the reward port[217], and thus link the outcome of the current trial to representations of distant places. Second, prior work has demonstrated that non-local representations during navigation are strongly associated with the theta rhythm, and sequences along the theta rhythm are thought to support learning by binding together elements of experience at a compressed timescale through plasticity mechanisms[90,91,95,208,218]. Thus, this system could be used to learn relationships between current experience and unchosen

alternatives within theta cycles (or other organizing codes[152]), in support of updating knowledge of remote alternatives during active navigation.

These possibilities are consistent with prior work showing that animals make internal model-based choices that rely on structural knowledge[40,46,185,219–222] rather than just stimulus-response associations. Internal model maintenance and updating is particularly beneficial for flexible decision making to keep up with changing environments. However, previous neurophysiological and computational work on how non-local representations in the hippocampus may support value updating has primarily focused on replay during "offline" rest[74,77,78,84,120,125,164,223–226], rather than active locomotion[227,228]. Thus, our findings suggest that "online" movement-associated non-local representations in the hippocampus also have the potential to contribute to this learning.

A Broad Range of Alternatives Are Flexibly Engaged

Our results also establish that a diverse range of alternatives are expressed by hippocampal non-local representations. In the context of a specific decision among options, previous work has often focused on representations that corresponded to possible future locations on the current trial[98,99,117]. Our findings show that representations of alternative paths not only ahead of the animal but also behind the animal—even when moving away from the associated choice point—are flexibly engaged across trials associated with decision-making and learning processes. We also found that this engagement included non-local representations of very remote locations,

including representations that occasionally jump far across the maze rather than sweep ahead at a constant rate.

The dynamic engagement of a broad range of alternative possibilities by the hippocampus[69] during locomotion is consistent with the idea that non-local population representations can support a range of cognitive functions that involve computations about alternatives[3,17,229,230]. In particular, the dynamic engagement of representations behind the animal[95] during navigation raises the possibility that the hippocampus supports computations about actual and counterfactual[41] past paths during navigation, not only during rest. Moreover, the relative prevalence of representations behind the animal predicted behavior on the subsequent trial, indicating that representations of alternatives behind the animal need not be strictly retrospective, but could also support future decision making.

Additional observations from prior studies exemplify that hippocampal activity during locomotion can correspond to various correlates, including representations of opposite directions[69], forward and reverse sequences[110], distant locations[70,98], alternative environmental contexts[109], and different spatial reference frames[108]; our results are consistent with the idea that any of these representations of alternatives may be differentially engaged and specifically curated across trials depending on the cognitive computations needed for the task at hand.

Non-Local Computations in Immediate and Long-Term Adaptive Behavior

Retrieving information about alternative options related to making a specific upcoming decision and retrieving information about alternatives in service of updating an internal

model are both computations that inform inference across space ("non-local inference"). In the former case, representations of alternatives may support non-local inference of which alternative has the greater expected value based on prior experience. And, in the latter, representations of alternatives may support non-local inference about how the expected value of remote alternatives should be altered based on outcomes observed elsewhere. This possibility is consistent with prior work indicating that the hippocampus is important for inference[45,46,164,168,231–236], and relational thinking more broadly[49,161,237–240].

The prolonged relationship across trials of non-local representations to Switching behavior serves as a reminder that the behavioral benefit of non-local inference need not be immediate. While we observed that non-local representations can predict immediate future and past Stay or Switch choices, we also found that representations of alternatives ramped across many trials before and after Switch choices, that they predicted choices an entire trial in advance, and that remote representations in alternative patches did not predict prior or subsequent patch choices. Thus, while non-local representations can have immediate behavioral benefit by supporting the immediately upcoming choice, they may also support decision making over the longer term, perhaps by integrating and accumulating[241,242] retrieved information across trials. This is consistent with the idea that inferring structure in an environment and relations between experiences may occur in the background of behavior[243], particularly as new information needs to be incorporated into an internal model, even if an animal does not yet know if or when it will later need to draw on that knowledge to support adaptive choices. This is thought to be especially important in complex and dynamic

environments to enable flexible decision making when confronted with new or unexpected circumstances.

While the circuit-level mechanisms that determine the flexible generation of representations of alternatives remain largely unstudied, our observations of modulation with respect to relative values and the need to learn them implicates brain regions like the prefrontal cortex as possible drivers[244–251]. Consistent with this possibility, activity in the medial prefrontal cortex (mPFC) and the hippocampus can be precisely coordinated[117,252,253] and mPFC spiking can predict the future engagement of non-local hippocampal spiking[70]. Further, observations of coordination of hippocampal representations and peripheral sensory-motor processes[254] suggest broad engagement of many brain areas at times when hippocampal non-local activity may be important for ongoing behavioral computations. The specific generation of non-local hippocampal activity patterns and associated representations of alternative possibilities could engage widely distributed computations to enable learning and adaptive decision making in a complex and dynamic world.

Methods

Experimental Methods

Subjects. All animal procedures were approved by the Institutional Animal Care and use Committee at University of California, San Francisco. Five adult male Long-Evans rats (Charles River Laboratories, 450-650g) were each pair housed in a temperature-and humidity-controlled environment on a 12-hour light-dark cycle with lights on 6 AM - 6 PM. Rats had *ad libitum* access to standard rat chow and water. Prior to behavioral

training, rats were transitioned to single housing and food restricted to 85% of their baseline weight.

Behavioral Pre-Training. All behavioral tasks were controlled via custom code written in Python and Statescript in combination with an Environmental Control Unit (Spikegadgets). Animals were pre-trained to run back and forth on a linear track with walls for liquid food reward (Carnation evaporated milk sweetened with 5% sucrose) from reward ports fixed at each end of the track. Upon entry of a reward port, 100 μL reward was delivered by syringe pump immediately and automatically, gated by an infrared beam. Animals learned to alternate between and obtain rewards from the two ports across two 40-minute and one 25-minute sessions. To familiarize animals with navigating an elevated maze, animals then performed the same task on an elevated linear track (1.1 m long, 84 cm high) until they performed at least 30 rewards in a 15-minute session. Pretraining took place in an environment with distal spatial cues on the walls. The pretraining environments were fully separate from the Spatial Bandit Task environment. Animals were subsequently returned to *ad libitum* food access prior to surgical implantation.

Neural Implant. Custom hybrid microdrives contained 24 independently movable 12.5 μm diameter nichrome tetrodes (California Fine Wire and Kanthal). The drive body was 3D printed (PolyJet HD, Stratasys) and funneled tetrodes into two cannulae to target 12 tetrodes to each hemisphere. Tetrodes were gold plated to reduce impedance to ~250 kOhms. Implants also contained a custom headstage (Spikegadgets) that coupled with

a stacking set of up to four custom 128-channel polymer probes[255] (Lawrence Livermore National Labs) per animal. Implant was housed in a custom 3D-printed enclosure with removable cap. Drive was sterilized before surgical implantation.

Custom hybrid microdrives were implanted stereotaxically under sterile conditions. In anesthetized animals, cannulae were implanted above dHC (+/-2.6-2.8 ML, -3.7-3.8 AP relative to skull Bregma) and polymer probes were implanted in mPFC and OFC. A stainless-steel ground screw was implanted over the cerebellum to serve as a global reference. Titanium screws were placed in the skull to help anchor the implant, which was secured with Metabond (Parkell, Inc) and dental cement (Henry Schein).

After surgery, tetrodes were advanced deeper into the brain daily over ~3-4 weeks. One tetrode per hemisphere was advanced to corpus callosum as a local reference and all other tetrodes were advanced to dorsal hippocampus CA1 stratum pyramidale, guided by physiological signatures of spiking activity and the local field potential.

After full recovery from surgery, animals were food deprived to 85% of their baseline weight and reintroduced to the elevated linear track, as described above. Animals reached criterion again over several days to refamiliarize them with obtaining reward on a track and to familiarize them with running with their implant. Two rats had additional experience on a fork maze[256]. Re-training animals to run after surgery recovery ensured that the animals were motivated enough to begin the main experiment.

Data Collection: Spatial Bandit Task. The Spatial Bandit task took place on a maze with three radially arranged Y-shaped foraging “patches” each containing a central hallway that bifurcated into two hallways that each terminated in a photogated reward port, resulting in two ports per patch and six total reward ports. Each port had a separate automated pump, as described above. Hallways were 6.5 cm wide and linear track segments were each 53 cm long. All hallways intersected at 120 degrees, such that the track had three-fold symmetry. The track was made of black acrylic (TAP Plastics) with walls that were 3 cm high, enabling animals to see distal spatial cues.

The behavior took place in a dimly lit room (2.4 m by 2.9 m) with black distal spatial cues on the white walls and a plastic black curtain separating the maze from the experimenter, rig, and computer. A ceiling-mounted camera (Allied Vision) centered over the maze recorded animal behavior at 30 frames/second and was synchronized with all other data via Precision Time Protocol. Before behavior began each day, a ring of red and green LEDs was mounted atop the implant to enable online head position and direction tracking via SpikeGadgets Trodes software.

Animals began the task only once tetrodes had reached stratum pyramidale. Neural data was recorded from the animals’ first experience of the Spatial Bandit maze and environment. The general structure of each day of recording began by moving rats from their home cage into a rest box in the same room as the Spatial Bandit maze. Run sessions were interleaved with rest sessions throughout the day; rats typically completed 7 or 8 ~20-minute run sessions per day, interleaved with rest sessions of at least 30 minutes each. Rest sessions helped to maintain stable motivation by providing

breaks throughout each day. Each day also started and ended with a rest session. These behavioral data were collected over 8-17 days per animal.

During each run session, an animal was first placed in the center of the track facing the same wall each time, and then navigated freely in the environment, which was otherwise fully automated by Trodes (SpikeGadgets) software and custom behavioral control scripts written in Statescript (SpikeGadgets) and Python. At a high level, each individual run session contained multiple reward contingency blocks. Each contingency defined the reward probability p(R) assigned to each of the six reward ports. Contingency blocks covertly changed when an animal completed a certain number of trials, or hit a 20-minute time limit per contingency, whichever came first. Contingency changes almost always occurred based on the trial limit, very rarely changed based on the time limit, and did not depend on any other performance metric. A trial was defined as the period from exiting one photogated reward port to exiting a different port. Therefore, each trial contained both the run between two ports and an outcome (100 μL reward or omission) at the chosen port. Reward was only available probabilistically if an animal poked into a distinct port; consecutive pokes at the same port were never rewarded.

Each animal's first run session began with reward probability p(R) of 1 at all reward ports for 100 trials, followed by an uncued contingency change to p(R) of 0.5 at all ports for the next 100 trials. This first session exposed the animal to the environment, encouraged the animal to visit all reward ports, and introduced probabilistic rewards. After this session, run sessions were each 180 trials long and contained three reward

contingencies that each lasted 60 trials for four animals. One animal had 160 trial sessions each containing two contingencies of 80 trials.

From the second run session onward, each reward contingency defined a p(R) of 0.2, 0.5, and 0.8 per port, such that each contingency had a best patch on average. Both ports within a patch could have the same or different reward probabilities, but the best patch combination was 0.5 and 0.8 (not 0.8 and 0.8) across the two ports in a patch. By having ports of two different values within a patch, animals were encouraged to spend some trials not only at high-value but also at nearby lower-value locations, enabling us to sample neural data as animals visited ports of a full range of values. Contingencies were pseudorandomized to counterbalance which patch was best, whether the left or right port within each patch was best, and whether the best, medium, and worst patch followed a clockwise or counterclockwise order around the track. Contingency changes and reward port values were never cued, so that animals had to make navigational choices based on memory of their prior experiences across the maze.

Continuous neural data, as well as digital inputs and outputs associated with beam breaks and reward pumps, were recorded during each rest and run session at 30 kHz in four animals and 20 kHz in one animal using Trodes version 1.8.2 (SpikeGadgets).

Histology. At the conclusion of behavioral experiments, animals were anesthetized with isoflurane, and small electrolytic lesions were made. One day later, animals were anesthetized and transcardially perfused with 4% paraformaldehyde. Brains were fixed

in situ overnight, after which tetrodes were retracted from the tissue and brains were extracted from the skull. Brain fixation continued for two days. After cryoprotection by 30% sucrose for 5-7 days, brains were blocked and sectioned with a cryostat into 50-100 μm sections. Nissl labeling enabled identification of tetrode tracks and localization of tetrode tips at the sites of electrolytic lesions (Figure 2.6).

Data Processing and Analyses

All data processing and analyses were carried out using Python and Julia using Spyglass[257].

Behavioral Analysis. Trials were parsed based on behavioral Statescript log files (SpikeGadgets). A trial was defined as the final exit from one port to the final exit from another port, and therefore included both the run from one port to the next as well as the entry into and outcome at the chosen port. Stay trials were defined as starting and ending within the same track patch, while Switch trials were defined by starting and ending in distinct track patches. The nominal reward probabilities assigned to each port and reward delivery times were also tracked by parsing log files. The nominal best patch on each trial was defined as the patch with the highest average of the nominal reward probabilities of the two ports within the patch (Figure 2.1). These reward probabilities are referred to as nominal because they are assigned by the reward contingency. Importantly, an animal's experienced reward values at each port at any time in behavior are related to but not necessarily identical to the location's nominal

reward probability, given the self-paced nature of the task and stochasticity of the rewards in the environment.

Behavioral Hidden Markov Model. To estimate the rats' expected belief about the value of each of the six ports, we used a Hidden Markov Model (HMM), which tracks belief across a discrete set of *states* of the world via a forward model of expected reward outcomes. The rationale for this modeling approach is to abstract away many under-constrained implementational decisions in a more mechanistic model, and instead take as a starting point the question of what inferences would be drawn by an ideal statistical observer experienced with the task environment's structure. More specifically, the HMM is capable of modeling the true dynamics of how rewards change during the experiment by having states of the HMM directly map onto the joint reward probabilities across the six reward ports, allowing it, for example, to capture the intuition that a change of reward probability at one patch should indicate to the animal that rewards have likewise changed at the other two patches, an inference that is not feasible with traditional Q-learning models. That said, we expect that a more mechanistic implementation would approximate the HMM computations by augmenting a Q-learning style model with non-local activations and updates[223,258].

Here, *state* is an assignment of expected reward probability to each of the six ports, and whose dynamics directly capture the ideal estimation strategy of an agent aware of the true underlying task structure, but without explicit knowledge of the precise timing of state changes. Notably, unlike traditional Q-learning models, the HMM captures the high-level idea that state configurations jointly change across each

session, and that each trial can be viewed as evidence of which among these possible state configurations the rat is currently experiencing.

The HMM tracks value of reward ports by assigning a probability, or belief, to each possible state configuration of the six reward ports, alongside dynamics to update these beliefs via per-trial outcomes. Here, we considered a *state* to be a particular assignment of reward probabilities to each of the 6 reward ports, such as:

$$\phi_i = [p_1 = 0.2, p_2 = 0.2, p_3 = 0.5, p_4 = 0.2, p_5 = 0.8, p_6 = 0.5]$$

Each trial can be viewed as evidence of which among these possible state configurations the rat is currently experiencing, allowing trial-by-trial state inference and thus value estimations of each of the reward ports.

Reward Port Value Estimation. In the HMM, we are modeling belief about which of n discrete states the rat is currently in. These states correspond to the contingencies that define reward probabilities for the rats. For computational convenience, and motivated by the informed ideal observer perspective, our set of contingencies, ϕ, was the subset of the $n = 72$ most frequently occurring of these possible contingencies in the actual experimental design, though we found our results to be relatively insensitive to the exact choice of contingencies.

Our belief state, α, is a probability distribution over these 72 contingencies such that $\sum \alpha_i = 1$. On any given trial, we can combine our belief state α with the set of contingencies ϕ to find an estimate of reward Q at each of the six ports that minimizes the MSE:

$$\boldsymbol{Q} = [Q_1, Q_2, Q_3, Q_4, Q_5, Q_6] = \sum_{i=1}^{n} \alpha_i \phi_i$$

Patch Value Estimation. Patch values were derived from the reward port estimates as follows: As the rat approaches a bifurcation and must choose between the upcoming port within a patch or the ports within the other patches, Q_{patch} for the current patch was defined as the Q-value of the upcoming reward port as estimated by the HMM. For the other two patches, it was assumed that their value was an average of their two reward ports as estimated by the HMM, or $Q_{patch} = (Q_a + Q_b)/2$, where a and b were the two reward ports of the respective patch.

Behavioral Model State Updates. We assumed that any contingency could transition to any other contingency with a fixed probability, or *hazard rate* h*,* which is the probability of transitioning to any new contingency, and $1 - h$ is the probability of remaining in the current contingency. This hazard rate was assumed to be identical among all state transitions, and was a free parameter inferred via fitting.

At each timestep, beliefs were updated via $\boldsymbol{\phi} \leftarrow \boldsymbol{T}\boldsymbol{\phi}$, where $\boldsymbol{T}$ was an $n \times n$ matrix whose diagonal entries were $1 - h$ and all other entries were $h/(n-1)$. This essentially results in a small diffusion of all reward port values around the maze towards baseline.

To incorporate observations (here, whether a port was rewarded), at the end of each trial we updated each entry of α based on the probability of the observed outcome for each respective contingency. If the outcome was positive, the probability given by a contingency was the value of ϕ for that contingency and port, and so $\alpha_i \leftarrow \alpha_i \, \phi_{ij}$ for

contingency i and port j. If the outcome was negative, this became $\alpha_i \leftarrow \alpha_i(1 - \phi_{ij})$. This update was performed for all contingencies.

Behavioral Choice Model. To connect this learning model to choices and assess our ability to capture animal behavior, we computed the likelihood of choice of patch on every trial as the *softmax* of the three patch values, given by:

$$p_{patch_i} = \frac{e^{V_{patch_i}}}{e^{V_{patch_1}+V_{patch_2}+V_{patch_3}}}$$

The value at the current patch was given by:

$$V_{patch_i} = \beta_{stay} Q_{patch_i} + b_{stay}$$

For the patch clockwise from the current patch, value was given by:

$$V_{patch_i} = \beta_{go} Q_{patch_i} + b_{turn}$$

The value of the remaining patch was given by:

$$V_{patch_i} = \beta_{go} Q_{patch_i}$$

where Q_{patch_i} was either an individual port estimate or average of ports as described above.

The inverse temperature parameters, β_{go} and β_{stay}, determined how sensitive choice behavior was to differences in expected reward – high temperatures indicate high sensitivity to reward, while low temperatures indicate insensitivity. To allow for differential sensitivity to reward estimates at the current patch versus alternate patches, we fit two independent softmax temperatures, β_{stay} for the current patch, and β_{go} for the two alternative patches.

The term b_{stay} reflected a fixed cost to changing patches due to the time and distance to reach a new patch (as opposed to the reward sensitivity fit by β_{stay}), and was added to the valuation of the current patch.

To incorporate tendencies of the animals to prefer to turn in certain directions when changing patches, we included a *turn bias* b_{turn}, an offset which was consistently applied to the patch to the left of the animal's current position. If the rat was in patch 1, this was added to patch 2. If the rat was in patch 2, this was added to patch 3, and if the rat was in patch 3, this was added to patch 1.

On trials where the animal switched patches, likelihood of port choice was modeled as a subsequent softmax between the two ports, with a separate softmax temperature β_{port} applied to the per-port reward estimates:

$$p_{port_{L/R}} = \frac{e^{V_{port_{L/R}}}}{e^{V_{port_L}} + V_{port_R}}$$

The value of the left port was given by:

$$V_{port_L} = \beta_{port} Q_{port_L} + b_{portbias_i}$$

where $b_{portbias_i}$ was fit independently for each of the three patches, reflecting biases in which reward port each rat tended to enter when switching to each of the three patches. The value of the right port was given by:

$$V_{port_R} = \beta_{port} Q_{port_R}$$

Local Behavioral Model. To test whether animal behavior reflected a belief that the outcomes at the six reward ports were linked, a core assumption of our Global behavioral model, or whether the animals instead estimated the value of the six reward

ports independently from one another, we considered an alternate model in which we estimated animal behavior via a set of six *independent* HMMs, one for each of the six reward ports (Figure 2.5). Each HMM individually contained the same distribution of rewards as in the distribution in our joint HMM model, as well as a belief distribution over those rewards. However, the belief distributions at the six reward ports were independent – while volatility would be modeled for all six reward ports after each trial, only the belief distribution of the visited reward port would be updated from observation via the forward model on each trial – the other five belief distributions would not incorporate this reward information. Importantly, the comparison between these two models (rather than, for instance, between the HMM and a local Q learning model) most directly isolates the question of local versus non-local updating, since the models are identical in all other features.

To compare the independent Local model with the original joint Global model, we computed a cross-validated approximation to the negative log marginal likelihood for each day. Specifically, we used leave-one-session-out cross validation for the population-level prior parameters and a Laplace approximation for the per-day parameters: for each day, we refit the population-level model omitting that day, then conditional on that prior, we computed a Laplace approximation to that day's log marginal likelihood. We aggregated these per-day scores to obtain a total score for each rat and model. Finally, we use paired tests on these scores across rats, between both models, to formally test whether either model fit consistently better over the population of rats (Figure 2.5).

Global Model Reward Sensitivity. To confirm that the HMM's behavioral predictions were driven by its reward estimates, and that incorporating these estimates improved its performance, we compared our Global behavioral model against two nested models, one model where β_{go} was fixed to 0 (such that choice predictions ignored the estimate of the current patch value) and one model where β_{stay} was fixed to zero (such that choice predictions ignored the estimate of alternative patch values). These two alternative models were compared against the full model (Figure 2.10) with an identical procedure to the Local model comparison (see above).

Behavioral Model Fitting Procedure. For our HMM, we optimized the free parameters of the model by embedding each of them within a hierarchical model to allow parameters to vary from day-to-day. Day-level parameters were themselves modeled as arising from a population-level Gaussian distribution over days, separately for each rat. We estimated the model for each rat to obtain best fitting day- and rat-level parameters to minimize the negative log likelihood of the data using an expectation-maximization algorithm with a Laplace approximation to the day-level marginal likelihoods in the M step[259].

Behavioral Model Variables. In subsequent analyses, we characterized value-guided Stay and Switch choices at the level of trials using variables estimated through the HMM, including the values associated with each port on each trial. We built a binomial family generalized linear model (Logit link function) for each animal to predict Stay or Switch choices from estimated expected values associated with the choice to Stay or to

Switch (Figure 2.1). The value of Staying was defined as the model-estimated value of the upcoming port within the patch on the current trial, and the value of Switching was defined as the greater of the average port values in the two alternative unoccupied patches into which the animal had the option to Switch on the current trial. The difference between the Switch and Stay values on each trial, without accounting for the cost of Switching or other biases, defined the relative value of Switching and Staying on each trial (Figure 2.1). The entropy over the reward belief state of the upcoming port (Figure 2.5) was calculated on each trial as follows. Given a belief state α, and reward probabilities ϕ, each port j had a marginal reward distribution $[p_1,p_2,p_3]$ where $p_1 = p(r = 0.2),\ p_2 = p(r = 0.5),\ p_3 = p(r = 0.8)$. For each port j, $reward\ entropy_j = -\sum_i p_i \log(p_i)$. The value update at each port on each trial was calculated as the difference between the port value before and after the outcome was observed on each trial, and absolute value updates across a patch, set of patches, or the entire maze, were summed across ports to determine the total update magnitude on each trial (Figure 2.5, Figure 2.10).

Position. LEDs mounted on the front and back of the implant were tracked via overhead camera at 30 frames/second using Trodes software (SpikeGadgets). Positions were converted from pixels to centimeters, smoothed, interpolated and upsampled to 500 Hz to match neural decoding resolution. The front and back LEDs were used to calculate head orientation and angular velocity. The head position and head speed were calculated from the centroid of the font and back LEDs. To enable faster decoding, head position was linearized into one dimension by projecting the two-dimensional head

position onto a graph representation of the track with nodes at the reward ports and track segment intersections, connected by edges. In between the track edges, which corresponded to the nine physical track segments, we introduced 15 cm gaps in 1D space to avoid inappropriate smoothing across adjacent linear positions that were not adjacent in physical space (Figure 2.6).

Spatial Decoding. To decode spatial position from hippocampal spiking at high spatial and temporal resolution, we used a previously established clusterless state space algorithm[212,213,260,261,84,262,254] from Denovellis et al., 2021[212]. Briefly, for each session, we built an encoding model to relate spike waveform features to the animal's linearized head position in 2 ms bins using data from the beginning of the first trial to the end of the final trial of the session. The peak amplitude of each spike on each dCA1 tetrode channel served as spike waveform features. Spikes were detected by a 100 μV threshold on the recorded neural signal filtered 600-6000 Hz. To remove potential artifacts, a 1 ms window around times in which at least 75% of channels had an amplitude of 2000 μV or greater was excluded.

For each session, we then decoded spatial position from this hippocampal population spiking in 2 cm linear position bins and 2 ms temporal bins (Figure 2.6). The state space decoding model[212] had two states, spatially continuous and spatially fragmented, to capture smooth and discontinuous movement dynamics, respectively, of the hippocampal representation through the maze. The fragmented state was modelled by a uniform transition matrix and the continuous state was modelled by a random-walk transition matrix with a 6 cm movement variance. The probability of staying in the

current state in the discrete transition matrix was 0.98. The initial conditions were set to be uniform across states and uniform across positions, as we did not have prior information about the most likely initial conditions. The model used kernel density estimators with a 24 μV bandwidth for waveform features and a 6 cm bandwidth for position to estimate the likelihood of a spike waveform feature predicting a position.

The model output the posterior probability of position across the two movement dynamic states. We marginalized the joint probability over the two states, and then determined the most likely decoded position in each 2 ms bin as the peak of the posterior probability distribution.

We assessed decoding quality based on the concentration of the posterior probability, and only included times that were decoded with high confidence in subsequent analyses, by requiring the highest 50% of the posterior probability values to be concentrated within 50 cm of the track, as has been done previously[254]. We further note that this decoding approach was conservative in the sense that we decoded using the same data used to build the encoding model. In contrast, a cross-validated approach would instead decode periods distinct from those used to encode, and thereby exclude the place cell spikes predominantly associated with the animal's actual current position during the decoded moments.

We measured the distance between the animal's actual position to the most likely decoded position as the shortest path distance in linear space along the track in each 2 ms bin based on Dijkstra's[263] algorithm, with the rat and most likely decoded position as nodes on the track graph. We also classified the animal position and most likely decoded position into one of nine track graph segments in each 2 ms bin. For

visualization, we projected linearized position back to the 2D track graph (Figs. 2.2, 2.3, 2.4).

Non-local representation of Stay and Switch paths. Non-local representation times were identified as times in which the most likely decoded position (hereafter "decoded position") was in a track segment distinct to the animal's actual position. We analyzed valid times of high speed (>10 cm/s) in order to limit our analyses to locomotor periods in which theta power is high and associated theta sequences are known to occur[89,200,204,264], and thereby excluded stationary rest periods in which sharp wave ripple replay events occur[53,74,80].

To analyze the content of non-local representations occurring as the animal approached the first choice point on each trial (Figure 2.2), we first identified which track segment the animal began in on each trial and used data from the time at which the animal exited the reward port to the first time at which the animal's position exited that initial track segment. During this period, the animal approached the first choice point in the trial (Figure 2.2). Out of this period, we limited analysis to only the times that satisfied the decoding confidence and speed thresholds described above. For these valid times, we then quantified the total duration of non-local representation. We also quantified the duration of representation of the single segment ahead consistent with a subsequent Stay path and the duration of representation of any of the other segments ahead consistent with a subsequent Switch path, as in Figure 2.2. This enabled us to then calculate the proportion of the total valid non-local representation duration that was consistent with a non-local Switch path representation. The per-trial proportion was

therefore calculated as: valid non-local Switch representation duration, divided by the sum of the valid non-local Switch representation duration and the valid non-local Stay representation duration.

A very similar approach was used to analyze the content of non-local representations occurring as the animal traversed the final track segment of each trial (Figure 2.3). Here, we identified which track segment the animal ended in on each trial and used data from the final time at which the animal crossed from another segment into this final segment to the time at which the animal entered the chosen reward port. During this period, the animal approached the reward port, rather than a bifurcation (Figure 2.3). Again, we limited analysis to only the times that satisfied the decoding confidence and speed thresholds described above. For these valid times, we then quantified the total duration of non-local representation, and what proportion of that total non-local duration corresponded to a non-local representation consistent with a Switch path behind the animal (rather than the neighboring Stay path within the patch), as in Figure 2.3.

Peri-Switch Linear Regressions and Shuffles. Ordinary least squares linear regressions were fit to Stay trial data for trials before or after Switch trials (Figure 2.2, Figure 2.3, Figure 2.7). Shuffled distributions were obtained within animal by randomly shifting trial labels 1000 times, where the label is the number of trials from either the prior or next Switch trial.

Predicting Stay and Switch Choices from Non-Local Representations. Generalized linear models with a logit link function (Figure 2.2, Figure 2.3) were used to predict Stay or Switch choices from the proportion of valid non-local representation corresponding to Switch paths. Models were fit separately per animal. Each model had one predictor, corresponding to the proportion of Switch path representation occurring while the animal occupied either the first or final segment of a trial t, the trial before t-1, or the average across the two trials t and t-1. Models were five-fold cross validated, and training data were balanced by resampling Switch trials. Predictions were binarized with a threshold of 0.5 and total accuracy was calculated across folds.

Theta Phase. Neural data from a reference electrode in corpus callosum was filtered 0-400 Hz and downsampled to 1 kHz. This local field potential signal was then further filtered in the theta band from 5-11 Hz. The analytic signal was calculated with a Hilbert transform. The instantaneous phases for all valid local and non-local decoded run times were assessed separately for periods in which the animal was approaching the first choice point in each Stay or Switch trial, or after crossing the final choice point in each Stay or Switch trial (Figure 2.7). Phase convention refers to descending phases of corpus callosum theta as early phases, 0 to π, and ascending phases as late phases, π to 2π. Data were binned for visualization only. Full theta phase distributions for local and non-local times were compared with Kuiper's non-parametric test for continuous circular data for each animal.

Non-Local Distance. Non-local distance was calculated as the distance between the animal's actual position and the most likely decoded position in each 2 ms bin (see spatial decoding methods). Using only high-quality valid running times, as described above, we identified the maximum absolute actual-to-decoded distance that occurred during the animal's approach to the first choice point on a trial (Figure 2.4) or during the animal's traversal of the final segment of a trial (Figure 2.4). Data in Figure 2.4 were calculated based only on representations that occurred in non-local, unoccupied, track segments to limit representations to those either ahead of or behind (Figure 2.4) the animal's trajectory.

To compare the non-local distance dynamics before and after Switches for each animal, we log-transformed the model $y = ae^{bx}$ into linear form and applied ordinary least squares regression to estimate $\log(a)$ and b, for both pre- and post-Switch data. Z tests were used to statistically compare the pre- and post-Switch parameters (Figure 2.4).

Figure 2.4 Switch trial maximum distances across track segments were not required to be in a distinct non-local track segment to the animal's actual position. Data occurring while the animal occupied the middle two segments of Switch trials (Figure 2.4) were calculated based on Switch trials in which exactly four segments were occupied. Wilcoxon rank sum tests were used for each animal to statistically compare distances observed between different segments.

Confidence Intervals and Statistics. Error bars, unless otherwise stated, correspond to 95% confidence intervals on the mean. To take a nonparametric approach to

determining the statistic, we randomly resampled with replacement from the underlying data distribution 1000 times and calculated the statistic's bootstrap distribution. Statistical tests are stated throughout the figure legends and text with sample sizes and p values.

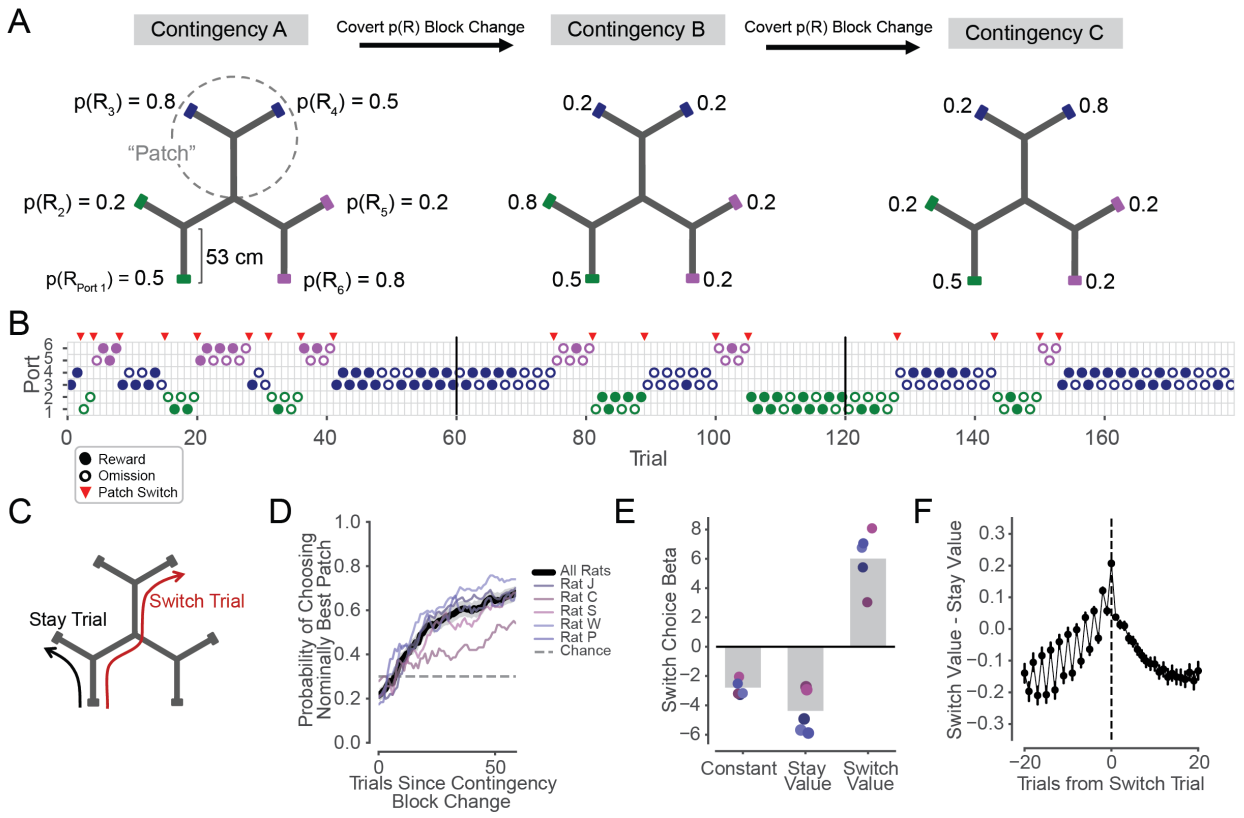

Figure 2.1 Rats make experience-guided decisions in the Spatial Bandit task
(A) Track schematic with example reward Contingencies A-C from one behavioral epoch. The track contains three Y-shaped "patches," each with 2 reward ports located at the ends of linear segments. Probability of reward at each of 6 identical reward ports is indicated as $p(R_{1\text{-}6})$. Ports are colored by patch for figure visualization only. Uncued reward contingency changes occur every 60 trials. **(B)** Example sequence of choices from one behavioral session. Vertical black lines indicate contingency changes shown in A. Circles indicate chosen port on each trial and are colored by patch. Filled circles represent rewarded trials, empty circles represent unrewarded trials. Red triangles indicate a patch Switch. **(C)** Schematic of example Stay (black) and Switch (red) trial trajectories. **(D)** Proportion of trials in which the animal chose a port in the patch with the highest average nominal reward probabilities, as a function of trial number within the contingency block. Black line and grey shading indicate average across rats and 95% CI on the mean. Distribution on 60th trial of blocks is significantly different from chance level of 1/3 in each animal (n=168, 87, 78, 131, 87 blocks for animals J, C, S, W, and P, respectively, p=2.7e-22, 8.7e-5, 6.4e-10, 2.5e-21, 6.1e-11, binomial test), and distribution on 60th trial is also significantly different from values on first trials of blocks (p=9.6e-20, 3.9e-4, 2.5e-9, 3.2e-18, 1.5e-11, Z test for proportions) **(E)** Regression coefficients on model-estimated values of Stay locations and Switch locations for the prediction of Stay or Switch choice on each trial. All individual coefficients are significantly different than zero, indicating that rats rely on reward history from both current "Stay" patch and alternative "Switch" patches to make Stay or Switch choices
(Figure caption continued on the next page.)

(Figure caption continued from the previous page.) (for each animal, $p_{Constant}$=3.8e-25, 1.6e-111, 5. 6e-98, 1.5e-39, 8.2e-40, p_{Stay}=5.6e-17, 7.3e-24, 3.8e-47, 1.5e-39, 8.2e-40, p_{Switch}=1.2e-23, 8.6e-9, 8.9e-30, 4.9e-25, 1.6e-22, from n = 5416, 8572, 9608, 7450, 4732 Stay trials and n = 344, 398, 583, 467, 488 Switch trials). Grey bars indicate averages and coefficient points are colored by subject per legend in D. **(F)** Relative value of Switching versus Staying increases across trials leading up to patch Switch choices, peaks on Switch trials, and decreases following patch Switch trials. Switch trial values are distinct from Stay trial values in the 20 trials before (p=4.3e-61, 2.1e-44, 4.5e-134, 6.6e-126, 4.3e-113, Wilcoxon rank sum test) and after (p=8.4e-81, 5.3e-63, 8.5e-162, 4.0e-149, 1.4e-132) Switch trials in each animal. Individual animal data in Figure 2.10.

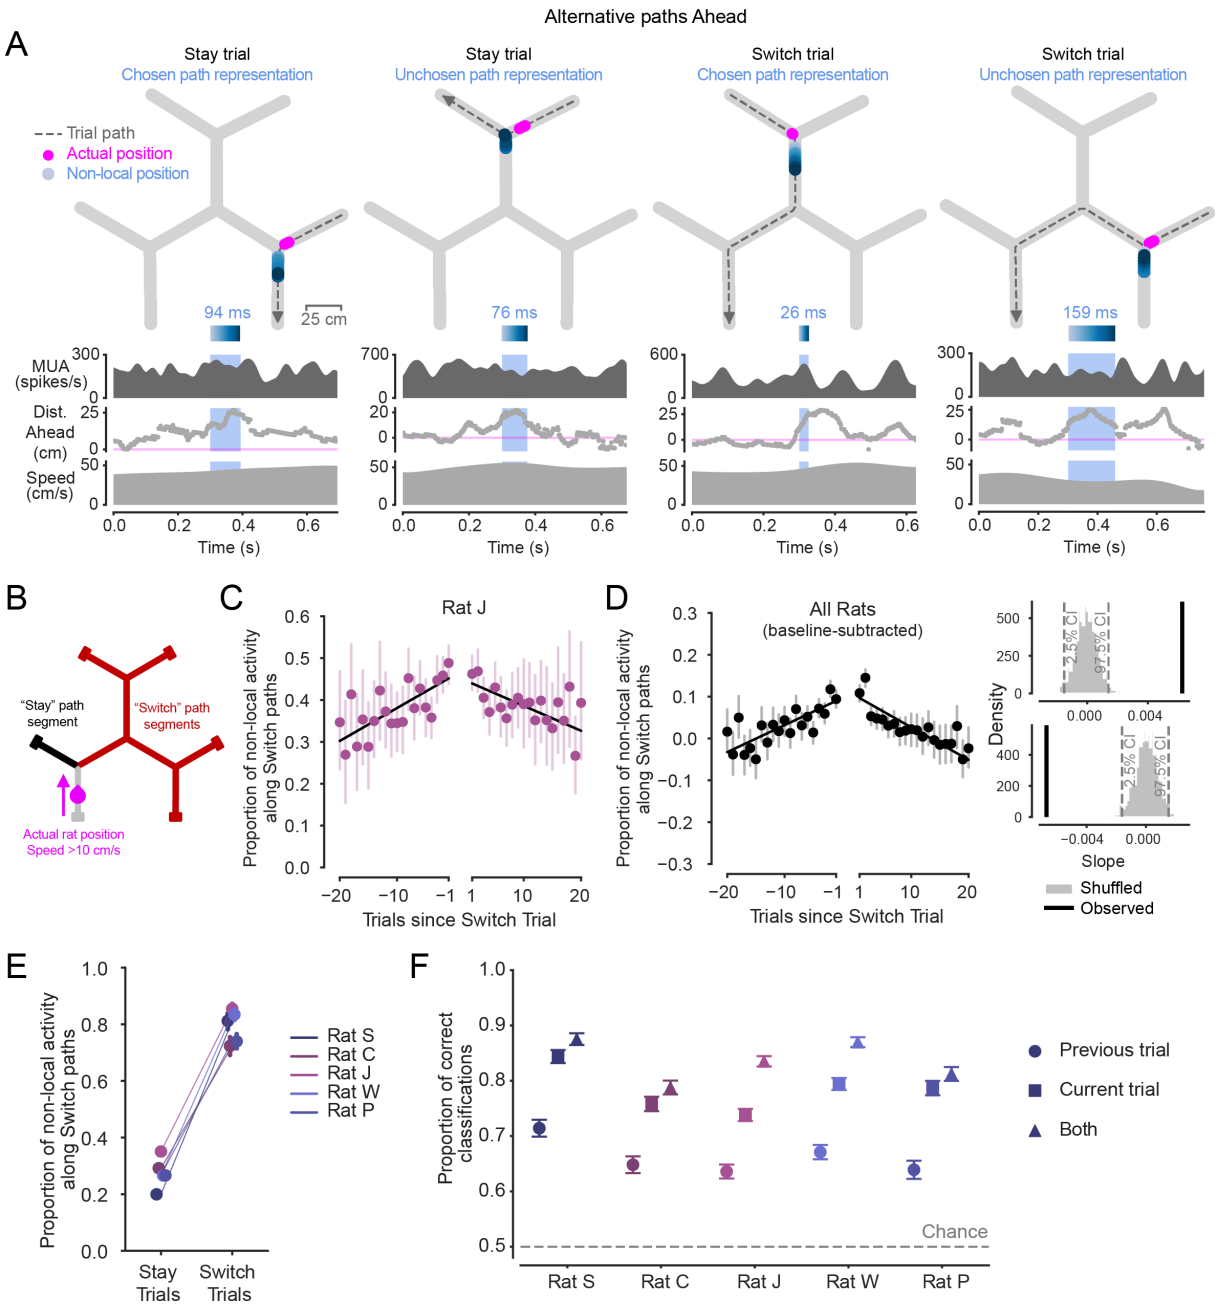

Figure 2.2 Non-local representations of alternative paths Ahead are enriched across trials before and after patch Switching

(A) Examples of non-local representations of alternative paths ahead on Stay and Switch trials. As the animal approaches the choice point, the hippocampus can represent the Stay or Switch path ahead. Top: Animal's actual position (magenta) and decoded position from hippocampal spiking (blue, shaded by time) during the light blue-highlighted period (below) are plotted on the track. Dashed grey line with arrow indicates animal's path and direction on the current trial. Highlighted non-local duration is labeled above blue time-shaded heatmap. Top-middle: multiunit spike rate from hippocampal tetrodes before, during, and after highlighted non-local representation. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) Bottom-middle: Distance from actual position to decoded position, with positive values corresponding to in front of the animal's heading, and negative numbers corresponding to behind the animal's heading. Note that blue non-local periods are required to have representations in a different track segment from the animal's actual position. Bottom: Animal speed. **(B)** Schematic defining Stay and Switch non-local activity ahead, analyzed in C-F. Non-local activity was analyzed from times in which the animal was running >10 cm/s from the reward port to the first choice point in the trial (magenta) by assessing the per-trial proportion of time in which non-local activity extended along the path ahead consistent with Switching (red) out of all non-local activity times along both the Stay (black) and Switch (red) paths. **(C)** Proportion of all non-local activity that represents paths consistent with Switching during the approach of the first choice point across Stay trials preceding and following Switch trials in one rat. Left and right x axes separated to reflect patch change. Error bars are 95% CIs on the mean. Linear regressions show increasing and decreasing proportions across trials before and after patch Switch trials. **(D)** Left: Proportion of all non-local activity that represents paths consistent with Switching during the approach of the first choice point across all animals (n=5 rats), normalized per animal by subtracting off the average baseline proportion on Stay trials. Error bars are 95% CIs on the mean. Pre- and post-Switch linear regressions overlaid. Upper right: observed slope of pre-Switch regression (black) is significantly different from 0 ($p<0.002$) based on 1000 shuffles of the underlying data (grey). Lower right: observed slope of post-Switch regression (black) is significantly different from 0 ($p<0.002$) based on 1000 shuffles of the underlying data (grey). Both slopes were individually significant in all five animals (Figure 2.7). **(E)** Proportion of all non-local activity that represents paths consistent with Switching during the approach of the first choice point in each animal on Stay trials and Switch trials. Error bars are 95% CIs on the mean. The proportion on Switch trials is greater than on Stay trials for all animals (p=2.7e-192, 4.2e-108, 1.0e-100, 2.1e-315, 1.6e-201, Wilcoxon rank sum test, for n=313, 335, 523, 422, 452 Switch trials, and n=3349, 3803, 5600, 5028, 3294 Stay trials). **(F)** Cross-validated accuracy of logistic regressions that predicted Stay or Switch choices based on the proportion of all non-local activity that represents paths consistent with Switching during the approach of the first choice point. Neural data from each of three trial types: previous trial (circle marker), current trial (square marker), or both (triangle marker). Error bars are 95% CIs on the proportion. Training data were balanced, so chance level was 0.5. Accuracy of model using previous trial neural data is significantly greater than chance level in all animals (p=3.2e-63, 1.1e-43, 4.0e-133, 9.7e-131, 3.3e-60, Z test for proportions), accuracy from model using current trial neural data is greater than accuracy from model using previous trial neural data (p=2.3e-39, 7.0e-27, 6.0e-33, 2.8e-46, 6.3e-43), and accuracy of model based on both trials is greater than accuracy of current trial model (p=1.1e-4, 1.2e-3, 4.3e-39, 4.9e-26, 5.6e-3).

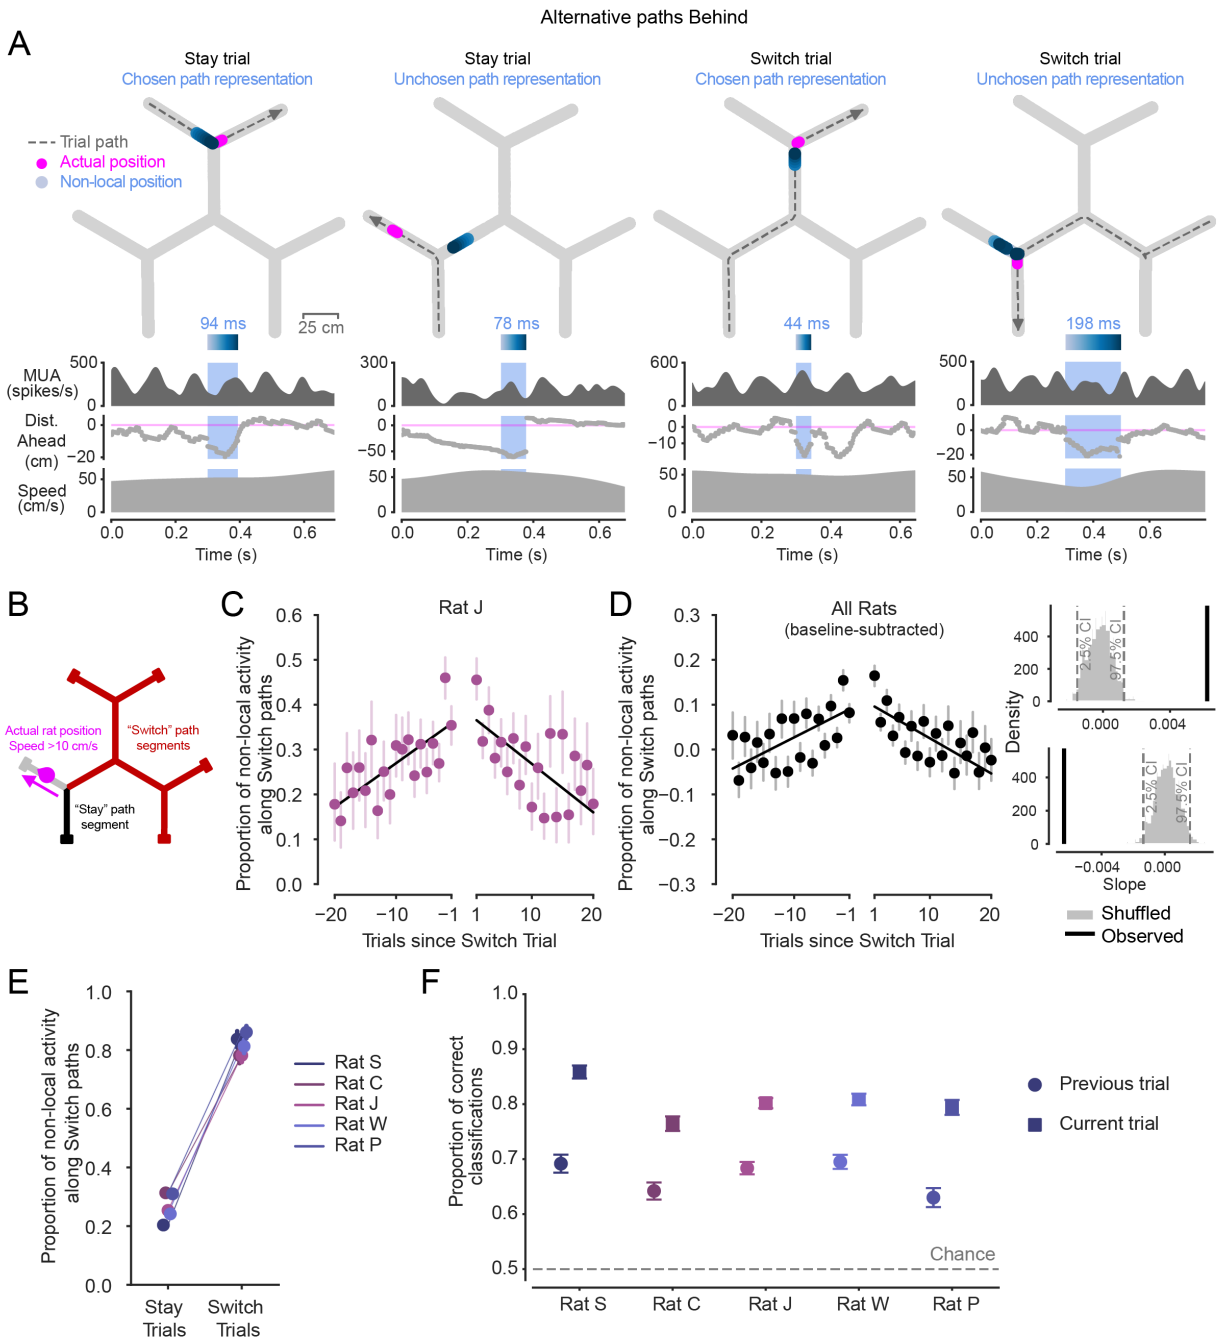

Figure 2.3 Non-local representations of alternative paths Behind are also enriched across trials before and after patch Switching

(A) Examples of non-local representations of paths behind in Stay and Switch trials. Top: Animal's actual position (magenta) and decoded position from hippocampal spiking (blue, shaded by time) during the blue-highlighted period (below) are plotted on the track. Dashed grey line with arrow indicates animal's path and direction on the current trial. Highlighted non-local duration is labeled above blue time-shaded heatmap. Top-middle: multiunit spike rate from hippocampal tetrodes before, during, and after highlighted non-local representation. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) Bottom-middle: Distance from actual position to decoded position, with positive values corresponding to in front of the animal's heading, and negative numbers corresponding to behind the animal's heading. Bottom: Animal speed. Note that periods outside the blue shaded region where the distance behind is large (e.g., second example) correspond to representations that are in the same track segment as the animal at that time rather than an alternative non-local segment. **(B)** Schematic defining Stay and Switch non-local activity behind, analyzed in C-F. **(C)** Proportion of all non-local activity that represents paths consistent with Switching during the approach of the reward port across Stay trials preceding and following Switch trials in one rat. Left and right x axes separated to reflect patch change. Error bars are 95% CIs on the mean. Linear regression lines show increasing and decreasing proportions before and after patch Switch trials occur. Individual data for all rats shown in Figure 2.7. **(D)** Left: Proportion of all non-local activity that represents paths consistent with Switching during the approach of the reward port across all animals (n=5 rats), normalized per animal by subtracting off the average baseline proportion on Stay trials. Error bars are 95% CIs on the mean. Pre- and post-Switch linear regressions overlaid. Upper right: observed slope of pre-Switch regression (black) is significantly different from 0 ($p<0.002$) based on 1000 shuffles of the underlying data (grey). Lower right: observed slope of post-Switch regression (black) is significantly different from 0 ($p<0.002$) based on 1000 shuffles of the underlying data (grey). Both slopes were individually significant in all five animals (Figure 2.7). **(E)** Proportion of all non-local activity that represents paths consistent with Switching during the approach of the reward port in each animal on Stay trials and Switch trials. Error bars are 95% CIs on the mean. Switch trial distributions are significantly greater than in Stay trials for all animals (p=5.6e-264, 1.1e-155, 3.2e-300, 1.0e-100, 1.1e-277, Wilcoxon rank sum test, for n=292, 344, 525, 424, 443 Switch trials and n=3079, 3660, 6582, 4924, 3001 Stay trials). **(F)** Accuracy of logistic regressions predicting Stay or Switch choices based on proportion of all non-local activity that represents paths consistent with Switching during the approach of the reward port. Neural data from two trial types: previous trial (circle marker), or current trial (square marker). Error bars are 95% CIs on the proportion. Training data were balanced, so chance level was 0.5. Accuracy of model using previous trial neural data is significantly greater than chance in all animals (p=2.4e-58, 7.1e-32, 1.2e-56, 1.3e-40, 3.1e-48, Z test for proportions), and current trial model accuracy is greater than previous trial model accuracy (p=2.4e-58, 7.1e-32, 1.2e-56, 1.3e-40, 3.1e-48).

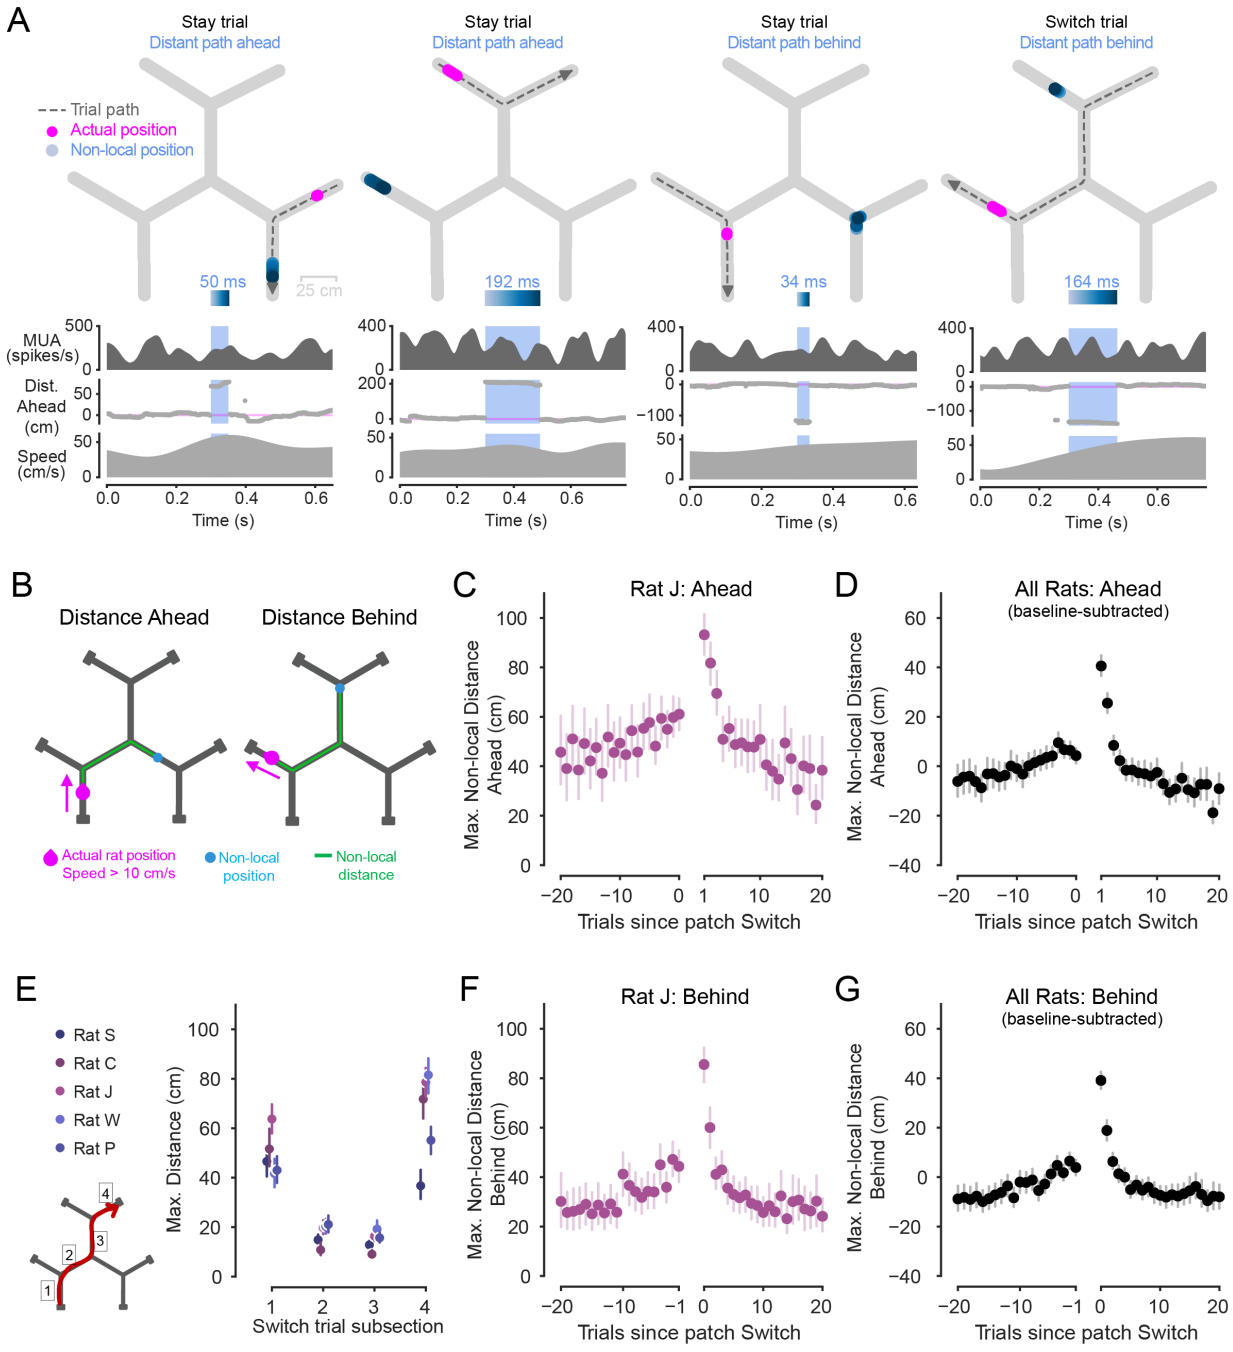

Figure 2.4 Non-local representations extend further early in patch experience
(A) Examples of distant non-local representations of paths ahead or behind, in track segments distinct from the animal's current segment. Top: Animal's actual position (magenta) and decoded position from hippocampal spiking (blue, shaded by time) during the blue-highlighted period (below) plotted on the track. Dashed grey line with arrow indicates animal's path and direction on the current trial. Highlighted non-local duration is labeled above blue time-shaded heatmap. Top-middle: multiunit spike rate from hippocampal tetrodes before, during, and after highlighted non-local representation. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) Bottom-middle: Distance from actual position to decoded position, with positive values corresponding to in front of the animal's heading, and negative numbers corresponding to behind the animal's heading. Bottom: Animal speed. Representations can correspond to distant locations within the current patch (left example) as well as in alternative, unoccupied patches (right three examples). **(B)** Schematic showing non-local distance, defined as distance in centimeters of the path along the track (green) from the animal's actual position (magenta) to the peak of the decoded posterior (blue). Example schematics correspond to representations of locations an example distance ahead of (left) or behind (right) the animal's actual position. **(C)** Maximum non-local distances represented on trials leading up to and following patch Switches for one animal. Data taken from during the period of each trial in which the animal approached the first choice point (as in B, left). Left and right x axes separated to reflect patch change. Error bars are 95% CIs on the mean. Individual animals shown in Figure 2.9. Exponential regression intercepts *a* and rate constants *b* are significantly larger magnitude for post-patch-change than pre-patch-change data in all individual animals except one (p_a = 1.3e-4, 4.6e-9, 1.1e-11, 3.1e-6, 0.8, p_b = 2.3e-4, 1.1e-8, 3.2e-10, 3.3e-6, 0.9, Z test). **(D)** Maximum non-local distance ahead represented on trials leading up to and following patch Switches, across all animals. Data taken from first segment of trials and normalized per animal by subtracting off the average across trials. Error bars are 95% CIs on the mean. **(E)** Maximum non-local distance represented for each animal while traversing each of the four segments of a Switch trial (segments 1-4 schematized in lower left). Error bars are 95% CIs on the mean. Non-local distance is significantly greater when each animal is in the first segment than the second (p=4.7e-52, 1.0e-55, 6.4e-59, 8.5e-24, 2.1e-42, Wilcoxon rank sum test), and is also significantly greater in the final segment than the third segment for each animal (p=2.3e-49, 2.4e-56, 3.3e-69, 1.5e-49, 3.0e-70). **(F)** Maximum non-local distance represented on trials leading up to and following patch Switches for one animal. Data taken from final segment of each trial, during approach of the reward port (as in B, right). Left and right x axes separated to reflect patch change. Error bars show 95% CIs on the mean. Individual animals shown in Figure 2.9. Exponential regression intercepts *a* and rate constants *b* are significantly larger magnitude for post-patch-change than pre-patch-change data in all five animals individually (p_a =8.8e-9, 2.9e-9, 1.4e-7, 3.2e-7, 9.2e-4, p_b =5.7e-6, 3.0e-7, 2.3e-5, 4.0e-5, 3.6e-2, Z test). **(G)** Maximum non-local distance behind represented on trials leading up to and following patch Switches, across all animals. Data taken from final segment of trials and normalized per animal by subtracting off the average across trials. Error bars are 95% CIs on the mean.

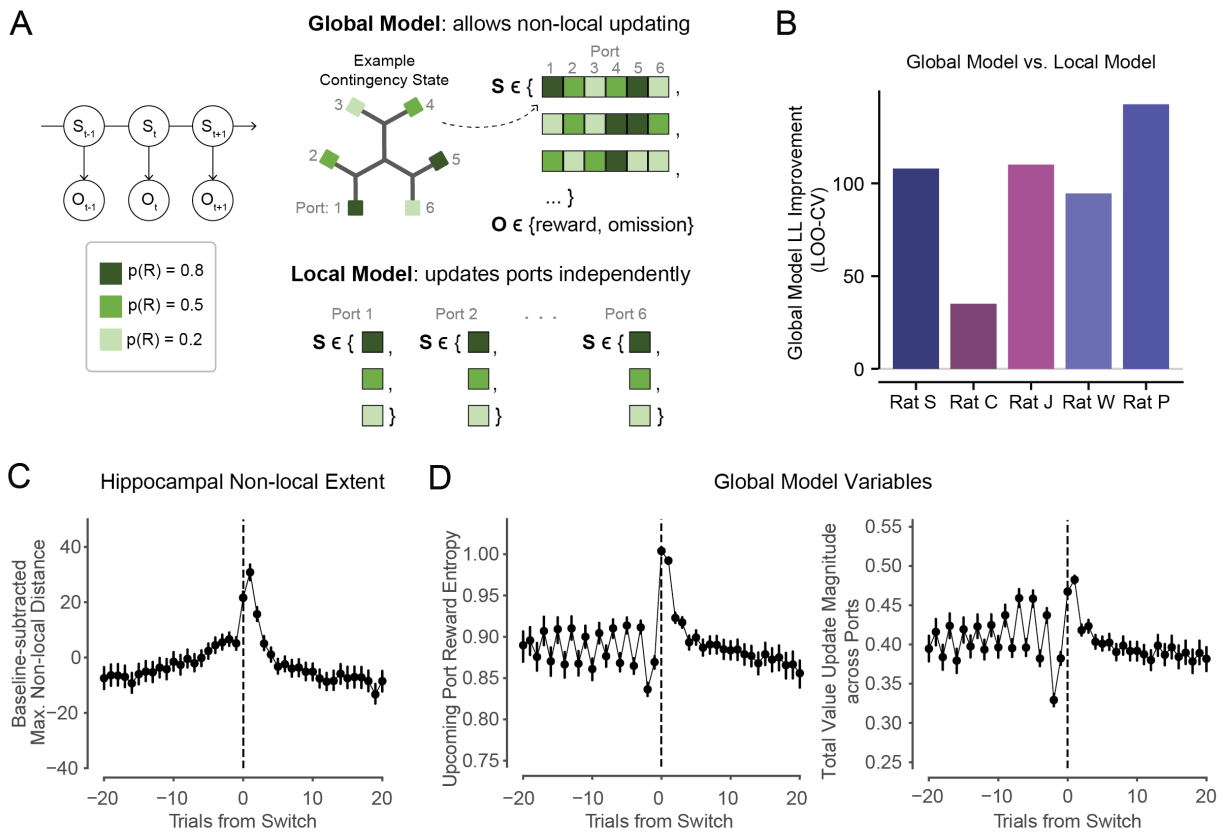

Figure 2.5 Behavioral modeling captures enhanced learning opportunities in early patch experience

(A) Left: schematic of Hidden Markov Model with hidden state S and observation O on each trial t. Colors correspond to potential nominal reward probabilities. Upper right: schematic of Global Model, showing example contingency states S with specific values for each port, and the two potential observations O of reward outcomes on each trial at the chosen port. Lower right: Example states S in Local Model, in which each port's value is estimated independently. The Global and Local models had matched reward distributions, though here we visualize unique example states. **(B)** Leave-one-out cross-validated log-likelihood improvement (higher is better) of Global Model versus Local Model. Global Model better fit behavior in all animals (p=.0035, .0285, .0006, .0013, .0104, t-test on cross-validated log-likelihoods per day). **(C)** All animals' normalized maximum hippocampal non-local distance represented on trials leading up to and following Switch trials, for non-local representations both ahead and behind (as in Figure 2.4). Trial-level quantification, rather than sub-trial-level quantification, shown here to match trial-level resolution of behavioral model. Error bars are 95% CIs on the mean. **(D)** Global Model variables related to learning opportunities in new patch and associated value updating across the maze are enhanced upon patch Switching. Left: Entropy over reward states at upcoming reward port shows asymmetrical pattern around Switch trials, becoming elevated on and after Switch trials, and decreasing across trials thereafter. Patterns were consistent across animals, shown individually in Figure 2.10. Right: Absolute value update resulting from each trial's reward outcome, summed across all ports in maze. Degree of value updating is elevated upon patch Switching, (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) and decreases across trials the longer the animal Stays within a patch. Patterns are similar for updates both within the current patch and across unoccupied patches (shown individually in Figure 2.10). Error bars are 95% CIs on the mean. Patterns were consistent across animals, shown individually in Figure 2.10.

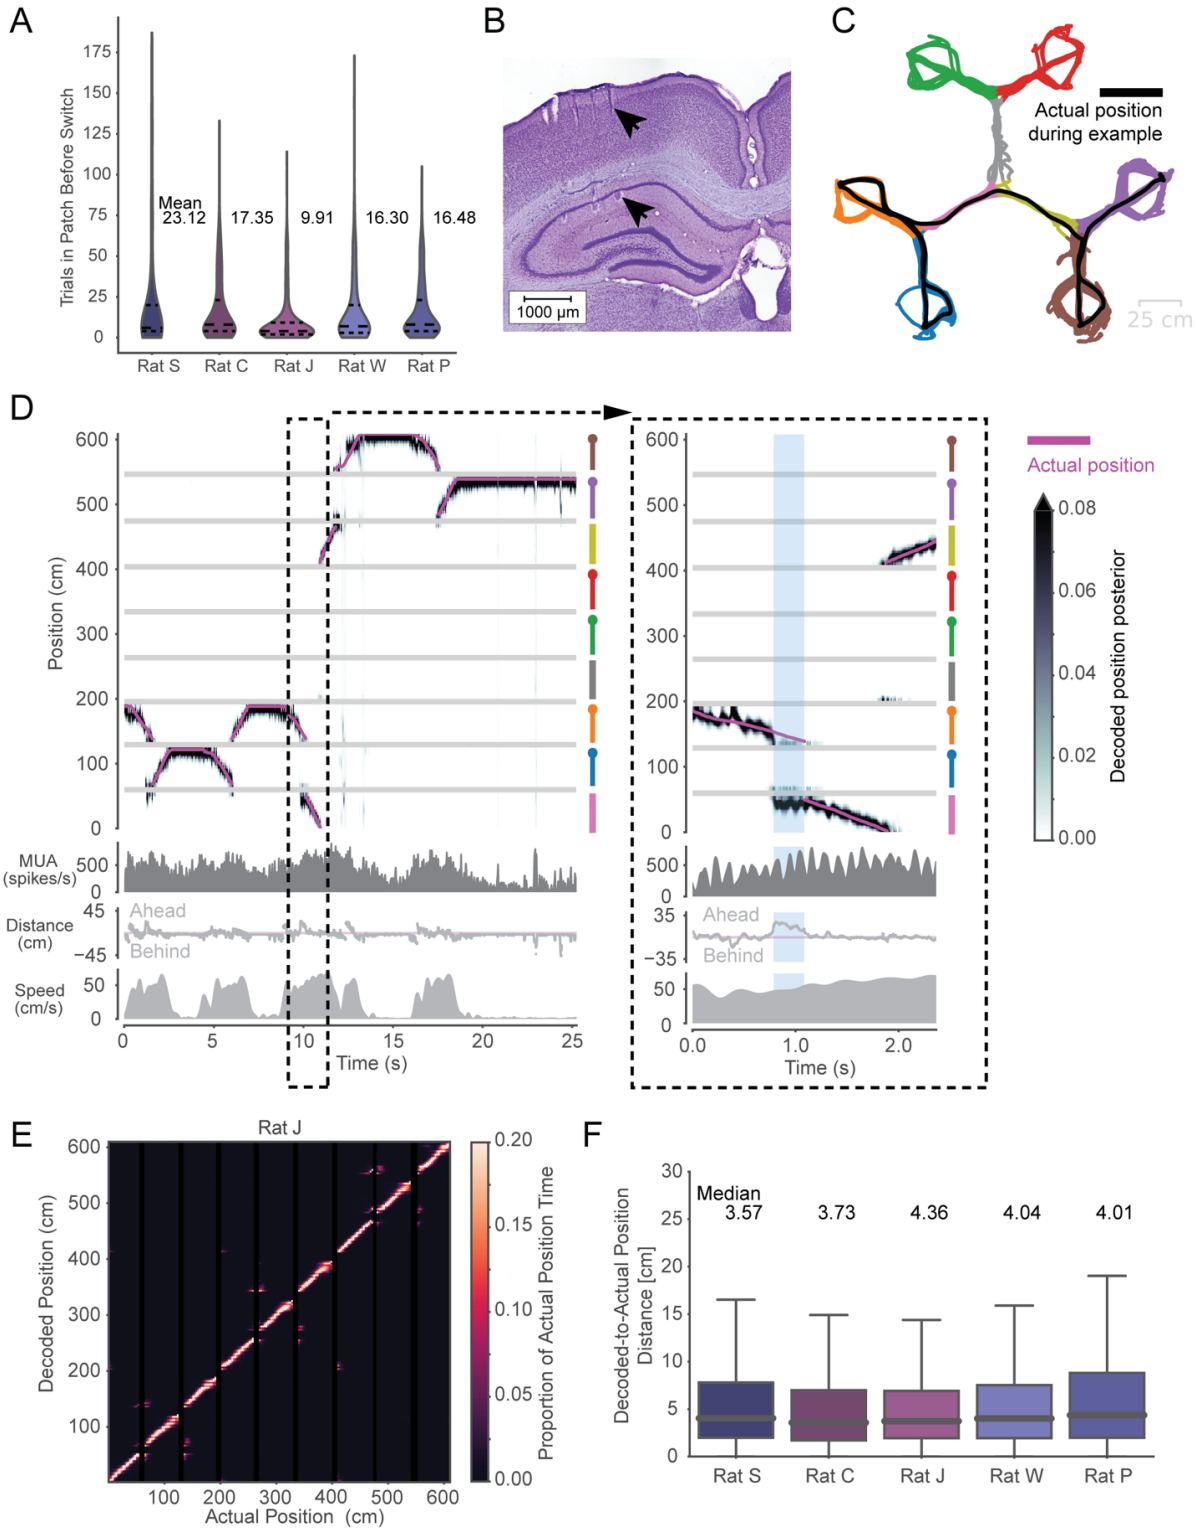

Supplemental Figure 2.6 Decoding animal position from Hippocampal spiking during patch foraging

(A) Violin plots showing distribution of number of trials spent in a patch before Switching, with horizontal dashed line (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) at median. Upper and lower dashed lines indicate quartiles. Mean trial bout durations are labeled per animal. **(B)** Nissl-labeled coronal brain tissue section representative example showing targeting of tetrodes to dCA1 of the hippocampus. Black arrows point to a tetrode track in dorsal cortex above and a lesion from the tetrode tip in the pyramidal cell layer below. **(C)** Animal head position over the course of one example session, with position data colored by the currently-occupied track segment. Black line highlights the animal's trajectory throughout the same ~25s period shown in D. Colored track segments correspond to the colored track segments in D. **(D)** Decoded position from hippocampal population spiking tracks the animal's actual position at a behavioral timescale and can sweep ahead or behind the animal to represent non-local positions at a sub-second timescale. Dashed lines highlight a small period that is enlarged at right. This period shows an example non-local representation (blue vertical bars) where the decoded position is in a segment distinct from the animal's actual position. Top: Actual head position in 1D (linearized) is shown in pink and decoded position posterior is shown in greyscale. Track segments are aligned on the right y-axis, and segment colors correspond to segments in C. Circles at the end of colored segment lines represent reward port locations. Note that actual and decoded position are stationary at reward port positions, as animal pokes into reward port. Horizontal grey lines correspond to 15 cm gaps introduced between track segments in linear position space. Top-middle: Multiunit spike rate across hippocampal tetrodes. Note fluctuations at roughly 8 Hz theta frequency, as expected during running. Bottom-middle: Distance of decoded position from animal's actual position, which can be either ahead (positive values), at (zero cm), or behind (negative values) the actual position. Bottom: Animal head speed. **(E)** Confusion matrix showing actual position and decoded position for one rat. Decoded position largely tracks animal's actual position across all track segments. Small amounts of off-diagonal density tend to occur at intersections of track segments, corresponding to adjacent positions in 2D space, as expected. **(F)** Boxplots of distribution of distance between decoded and actual position across valid decoded run times for each animal. Boxes show quartiles, whiskers correspond to the data range, and horizontal lines indicate medians, which are also labeled above.

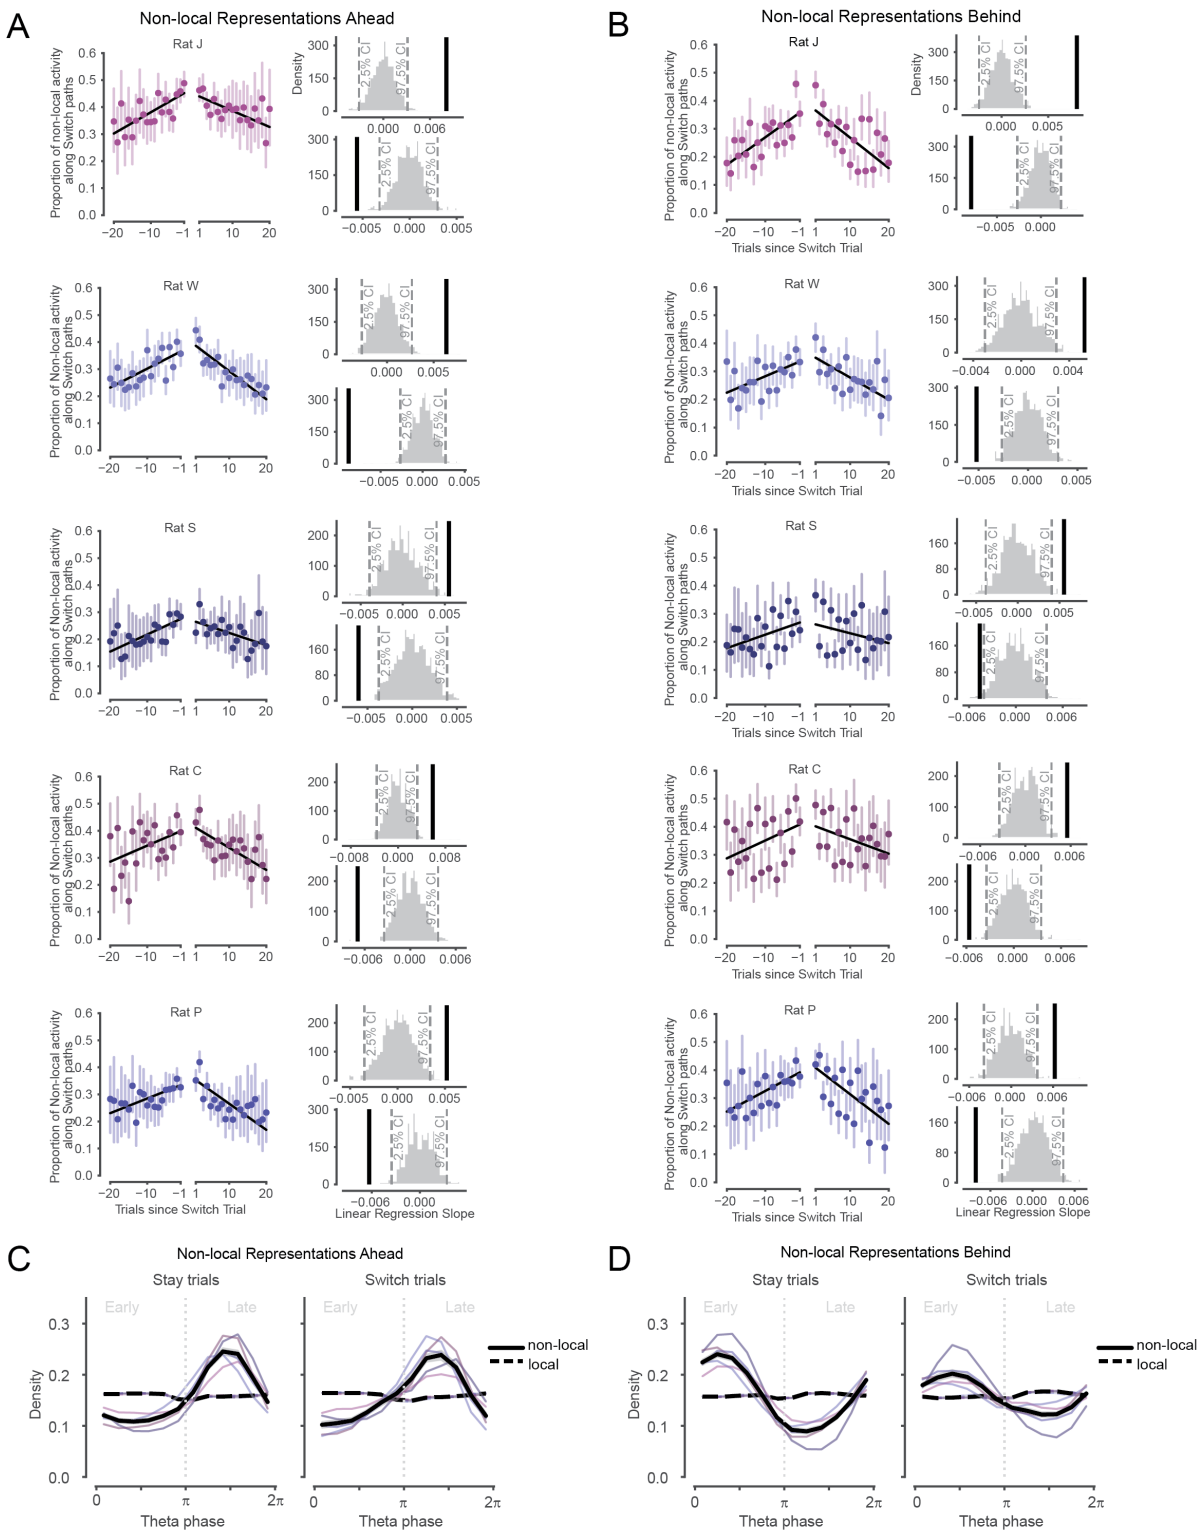

Supplemental Figure 2.7 Non-local representations of alternative paths Ahead and Behind are enriched across trials before and after patch Switching

(A) Proportion of all non-local activity that represents paths consistent with Switching on Stay trials before and after Switch trial (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) for each animal. Data correspond to Stay trials when animals were located in the first track segment and approaching the choice point, as in Figure 2.2. Error bars are 95% CIs on the mean. Pre- and post-Switch linear regressions overlaid in black. All slopes are significantly different than 0 (p_{pre}=0.002, 0.002, .0.014, 0.002, 0.004, p_{post}=0.002, 0.002, 0.008, 0.006, 0.004). Upper right plot: slope of pre-switch linear regression (black) is greater than the 97.5% CI on the slopes from 1000 shuffles of the underlying data (grey). Lower right plot: slope of post-switch linear regression (black) is less than the 2.5% CI on the slopes from 1000 shuffles of the underlying data (grey). **(B)** Same as A, but for non-local activity that represents paths consistent with Switching as the animal traverses the final segment of each trial and approaches the reward port, as in Figure 2.3. All slopes are significantly different than 0 (p_{pre}=0.002, 0.002, 0.004, 0.01, 0.008, p_{post}=0.002, 0.032, 0.004, 0.002, 0.002). **(C)** Non-local representations of paths ahead (solid lines) are concentrated in late phases of the theta rhythm, compared to local representations corresponding to the current track segment (dashed lines), on both Stay trials (left) and Switch trials (right). All animal data in black with 95% CI on the mean in grey band, and individual animal data colored per A. Early and late phases are separated with a vertical dotted grey line, and labeled in grey text above. Local and non-local distributions are significantly different in each animal for Stay trials (p = 4.4e-64, 1.0e-16, 2.5e-175,8.2e-129, 2.9e-18, Kuiper test) and Switch trials (p = 1.2e-7, 1.4e-11, 3.3e-12, 3.8-19, 9.3e-14). **(D)** Same as C, but for non-local representations of paths behind (solid lines) and local representations corresponding to the current track segment (dashed lines). Non-local paths behind are concentrated in early phases of theta. Local and non-local distributions are significantly different in each animal for Stay trials (p = 2.2e-7, 1.3e-52, 2.9e-202, 3.5e-23, 1.0e-14, Kuiper test) and Switch trials (p = 8.7e-4, 0.040, 1.3e-16, 0.015, 0.016).

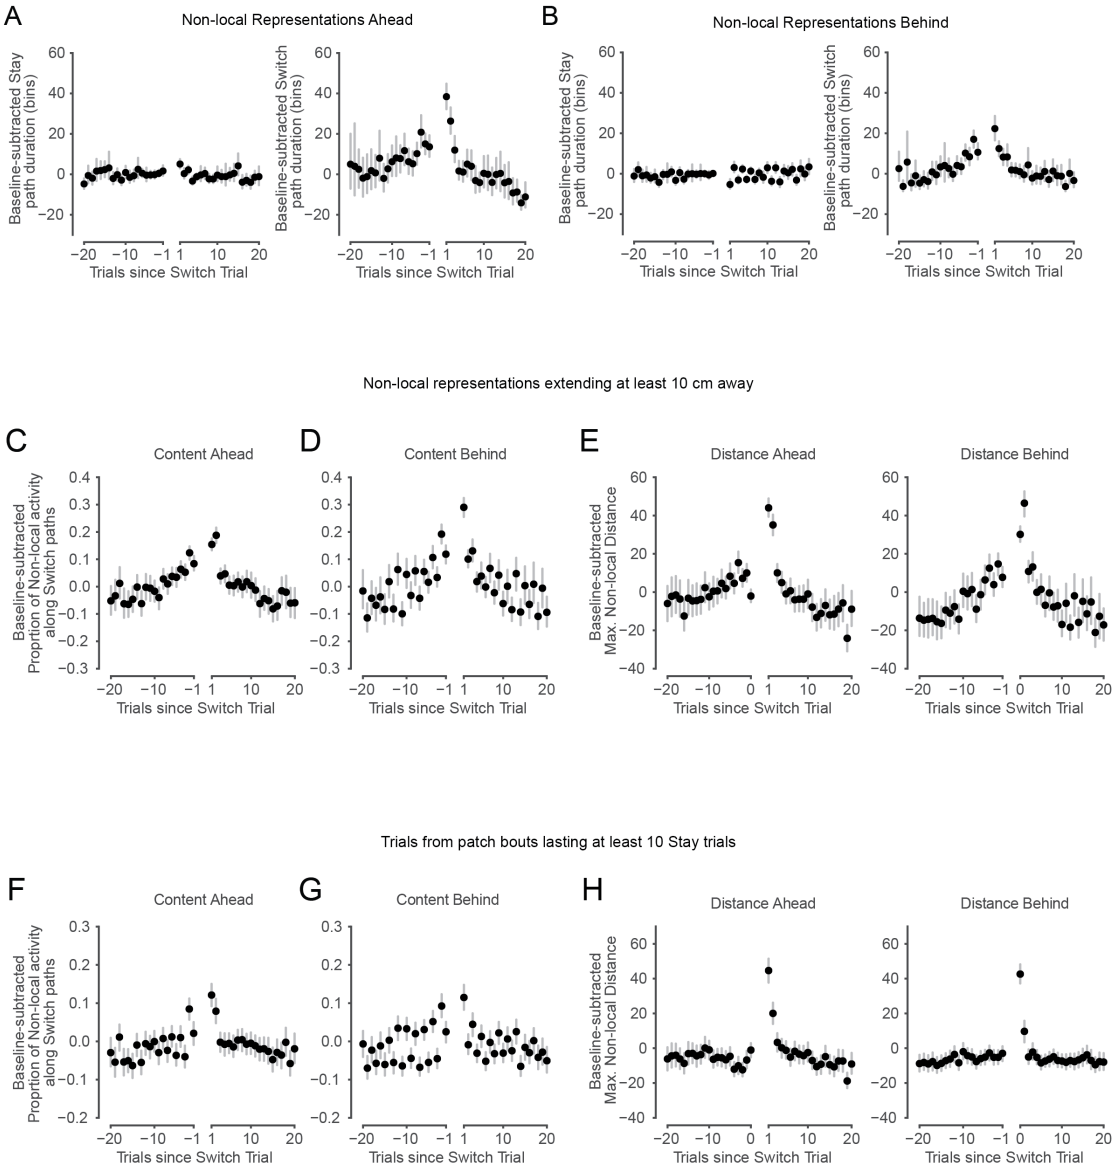

Supplemental Figure 2.8 Non-local representations Ahead and Behind are flexibly engaged around Switch trials

(A) Baseline-subtracted Stay path duration (left) and Switch path duration (right) across all animals. Durations are in 2 ms bins and quantify the number of bins where non-local representations were present during the approach of the first choice point across trials before and after Switch trials. Error bars are 95% CIs on the mean. Switch path durations were modulated more than were Stay path durations leading up to and following Switch trials. Related to Figure 2.2. Baseline durations per animal ranged 21.12-25.68 bins for Stay path representations and 14.61-28.03 bins for Switch path representations. **(B)** Same as A, but for non-local representations occurring during traversal of the final track segment and approach of the reward port across trials before and after Switch trials. Related to Figure 2.3. Baseline durations per animal ranged 13.99 – 26.65 bins for Stay path representations and 5.7-14.17 bins for Switch path representations. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) **(C)** Baseline-subtracted proportion of non-local representations along Switch paths, across all animals. These representations were expressed during the approach of the first choice point across trials before and after Switch trials. Only non-local representations at least 10 cm from the animal are included. Error bars are 95% CIs on the mean. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Figure 2.7 ($slope_{pre}$=0.0073, p=0.002; $slope_{post}$=-0.0094, p=0.002). The similarity to Figure 2.2 demonstrates that our results are not driven by representations very close to the animal. **(D)** Same as C, but for non-local representations occurring during traversal of the final track segment and approach of the reward port. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Figure 2.7 ($slope_{pre}$=0.0099, p=0.002; $slope_{post}$=-0.011, p=0.002). The similarity to Figure 2.3 demonstrates that our results are not driven by representations very close to the animal. **(E)** Baseline-subtracted maximum non-local distance occurring during approach of the first choice point (left) and traversal of the final track segment and approach of the reward port (right), across all animals. Only non-local representations at least 10 cm from the animal are included. Error bars are 95% CIs on the mean. Distances are especially elevated upon Switching patches and for a couple trials thereafter, and then decrease toward baseline levels. The similarity to Figure 2.4 demonstrates that our results are not driven by representations very close to the animal. **(F)** Baseline-subtracted proportion of non-local representations along Switch paths, across all animals. Data correspond to the approach of the first choice point across trials before and after Switch trials. Here, non-local representations are only included from trials in bouts within a patch lasting at least 10 Stay trials. Error bars are 95% CIs on the mean. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Figure 2.7 ($slope_{pre}$=0.0033, p=0.002; $slope_{post}$=-0.0047, p=0.002). The approximately symmetrical pattern around the Switch is similar to that seen in Figure 2.2, indicating that the results are not only driven by periods where animals Stayed in a patch for a small number of trials. **(G)** Same as F, but for non-local representations occurring during traversal of the final track segment and approach of the reward port across trials before and after Switch trials. Error bars are 95% CIs on the mean. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Figure 2.7 ($slope_{pre}$=0.0034, p=0.002; $slope_{post}$=-0.0036, p=0.002). The approximately symmetrical pattern around the Switch is similar to that seen in Figure 2.3, indicating that the results are not only driven by periods where animals Stayed in a patch for a small number of trials. **(H)** Baseline-subtracted maximum non-local distance during the approach of the first choice point (left) and traversal of the final track segment and approach of the reward port (right) across all animals. Here, non-local representations are only included from trials in bouts within a patch lasting at least 10 Stay trials. Distances are asymmetric around Switch trials, and are especially enhanced upon Switching patches, and then decrease across trials after the Switch. Related to Figure 2.4.

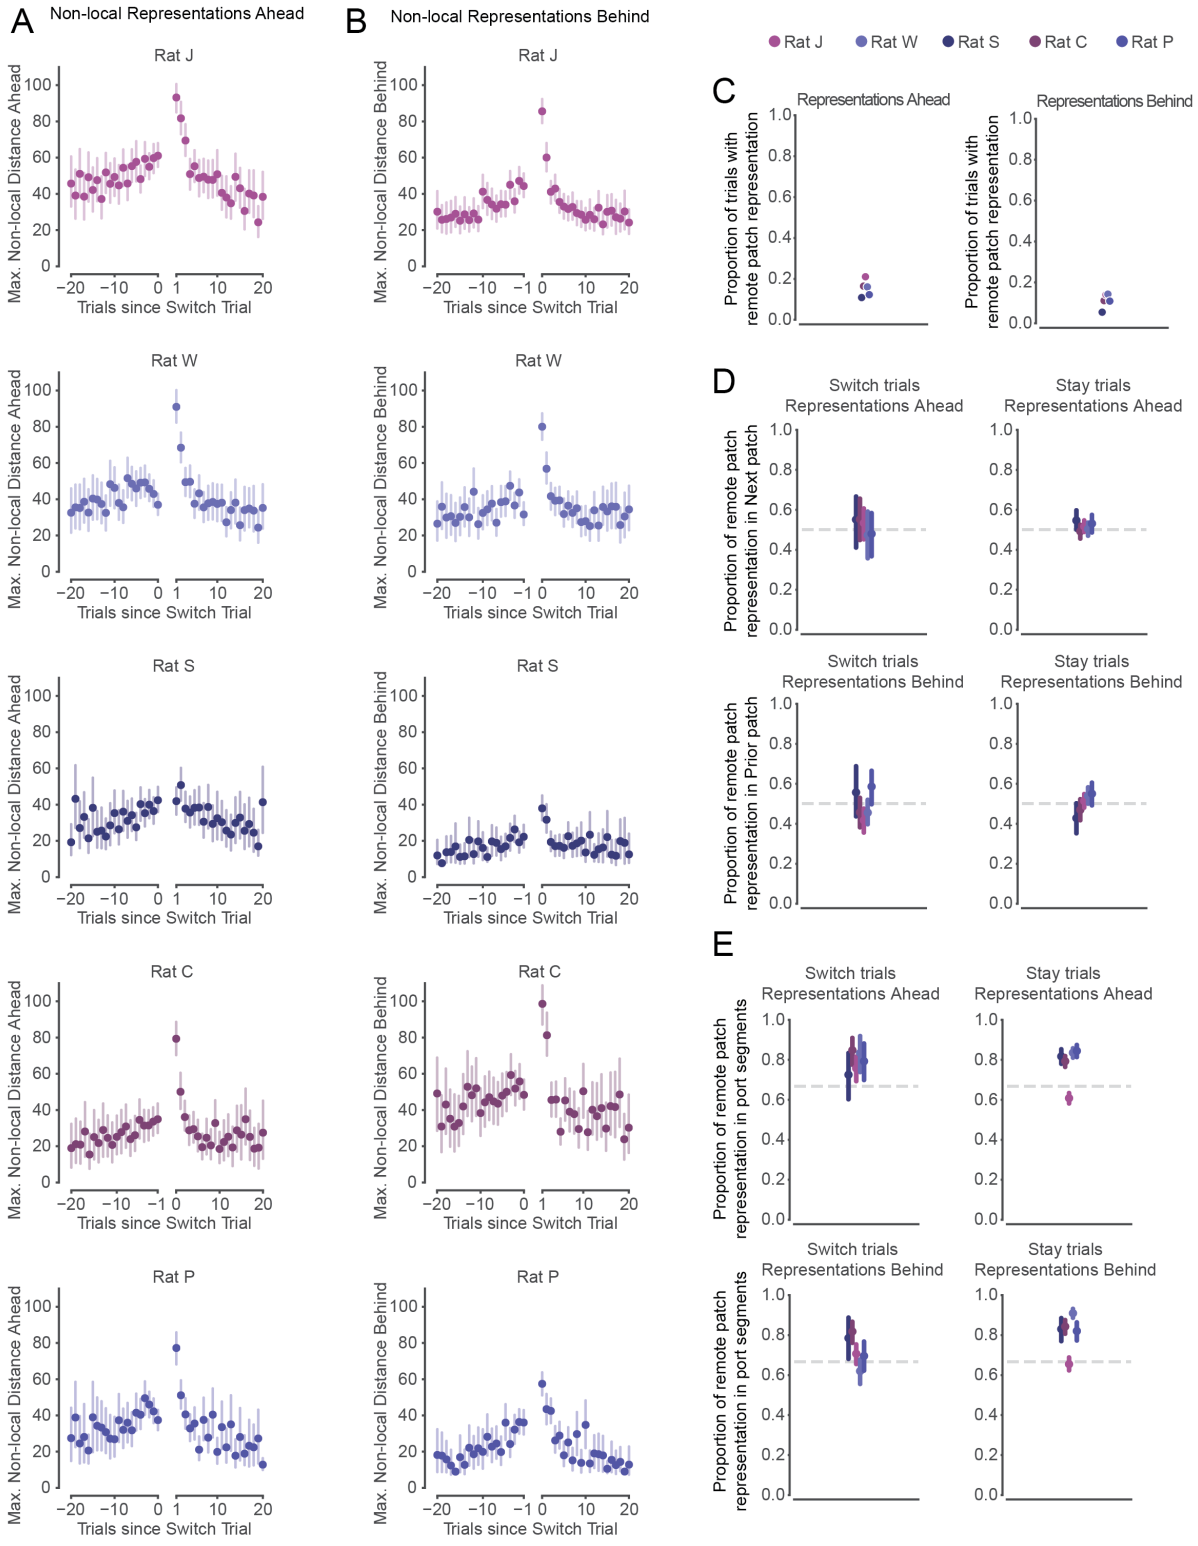

Supplemental Figure 2.9 Non-local representations of distant locations
(A) Maximum non-local distance represented on trials leading up to and following patch Switches for each animal. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) Data are from the period of each trial in which the animal approached the first choice point. Error bars are 95% CIs on the mean. Related to Figure 2.4. **(B)** Maximum non-local distance represented on trials leading up to and following patch Switches, during the period of each trial in which the animal traversed the final track segment towards the reward port, for each animal. Error bars are 95% CIs on the mean. Related to Figure 2.4. **(C)** Proportion of trials with non-local representations corresponding to remote, unoccupied patches, out of all trials with any non-local representations, for each animal. Left: calculated for non-local representations occurring as animals approached first choice point. Right: calculated for non-local representations occurring as animals traversed final track segment and approached reward port. In both cases, ~10-20% of trials contain representations corresponding to locations in remote patches. **(D)** Non-local representations corresponding to locations in remote patches are not biased to represent subsequently chosen or immediately previous patches. Top row: proportion of remote patch representations occurring during approach of first choice point on Switch trials (left) and Stay trials (right) that correspond to subsequently chosen (next) patch. Bottom row: proportion of remote patch representations occurring during traversal of final track segment and approach of reward port on Switch trials (left) and Stay trials (right) that correspond to most recently chosen (prior) patch. Error bars are 95% CIs on the mean. Grey dashed lines correspond to an equal (50%) representation of the chosen and non-chosen patches. **(E)** Non-local representations in remote patches are approximately equally distributed across track segments. Top row: proportion of remote patch representations occurring during approach of first choice point on Switch trials (left) and Stay trials (right) that correspond to track segments containing reward ports. Bottom row: proportion of remote patch representations occurring during traversal of final track segment and approach of reward port on Switch trials (left) and Stay trials (right) that correspond to track segments containing reward ports. Error bars are 95% CIs on the mean. Grey dashed lines correspond to a chance (66.67%) representation of the 4 segments containing reward ports out of 6 total track segments in remote patches.

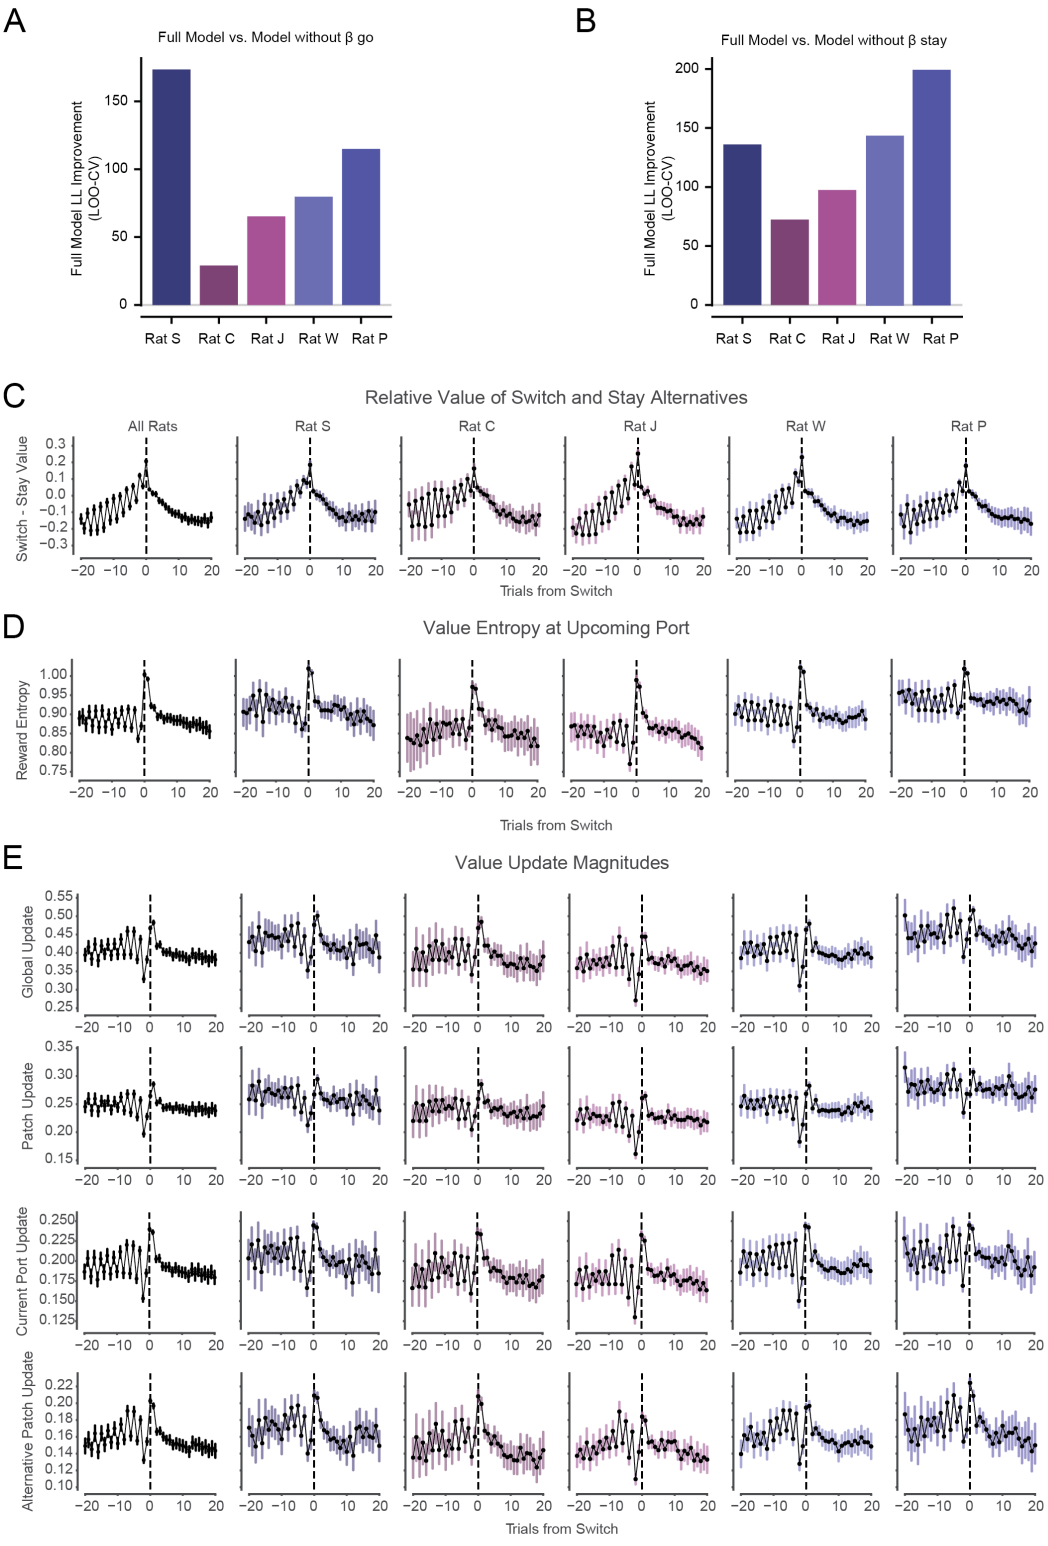

Supplemental Figure 2.10 Behavioral model reward sensitivity and variables for individual animals

(A) Leave-one-out cross-validated log-likelihood improvement (higher is better) of Full Model versus an alternative model. (Figure caption continued on the next page.)

(Figure caption continued from the previous page.) In alternative model, β_{go} was fixed to 0 (such that choice predictions ignored the estimate of the current patch value). Full Model significantly better fit behavior in all animals but one (p=1e-4, 0.113, 0.002, 0.016, 0.001, t-test on cross-validated log-likelihoods per day). **(B)** Leave-one-out cross-validated log-likelihood improvement (higher is better) of Full Model versus a model where β_{stay} was fixed to 0 (such that choice predictions ignored the estimate of the alternative patch values). Full Model significantly better fit behavior in all animals (p=0.002, 0.002, 0.002, 0.0002, 0.003, t-test on cross-validated log-likelihoods per day). **(C)** Relative value of Switching and Staying increases leading up to Switch trials and decreases afterwards. The value of Staying is the behavioral model-estimated value of the upcoming port within the current patch. The value of Switching is the behavioral model-estimated value of the greater value patch, where patch value is the average of the two ports within. Across-trial dynamics are roughly symmetric around the Switch. Error bars are 95% CIs on the mean. Related to Figure 2.1. **(D)** Entropy over behavioral model-estimated value states in the upcoming (chosen) port on each trial, increasing on and after patch Switches. Across-trial dynamics are asymmetric around the Switch. Error bars are 95% CIs on the mean. Related to Figure 2.5. **(E)** Value updates on each trial increase leading up to and especially on and after Switch trials, then decay across trials back to baseline. This pattern is observed across ports in the maze, not only at the currently visited port, indicating a period of enhanced learning by value updating upon patch Switching. Value updates are calculated as the absolute magnitude of the change in value across all ports or a subset of ports as follows. Top: Global value updates across all ports in the environment, as in Figure 2.5. Top-middle: Value updates at the current port. Bottom-middle: Value updates at ports in the current patch. Bottom: Value updates at ports in unoccupied, alternative patches. Error bars are 95% CIs on the mean. Related to Figure 2.5.

CHAPTER 3: CONCLUSIONS AND IMPLICATIONS

Conclusions

This work has described neural activity patterns generated in the hippocampus that correspond to alternative “non-local” possibilities. While a long and fruitful tradition of neuroscientific research has aimed to relate neuronal firing patterns to current external stimuli[265] and immediate behavior[266], the brain also has the remarkable abilities to store experiences, construct new knowledge from limited experiences, and later apply these to behave adaptively at later times and in distant places. In Chapter 1, we synthesized prior evidence that the hippocampus essentially contributes to such functions through generative neural activity patterns that reflect alternatives to an animal’s current experience. In Chapter 2, we presented evidence that during active navigation hippocampal representations of a wide range of alternatives are flexibly generated in patterns that match current demands for both experience-guided decision making and learning from limited experience. Together, this work suggests that the hippocampus not only regularly generates representations that correspond to alternative possibilities, but that mechanisms exist in the brain that tailor these representations to meet an animal’s changing functional needs. Critically, this work leaves open several lines of investigation with the potential to reveal the neural mechanisms underlying these flexibly generated neural activity patterns, as well as their possible roles in cognition and behavior.

Generative Representations Across Brain States

A central finding of this work is that the hippocampus regularly represents different alternatives across trials of experience particularly during active locomotion. We focused on movement times during our foraging task because the natural world often requires – sometimes as a matter of survival – that animals make adaptive choices during active behavior. For instance, a mouse escaping from a looming predator must make a quick choice about which way to navigate towards shelter[267]. Or, in the case of learning, after running by a newly available route[38], an animal may later be able to return to use this route, having previously updated its internal model of the environment while on the move. That said, the hippocampus also generates representations of alternatives during rest[204], as in sharp wave ripple replay events that are observed during immobility, wakeful rest, and sleep[76]. Thus, a relevant question for future work is whether these replay events are engaged in similar or different patterns to those we observed during movement.

One influential hypothesis, that makes predictions about neural reactivation during rest in the spatial bandit task, posits that representations in the hippocampus during active behavior are later consolidated during rest[217]. In the context of our work, this two-stage model raises the possibility that the neural activity corresponding to specific alternatives represented during running in the spatial bandit task are later reactivated during rest, either for use in subsequent behavior, consolidation for longer-term storage, or both[62]. Reactivation of reward-related locations would be expected based previous work indicating that individual episodes are replayed following behavioral experiences[268,269], particularly in association with reward experiences[74,77,78].

Some reports, taking both experimental and computational approaches, further link replay with value coding and value learning, for memory maintenance and updating[53,84,119,120,125,223,270]. These lines of research suggest that in our foraging task, replay may be biased to represent locations in the environment in association with their values. Whether replay does so, and whether replay is also enhanced on trials with greater value updating, may further uncover its contribution to value-guided foraging and decision making. This is the subject of ongoing work. Additionally, hippocampal spatial sequence representations during running are thought to directly relate to later replay during rest[199,271]. Given that we observed dynamic non-local representations across individual trials of running throughout our foraging task, this provides an opportunity to test whether the content of hippocampal replay events either follows along with or systematically diverges from the content we observed on the same trials during navigation. Similarly, more- or less-rewarding[68] locations may be preferentially reactivated during post-behavior sleep, which may in turn predict later biases in foraging choices[272].

The Coordination of Generative Representations Across Brain Regions

Another finding from this work is that hippocampal activity corresponding to specific locations in the environment can change in conjunction with behaviorally relevant variables, as they evolve across trials of experience. Specifically, we found that the relative non-local content along different paths changed with the relative value of those paths, and that the extent of these non-local representations was greater during periods of enhanced value updating about locations across the environment. Our findings

indicate that mechanisms exist that regulate non-local representations, in terms of both content and extent.

While the neural mechanisms underlying these two forms of modulation of non-local representations are not known, existing literature points to several candidate processes that may be involved. An initial clue is that our findings suggest that such mechanisms have access to value signals. Neural activity directly related to reward value has indeed been reported in the hippocampus[228,239], raising the possibility that within-hippocampus value coding contributes to the content and extent of non-local representations. Evidence for neural value signals in other brain systems that communicate with the hippocampus, however, is more abundant. These systems include, but are not limited to, the frontal cortex, striatum, and dopaminergic system. Although a review of these systems in relation to hippocampal non-local activity is beyond the scope of this work, some initial points can be made in relation to our experiments.

Areas in frontal cortex are of particular interest for the potential neural mechanisms underlying the dynamic generation of non-local hippocampal representations. Neural activity in prefrontal cortex can represent expected value and is modulated by reward[273–278]. Prefrontal cortical activity is also modulated by hippocampal local field potential features associated with non-local representations, including the theta rhythm and sharp wave ripples[253,279–282]. Further, prefrontal cortical activity can predict nonlocal hippocampal firing during movement[70]. Importantly, because prefrontal cortex also receives input from the hippocampus[283–285], prefrontal cortex may also be a site of the later evaluation of paths represented by the hippocampus. And, given that

roles in generating and evaluating hippocampal representations of different spatial paths are not mutually exclusive, frontal cortex may, speculatively, subserve both. Along these lines, prior work showed a cortical-hippocampal-cortical loop between sensory cortical activity and non-local hippocampal activity[123].

We found that non-local hippocampal activity preferentially represents different paths and locations as the values of those locations evolve across trials of experience. Because our task design dissociates value from location by regularly changing the reward probabilities associated with different reward ports, monitoring of frontal cortical neuronal activity would enable us to test whether it concurrently reflects the values of the specific non-local paths represented by the hippocampus. The above reasons and others motivated us to simultaneously record from prefrontal cortical brain areas alongside the hippocampus in our experiments (and to develop the required technology to do so[256]). Therefore, guided by our computational model of behavior, ongoing work[286] aims to address how value, uncertainty, and other decision variables in prefrontal cortex are coordinated with non-local hippocampal representations during learning and decision making in our foraging task. Examining whether value- and value updating-related patterns in the hippocampus are concurrently observed in or coordinated with dopamine release in frontal cortex[247], as well as in nucleus accumbens[185], may further elucidate the processes that underlie the dynamic generation of non-local hippocampal representations reported in Chapter 2.

Further motivation to investigate non-local representations in coordination with neural activity across brain areas in our experiments comes from critical evidence that non-local activity in the hippocampus is internally generated, rather than primarily driven

by immediate external stimuli. As described in Chapter 1, non-local representations are aligned to the theta rhythm[55,287], which is itself internally generated[89–92]. Sharp wave ripples, associated with replay, are also thought to be internally generated[76]. Additionally, while non-local representations associated with both theta and sharp wave ripples have been correlated with immediately upcoming behavioral choices[98,100–102], there are also a variety of reports in which they are not[69,84,99,115]. Prior work further suggests that these internally generated representations can be more or less coupled to ongoing behavior depending on the memory demands of the task[254]. Together, these findings provide evidence that generative representations can be expressed without immediate sensory or behavioral correlates. Consistent with this, we found that two potential path options can be equally likely to be represented through non-local activity (and therefore would not discriminate immediately upcoming choice[69]), while at other times in behavior the previously or subsequently taken path dominated non-local representations. Thus, because internal representations can be dissociated from immediate behavior, one strategy to understand the computations performed by hippocampal representations is to monitor the activity of other brain areas or neurobiological processes simultaneously. In our work this approach could, for example, reveal an increase in the value represented in one brain area upon hippocampal reactivation of a just-rewarded location. Studying representational content in one area as it relates to that of another area has become possible with technology that enables recording of large populations across brain regions[256,288,289], and has already been applied to reveal coordinated spatial content across the hippocampus and other brain areas[101,115–117,124]. In our data, this approach may help to address how the brain

coordinates spatial and value-related information, and how this coordination may develop through experience. Indeed, the need for understanding internally generated neural activity patterns in the hippocampus and beyond has also been converged upon not only from a representational view of the brain, but by nonrepresentational views alike[290].

The Role of Generative Activity in Experience-Guided Decision Making

Our results describe relationships between hippocampal population activity and value-guided behavior during navigation in our foraging task, but our conclusions are limited by an absence of causal evidence for these relationships. Previous studies showed that the rodent hippocampus is important for memory-guided decision making and inference[45,46,192], as well as that hippocampal theta rhythmic activity plays a role in specific aspects of memory encoding and retreival[198,291,292]. Recent work also identified that the sequential nature of hippocampal firing along the theta rhythm, which is thought to give rise to non-local representations, is also critical for behavior in a flexible memory-guided navigation task[199]. To address whether the coordinated firing of place cells along the theta rhythm specifically during locomotion is required for learning a memory-guided task, we developed a closed-loop manipulation based on the phase of the theta rhythm to disrupt hippocampal population coordination during movement via the medial septum[293]. Preliminary results[293] suggest a learning deficit in a memory-guided navigation task. Experiments such as these complement a growing literature on closed-loop manipulations of sharp wave ripples during rest, which can be interrupted to impair, and lengthened to improve, learning experience-guided behaviors[157,158].

Together, these observations suggest a direct role of theta- and sharp wave ripple-associated neural activity in encoding experience for later use. Whether the specific representational content of these non-local events on individual trials is causally related to behavior, however, is not well understood. Closed-loop experiments that involve decoding and manipulating hippocampal spatial representations in real-time[294–296], and in complex and dynamic environments[297], will continue to advance our understanding of whether and how the different kinds of non-local events we observed are required for flexible experience-guided decision making.

Generative Representations in Natural Behavior

Our work also highlights the importance of studying brain activity during naturalistic behavior. Significant prior work has often investigated hippocampal non-local activity in simpler environments, with well learned alternation rules, or with changes that must be learned directly rather than inferred through learned structure across trials[298]. Despite important progress using these approaches, animals in the natural world live in environments that are more vast, dynamic, and uncertain than those in traditional rodent spatial memory tasks[186,190]. To begin to account for this, our spatial bandit task design took inspiration from complex real-world foraging[211], borrowed the design of controlled blocks of stochastic rewards from decision-making experiments[299], and incorporated branching maze structures from spatial learning and memory experiments[298]. While the task is by no means comparable to wild foraging environments, it enabled us to ask how spatial possibilities are internally selected among to reflect just the most relevant non-local possibilities across many locations near and far, across many potential choice

points, and throughout regularly changing reward experiences that encouraged constant adaptation. This in turn enabled us to discover several patterns of modulation of non-local representations, including representations extending along unchosen alternative paths ahead and behind the animal and representations that occasionally jumped unexpectedly far across the environment into alternative patches. Perhaps one of the more surprising results we observed was the significant increase in non-local extent as animals traversed the final track segments of switch trials. This indicates that factors above and beyond distances to choice points[300], goals[98], and novelty can modulate non-local representations during locomotion. This finding was neither predicted by prior work, nor would have been found in a similar probabilistically rewarded yet simpler Y-maze. Thus, by studying the brain in more complex, naturalistic, and ethologically relevant settings, we can make unexpected observations about how the brain rises to the challenges of the real world—challenges it has evolved to meet. Indeed, the field of hippocampal spatial memory was launched in part by the observation of spatially specific firing of hippocampal neurons in freely foraging rodents[43,47]. Now, equipped with larger scale neural and behavioral recording and analysis technologies[301–304], future studies in naturally behaving animals and humans alike are positioned to advance our understanding of the circuit mechanisms underlying generative representations in the hippocampus and their roles in cognition and adaptive behavior.

Summary

Understanding how the hippocampus flexibly generates representations of alternative possibilities contributes to our knowledge of how prior experience may be used to

update and adaptively apply an internal model of the world to make advantageous choices in complex and dynamic environments. In this study we described a neurophysiological basis for the systematic generation of a range of alternative possibilities that appropriately meet multiple systematically changing cognitive demands throughout experience-based decision making and learning. Building from this work and others, future investigations of the mechanisms that give rise to and coordinate with these neural activity patterns have the opportunity to dissect how the brain's memory and decision systems may interact to support cognition and experience-guided adaptive behavior.

REFERENCES

1. Pezzulo, G., Kemere, C., and van der Meer, M.A.A. (2017). Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. Ann N Y Acad Sci *1396*, 144–165. https://doi.org/10.1111/nyas.13329.

2. Corballis, M.C. (2019). The Origins of Generativity. In Handbook of Cognitive Archaeology.

3. McNamee, D.C., Stachenfeld, K.L., Botvinick, M.M., and Gershman, S.J. (2021). Flexible modulation of sequence generation in the entorhinal-hippocampal system. Nat Neurosci *24*, 851–862. https://doi.org/10.1038/s41593-021-00831-7.

4. Carey, S., Leahy, B., Redshaw, J., and Suddendorf, T. (2020). Could It Be So? The Cognitive Science of Possibility. Trends Cogn Sci *24*, 3–4. https://doi.org/10.1016/j.tics.2019.11.007.

5. Cisek, P. (2012). Making decisions through a distributed consensus. Curr Opin Neurobiol *22*, 927–936. https://doi.org/10.1016/j.conb.2012.05.007.

6. Talland, G.A. (1965). Deranged Memory: A Psychonomic Study of the Amnesic Syndrome (Academic Press).

7. Tulving, E. (1985). Memory and consciousness. Canadian Psychology, 1–12.

8. Korsakoff, S.S. (1996). Medico-psychological study of a memory disorder. Conscious Cogn *5*, 2–21.

9. Klein, S.B. (2002). Memory and temporal experience: the effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. Soc. Cogn. *20*, 353–379. https://doi.org/10.1521/soco.20.5.353.21125.

10. Rosenbaum, R.S., Köhler, S., Schacter, D.L., Moscovitch, M., Westmacott, R., Black, S.E., Gao, F., and Tulving, E. (2005). The case of K.C.: contributions of a memory-impaired person to memory theory. Neuropsychologia *43*, 898–1021. https://doi.org/10.1016/j.neuropsychologia.2004.10.007.

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

12. Corkin, S. (2002). What's new with the amnesic patient H.M.? Nat Rev Neurosci *3*, 153–160. https://doi.org/10.1038/nrn726.

13. Hassabis, D., and Maguire, E.A. (2007). Deconstructing episodic memory with construction. Trends Cogn. Sci. *11*, 299–306. https://doi.org/10.1016/j.tics.2007.05.001.

14. Hassabis, D., Kumaran, D., Vann, S.D., and Maguire, E.A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. Proc Natl.Acad.Sci.U.S.A *104*, 1726–1731.

15. Suddendorf, T., and Corballis, M.C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? Behav Brain Sci *30*, 299–313. https://doi.org/10.1017/S0140525X07001975.

16. Rosenbaum, R.S., Gilboa, A., Levine, B., Winocur, G., and Moscovitch, M. (2009). Amnesia as an impairment of detail generation and binding: evidence from personal, fictional, and semantic narratives in K.C. Neuropsychologia *47*, 2181–2187.

17. Buckner, R.L. (2010). The role of the hippocampus in prediction and imagination. Annual review of psychology *61*, 27-48-C1-8. https://doi.org/10.1146/annurev.psych.60.110707.163508.

18. Mullally, S.L., and Maguire, E.A. (2014). Memory, Imagination, and Predicting the Future: A Common Brain Mechanism? Neuroscientist *20*, 220–234. https://doi.org/10.1177/1073858413495091.

19. Buckner, R.L., and Carroll, D.C. (2007). Self-projection and the brain. Trends Cogn Sci *11*, 49–57. https://doi.org/10.1016/j.tics.2006.11.004.

20. Buckner, R.L., Andrews-Hanna, J.R., and Schacter, D.L. (2008). The brain's default network: anatomy, function, and relevance to disease. Ann N Y Acad Sci *1124*, 1–38. https://doi.org/10.1196/annals.1440.011.

21. Spreng, R.N., Mar, R.A., and Kim, A.S. (2009). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. J Cogn Neurosci *21*, 489–510. https://doi.org/10.1162/jocn.2008.21029.

22. Addis, D.R., and Schacter, D.L. (2011). The hippocampus and imagining the future: where do we stand? Front Hum Neurosci *5*, 173. https://doi.org/10.3389/fnhum.2011.00173.

23. Addis, D.R. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. Neuropsychologia *45*, 1363–1377. https://doi.org/10.1016/j.neuropsychologia.2006.10.016.

24. Addis, D.R., Pan, L., Vu, M.A., Laiser, N., and Schacter, D.L. (2009). Constructive episodic simulation of the future and the past: distinct subsystems of a core brain network mediate imagining and remembering. Neuropsychologia *47*, 2222–2238. https://doi.org/10.1016/j.neuropsychologia.2008.10.026.

25. Andrews-Hanna, J.R., Reidler, J.S., Sepulcre, J., Poulin, R., and Buckner, R.L. (2010). Functional-anatomic fractionation of the brain's default network. Neuron *65*, 550–562. https://doi.org/10.1016/j.neuron.2010.02.005.

26. Weiler, J.A., Suchan, B., and Daum, I. (2010). When the future becomes the past: Differences in brain activation patterns for episodic memory and episodic future thinking. Behav Brain Res *212*, 196–203. https://doi.org/10.1016/j.bbr.2010.04.013.

27. Zeidman, P., Mullally, S.L., and Maguire, E.A. (2015). Constructing, Perceiving, and Maintaining Scenes: Hippocampal Activity and Connectivity. Cereb Cortex *25*, 3836–3855. https://doi.org/10.1093/cercor/bhu266.

28. Hassabis, D., and Maguire, E.A. (2009). The construction system of the brain. Philos Trans R Soc Lond B Biol Sci *364*, 1263–1271. https://doi.org/10.1098/rstb.2008.0296.

29. Maguire, E.A., and Mullally, S.L. (2013). The Hippocampus: A Manifesto for Change. J Exp Psychol Gen *142*, 1180–1189. https://doi.org/10.1037/a0033650.

30. Tulving, E. (1983). Elements of episodic memory (Clarendon Press ; Oxford University Press).

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

32. Burgess, N., Maguire, E.A., and O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. Neuron *35*, 625–641.

33. Mullally, S.L., Hassabis, D., and Maguire, E.A. (2012). Scene construction in amnesia: An fMRI study. Journal of Neuroscience *32*, 5646–5653.

34. Mullally, S.L., Vargha-Khadem, F., and Maguire, E.A. (2014). Scene construction in developmental amnesia: An fMRI study. Neuropsychologia *52*, 1–10. https://doi.org/10.1016/j.neuropsychologia.2013.11.001.

35. Lazareva, O., and Wasserman, E. (2017). Categories and Concepts in Animals. In Reference Module in Neuroscience and Biobehavioral Psychology https://doi.org/10.1016/B978-0-12-809324-5.21008-0.

36. Colbert-White, E., and Kaufman, A. (2019). Animal Cognition 101. In (Springer Publishing Company), pp. 1–30.

37. Dona Hiruni Samadi, G., and Chittka, L. (2020). Charles H. Turner, pioneer in animal cognition. Science *370*, 530–531. https://doi.org/10.1126/science.abd8754.

38. Tolman, E.C., Ritchie, B.F., and Kalish, D. (1946). Studies in spatial learning: Orientation and the short-cut. Journal of experimental psychology *36*, 13–24.

39. Tolman, E.C., and Gleitman, H. (1949). Studies in learning and motivation: I. Equal reinforcements in both end-boxes, followed by shock in one end-box. Journal of Experimental Psychology *39*, 810–819. https://doi.org/10.1037/h0062845.

40. Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol Rev *55*, 189–208.

41. Steiner, A.P., and Redish, A.D. (2014). Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. Nature neuroscience *17*, 995–1002. https://doi.org/10.1038/nn.3740.

42. Redish, A.D. (2016). Vicarious trial and error. Nat Rev Neurosci *17*, 147–159. https://doi.org/10.1038/nrn.2015.30.

43. O'Keefe, J., and Nadel, L. (1978). The hippocampus as a cognitive map (Oxford University Press).

44. Knierim, J.J. (2015). The hippocampus. Current Biology *25*, R1116–R1121.

45. Dusek, J.A., and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. Proc.Natl.Acad.Sci.U.S.A. *94*, 7109–7114.

46. Miller, K.J., Botvinick, M.M., and Brody, C.D. (2017). Dorsal hippocampus contributes to model-based planning. Nat Neurosci *20*, 1269–1276. https://doi.org/10.1038/nn.4613.

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

48. Moser, E.I., Kropff, E., and Moser, M.B. (2008). Place Cells, Grid Cells, and the Brain's Spatial Representation System. Annu.Rev.Neurosci. *%19;.*

49. Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. J Neurophysiol *117*, 1785–1796. https://doi.org/10.1152/jn.00005.2017.

50. McNaughton, B.L., Barnes, C.A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp Brain Res *52*, 41–49.

51. Pavlides, C., and Winson, J. (1989). Influences of Hippocampal Place Cell Firing in the Awake State on the Activity of These Cells during Subsequent Sleep Episodes. J Neurosci *9*, 2907–2918.

52. Lisman, J., and Redish, A.D. (2009). Prediction, sequences and the hippocampus. Philos.Trans.R.Soc.Lond B Biol.Sci. *364*, 1193–1201.

53. Foster, D.J. (2017). Replay Comes of Age. Annu Rev Neurosci *40*, 581–602. https://doi.org/10.1146/annurev-neuro-072116-031538.

54. Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679. https://doi.org/10.1126/science.8036517.

55. Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus *6*, 149–172. https://doi.org/10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K.

56. Deng, X., Liu, D.F., Kay, K., Frank, L.M., and Eden, U.T. (2015). Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter. Neural Comput *27*, 1438–1460. https://doi.org/10.1162/NECO_a_00744.

57. Tingley, D., and Peyrache, A. (2020). On the methods for reactivation and replay analysis. Philos Trans R Soc Lond B Biol Sci *25*, 1799.

58. van der Meer, M.A.A., Kemere, C., Diba, K., and . (2020). Progress and issues in second-order analysis of hippocampal replay. Philos Trans R Soc Lond B Biol Sci *25*, 1799. https://doi.org/10.1098/rstb.2019.0238.

59. Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C., and Wilson, M.A. (1998). A Statistical Paradigm for Neural Spike Train Decoding Applied to Position Prediction from Ensemble Firing Patterns of Rat Hippocampal Place Cells. J Neurosci *18*, 7411–7425. https://doi.org/10.1523/JNEUROSCI.18-18-07411.1998.

60. Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. J.Neurophysiol. *79*, 1017–1044.

61. Denovellis, E.L., Gillespie, A.K., Coulter, M.E., Sosa, M., Chung, J.E., Eden, U.T., and Frank, L.M. (2021). Hippocampal replay of experience at real-world speeds. eLife *10*. https://doi.org/10.7554/eLife.64505.

62. Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. Nat Rev Neurosci *19*, 744–757. https://doi.org/10.1038/s41583-018-0077-1.

63. Drieu, C., and Zugaro, M. (2019). Hippocampal Sequences During Exploration: Mechanisms and Functions. Front Cell Neurosci *13*, 232. https://doi.org/10.3389/fncel.2019.00232.

64. Findlay, G., Tononi, G., and Cirelli, C. (2020). The evolving view of replay and its functions in wake and sleep. Sleep Adv *1*, zpab002. https://doi.org/10.1093/sleepadvances/zpab002.

65. Wikenheiser, A.M., and Redish, A.D. (2014). Decoding the cognitive map: ensemble hippocampal sequences and decision making. Current opinion in neurobiology *32C*, 8–15. https://doi.org/10.1016/j.conb.2014.10.002.

66. Olafsdottir, H.F., Bush, D., and Barry, C. (2018). The Role of Hippocampal Replay in Memory and Planning. Current biology : CB *28*, R37-r50. https://doi.org/10.1016/j.cub.2017.10.073.

67. Robinson, J.C., and Brandon, M.P. (2021). Skipping ahead: A circuit for representing the past, present, and future. Elife *10*. https://doi.org/10.7554/eLife.68795.

68. Carey, A.A., Tanaka, Y., and van der Meer, M.A.A. (2019). Reward revaluation biases hippocampal replay content away from the preferred outcome. Nature Neuroscience *22*, 1450–1459. https://doi.org/10.1038/s41593-019-0464-6.

69. Kay, K., Chung, J.E., Sosa, M., Schor, J.S., Karlsson, M.P., Larkin, M.C., Liu, D.F., and Frank, L.M. (2020). Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. Cell *180*, 552-567 e25. https://doi.org/10.1016/j.cell.2020.01.014.

70. Yu, J.Y., and Frank, L.M. (2021). Prefrontal cortical activity predicts the occurrence of nonlocal hippocampal representations during spatial navigation. PLOS Biology *19*, e3001393. https://doi.org/10.1371/journal.pbio.3001393.

71. Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. J Neurosci *19*, 9497–9507.

72. Louie, K., and Wilson, M.A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron *29*, 145–156. https://doi.org/10.1016/s0896-6273(01)00186-6.

73. Lee, A.K., and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. Neuron *36*, 1183–1194.

74. Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

75. Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat Neurosci *10*, 1241–1242. https://doi.org/10.1038/nn1961.

76. Buzsaki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188.

77. Singer, A.C., and Frank, L.M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. Neuron *64*, 910–921.

78. Ambrose, R.E., Pfeiffer, B.E., and Foster, D.J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. Neuron *91*, 1124–1136. https://doi.org/10.1016/j.neuron.2016.07.047.

79. Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.

80. Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. Nat Neurosci *12*, 913–918.

81. O'Neill, J., Pleydell-Bouverie, B., Dupret, D., and Csicsvari, J. (2010). Play it again: reactivation of waking experience and memory. Trends Neurosci *33*, 220–229. https://doi.org/10.1016/j.tins.2010.01.006.

82. Gupta, A.S., van der Meer, M.A., Touretzky, D.S., and Redish, A.D. (2010). Hippocampal replay is not a simple function of experience. Neuron *65*, 695–705.

83. Wu, C.T., Haggerty, D., Kemere, C., and Ji, D. (2017). Hippocampal awake replay in fear memory retrieval. Nature neuroscience *20*, 571–580. https://doi.org/10.1038/nn.4507.

84. Gillespie, A.K., Astudillo Maya, D.A., Denovellis, E.L., Liu, D.F., Kastner, D.B., Coulter, M.E., Roumis, D.K., Eden, U.T., and Frank, L.M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. Neuron *109*, 3149-3163 e6. https://doi.org/10.1016/j.neuron.2021.07.029.

85. Wu, X., and Foster, D.J. (2014). Hippocampal replay captures the unique topological structure of a novel environment. J Neurosci *34*, 6459–6469.

86. Roux, L., Hu, B., Eichler, R., Stark, E., and Buzsaki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. Nature neuroscience *20*, 845–853. https://doi.org/10.1038/nn.4543.

87. Stella, F., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. Neuron *102*, 450-461 e7. https://doi.org/10.1016/j.neuron.2019.01.052.

88. Grosmark, A.D., Sparks, F.T., Davis, M.J., and Losonczy, A. (2021). Reactivation predicts the consolidation of unbiased long-term cognitive maps. Nature Neuroscience *24*, 1574–1585. https://doi.org/10.1038/s41593-021-00920-7.

89. Vanderwolf, C.H. (1969). Hippocampal electrical activity and voluntary movement in the rat. Electroencephalogr Clin Neurophysiol *26*, 407–418.

90. Buzsáki, G., and Moser, E.I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat Neurosci *16*, 130–138. https://doi.org/10.1038/nn.3304.

91. Colgin, L.L. (2013). Mechanisms and functions of theta rhythms. Annual review of neuroscience *36*, 295–312.

92. Hasselmo, M.E., and Stern, C.E. (2014). Theta rhythm and the encoding and retrieval of space and time. Neuroimage *85 Pt 2*, 656–666. https://doi.org/10.1016/j.neuroimage.2013.06.022.

93. O'Keefe, J. ;Recce M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus, 317–330.

94. Dragoi, G., and Buzsáki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. Neuron *50*, 145–157. https://doi.org/10.1016/j.neuron.2006.02.023.

95. Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2012). Segmentation of spatial experience by hippocampal θ sequences. Nature neuroscience *15*, 1032–1039. https://doi.org/10.1038/nn.3138.

96. Ferbinteanu, J., and Shapiro, M.L. (2003). Prospective and retrospective memory coding in the hippocampus. Neuron *40*, 1227–1239.

97. Itskov, V., Pastalkova, E., Mizuseki, K., Buzsaki, G., and Harris, K.D. (2008). Theta-mediated dynamics of spatial information in hippocampus. J Neurosci *28*, 5959–5964.

98. Wikenheiser, A.M., and Redish, A.D. (2015). Hippocampal theta sequences reflect current goals. Nature neuroscience *18*, 289–294. https://doi.org/10.1038/nn.3909.

99. Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J Neurosci *27*, 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007.

100. Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. Science *321*, 1322–1327. https://doi.org/10.1126/science.1159775.

101. Berners-Lee, A., Wu, X., and Foster, D.J. (2021). Prefrontal Cortical Neurons Are Selective for Non-Local Hippocampal Representations during Replay and Behavior. The Journal of Neuroscience *41*, 5894. https://doi.org/10.1523/JNEUROSCI.1158-20.2021.

102. Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.

103. Xu, H., Baracskay, P., O'Neill, J., and Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze. Neuron *101*, 119-132 e4. https://doi.org/10.1016/j.neuron.2018.11.015.

104. Igata, H., Ikegaya, Y., and Sasaki, T. (2021). Prioritized experience replays on a hippocampal predictive map for learning. Proceedings of the National Academy of Sciences *118*, e2011266118. https://doi.org/10.1073/pnas.2011266118.

105. Buzsáki, G. (2006). Rhythms of the brain (Oxford University Press).

106. Kay, K., Sosa, M., Chung, J.E., Karlsson, M.P., Larkin, M.C., and Frank, L.M. (2016). A hippocampal network for spatial coding during immobility and sleep. Nature *531*, 185–190. https://doi.org/10.1038/nature17144.

107. Yu, J.Y., Kay, K., Liu, D.F., Grossrubatscher, I., Loback, A., Sosa, M., Chung, J.E., Karlsson, M.P., Larkin, M.C., and Frank, L.M. (2017). Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility. Elife *6*, e27621. https://doi.org/10.7554/eLife.27621.

108. Kelemen, E., and Fenton, A.A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. PLoS.Biol. *8*, e1000403.

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

110. Wang, M., Foster David, J., and Pfeiffer Brad, E. (2020). Alternating sequences of future and past behavior encoded within hippocampal theta oscillations. Science *370*, 247–250. https://doi.org/10.1126/science.abb4151.

111. Straube, B. (2012). An overview of the neuro-cognitive processes involved in the encoding, consolidation, and retrieval of true and false memories. Behavioral and Brain Functions : BBF *8*, 35. https://doi.org/10.1186/1744-9081-8-35.

112. Simons, J.S., Garrison, J.R., and Johnson, M.K. (2017). Brain Mechanisms of Reality Monitoring. Trends in cognitive sciences *21*, 462–473. https://doi.org/10.1016/j.tics.2017.03.012.

113. Remondes, M., and Wilson, M.A. (2013). Cingulate-hippocampus coherence and trajectory coding in a sequential choice task. Neuron *80*, 1277–1289. https://doi.org/10.1016/j.neuron.2013.08.037.

114. Jadhav, S.P., Rothschild, G., Roumis, D.K., and Frank, L.M. (2016). Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. Neuron *90*, 113–127. https://doi.org/10.1016/j.neuron.2016.02.010.

115. Shin, J.D., Tang, W., and Jadhav, S.P. (2019). Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. Neuron *104*, 1110-1125.e7. https://doi.org/10.1016/j.neuron.2019.09.012.

116. Zielinski, M.C., Tang, W., and Jadhav, S.P. (2017). The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition. Hippocampus. https://doi.org/10.1002/hipo.22821.

117. Tang, W., Shin, J.D., and Jadhav, S.P. (2021). Multiple time-scales of decision-making in the hippocampus and prefrontal cortex. eLife *10*, e66227. https://doi.org/10.7554/eLife.66227.

118. Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci *10*, 100–107. https://doi.org/10.1038/nn1825.

119. Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., and Pennartz, C.M. (2009). Hippocampus leads ventral striatum in replay of placereward information. Plos Biol *7*, 1000173.

120. Gomperts, S.N., Kloosterman, F., and Wilson, M.A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. Elife *4*. https://doi.org/10.7554/eLife.05360.

121. Girardeau, G., Inema, I., and Buzsaki, G. (2017). Reactivations of emotional memory in the hippocampus-amygdala system during sleep. Nature neuroscience *20*, 1634–1642. https://doi.org/10.1038/nn.4637.

122. Joshi, A., Salib, M., Viney, T.J., Dupret, D., and Somogyi, P. (2017). Behavior-Dependent Activity and Synaptic Organization of Septo-hippocampal GABAergic Neurons Selectively Targeting the Hippocampal CA3 Area. Neuron *96*, 1342-1357 e5. https://doi.org/10.1016/j.neuron.2017.10.033.

123. Rothschild, G., Eban, E., and Frank, L.M. (2017). A cortical–hippocampal–cortical loop of information processing during memory consolidation. Nat Neurosci *20*, 251–259. https://doi.org/10.1038/nn.4457.

124. Tingley, D., and Buzsáki, G. (2018). Transformation of a Spatial Map across the Hippocampal-Lateral Septal Circuit. Neuron.

125. Sosa, M., Joo, H.R., and Frank, L.M. (2020). Dorsal and Ventral Hippocampal Sharp-Wave Ripples Activate Distinct Nucleus Accumbens Networks. Neuron *105*, 725-741 e8. https://doi.org/10.1016/j.neuron.2019.11.022.

126. Nuñez, A., and Buño, W. (2021). The Theta Rhythm of the Hippocampus: From Neuronal and Circuit Mechanisms to Behavior. Frontiers in Cellular Neuroscience *15*. https://doi.org/10.3389/fncel.2021.649262.

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

128. Goutagny, R., Jackson, J., and Williams, S. (2009). Self-generated theta oscillations in the hippocampus. Nature Neuroscience *12*, 1491–1493. https://doi.org/10.1038/nn.2440.

129. Hasselmo, M.E., Bodelon, C., and Wyble, B.P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput *14*, 793–817. https://doi.org/10.1162/089976602317318965.

130. Deshmukh, S.S., and Knierim, J.J. (2012). Hippocampus. Wiley Interdiscip Rev Cogn Sci *3*, 231–251. https://doi.org/10.1002/wcs.1164.

131. Pfeiffer, B.E. (2017). The content of hippocampal “replay.” Hippocampus. https://doi.org/10.1002/hipo.22824.

132. Craver, C.F., Kwan, D., Steindam, C., and Rosenbaum, R.S. (2014). Individuals with episodic amnesia are not stuck in time. Neuropsychologia *57*, 191–195. https://doi.org/10.1016/j.neuropsychologia.2014.03.004.

133. De Brigard, F., and Gessell, B.S. (2016). Time Is Not of the Essence: Understanding the Neural Correlates of Mental Time Travel. In Seeing the Future (Oxford University Press). https://doi.org/10.1093/acprof:oso/9780190241537.003.0008.

134. D’Argembeau, A., Xue, G., Lu, Z.-L., Van der Linden, M., and Bechara, A. (2008). Neural correlates of envisioning emotional events in the near and far future. NeuroImage *40*, 398–407. https://doi.org/10.1016/j.neuroimage.2007.11.025.

135. Johnson, M.K., and Raye, C.L. (1981). Reality monitoring. Psychological Review *88*, 67–85. https://doi.org/10.1037/0033-295X.88.1.67.

136. Cabeza, R., Rao, S.M., Wagner, A.D., Mayer, A.R., and Schacter, D.L. (2001). Can medial temporal lobe regions distinguish true from false? An event-related functional MRI study of veridical and illusory recognition memory. Proceedings of the National Academy of Sciences *98*, 4805–4810. https://doi.org/10.1073/pnas.081082698.

137. Buzsaki, G., Logothetis, N., and Singer, W. (2013). Scaling brain size, keeping timing: evolutionary preservation of brain rhythms. Neuron *80*, 751–764. https://doi.org/10.1016/j.neuron.2013.10.002.

138. Kanamori, N. (1985). A spindle-like wave in the cat hippocampus: a novel vigilance level-dependent electrical activity. Brain research *334*, 180–182.

139. Eguchi, K., and Satoh, T. (1987). Relationship between positive sharp wave bursts and unitary discharges in the cat hippocampus during slow wave sleep. Physiology & behavior *40*, 497–499.

140. Ulanovsky, N., and Moss, C.F. (2007). Hippocampal cellular and network activity in freely moving echolocating bats. Nat Neurosci *10*, 224–233.

141. Nokia, M.S., Penttonen, M., and Wikgren, J. (2010). Hippocampal ripple-contingent training accelerates trace eyeblink conditioning and retards extinction in rabbits. The Journal of neuroscience : the official journal of the Society for Neuroscience *30*, 11486–11492. https://doi.org/10.1523/jneurosci.2165-10.2010.

142. Nokia, M.S., Mikkonen, J.E., Penttonen, M., and Wikgren, J. (2012). Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning. Front Behav Neurosci *6*, 84. https://doi.org/10.3389/fnbeh.2012.00084.

143. Leonard, T.K., Mikkila, J.M., Eskandar, E.N., Gerrard, J.L., Kaping, D., Patel, S.R., Womelsdorf, T., and Hoffman, K.L. (2015). Sharp Wave Ripples during Visual Exploration in the Primate Hippocampus. The Journal of neuroscience : the official

journal of the Society for Neuroscience *35*, 14771–14782. https://doi.org/10.1523/jneurosci.0864-15.2015.

144. Payne, H.L., Lynch, G.F., and Aronov, D. (2021). Neural representations of space in the hippocampus of a food-caching bird. Science *373*, 343–348. https://doi.org/10.1126/science.abg2009.

145. Kurth-Nelson, Z., Economides, M., Dolan, R.J., and Dayan, P. (2016). Fast Sequences of Non-spatial State Representations in Humans. Neuron *91*, 194–204. https://doi.org/10.1016/j.neuron.2016.05.028.

146. Schapiro, A.C., McDevitt, E.A., Rogers, T.T., Mednick, S.C., and Norman, K.A. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. Nature Communications *9*, 3920. https://doi.org/10.1038/s41467-018-06213-1.

147. Liu, Y., Dolan, R.J., Kurth-Nelson, Z., and Behrens, T.E. (2019). Human replay spontaneously reorganizes experience. Cell *178*, 640-652. e14.

148. Schuck, N.W., and Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. Science *364*. https://doi.org/10.1126/science.aaw5181.

149. Vaz, A.P., Wittig, J.H., Inati, S.K., and Zaghloul, K.A. (2020). Replay of cortical spiking sequences during human memory retrieval. Science *367*, 1131–1134.

150. Kahana, M.J., Sekuler, R., Caplan, J.B., Kirschen, M., and Madsen, J.R. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. Nature *399*, 781–784.

151. Killian, N.J., Jutras, M.J., and Buffalo, E.A. (2012). A map of visual space in the primate entorhinal cortex. Nature *491*, 761–764. https://doi.org/10.1038/nature11587.

152. Eliav, T., Geva-Sagiv, M., Yartsev, M.M., Finkelstein, A., Rubin, A., Las, L., and Ulanovsky, N. (2018). Nonoscillatory Phase Coding and Synchronization in the Bat Hippocampal Formation. Cell *175*, 1119-1130. e15.

153. Courellis, H.S., Nummela, S.U., Metke, M., Diehl, G.W., Bussell, R., Cauwenberghs, G., and Miller, C.T. (2019). Spatial encoding in primate hippocampus during free navigation. PLOS Biology *17*, e3000546. https://doi.org/10.1371/journal.pbio.3000546.

154. Qasim, S.E., Fried, I., and Jacobs, J. (2021). Phase precession in the human hippocampus and entorhinal cortex. Cell *184*, 3242-3255.e10. https://doi.org/10.1016/j.cell.2021.04.017.

155. Reddy, L., Self, M.W., Zoefel, B., Poncet, M., Possel, J.K., Peters, J.C., Baayen, J.C., Idema, S., VanRullen, R., and Roelfsema, P.R. (2021). Theta-phase dependent neuronal coding during sequence learning in human single neurons. Nat Commun *12*, 4839. https://doi.org/10.1038/s41467-021-25150-0.

156. Tulving, E., Donaldson, W., Bower, G.H., and United States. Office of Naval Research. (1972). Organization of memory (Academic Press).

157. Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. Science *336*, 1454–1458.

158. Fernández-Ruiz, A., Oliva, A., Fermino de Oliveira, E., Rocha-Almeida, F., Tingley, D., and Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. Science *364*, 1082–1086. https://doi.org/10.1126/science.aax0758.

159. Hunt, L.T., Daw, N.D., Kaanders, P., MacIver, M.A., Mugan, U., Procyk, E., Redish, A.D., Russo, E., Scholl, J., Stachenfeld, K., et al. (2021). Formalizing planning and information search in naturalistic decision-making. Nature Neuroscience *24*, 1051–1064. https://doi.org/10.1038/s41593-021-00866-w.

160. Eichenbaum, H., and Cohen, N.J. (2014). Can we reconcile the declarative memory and spatial navigation views on hippocampal function? Neuron *83*, 764–770.

161. Behrens, T.E., Muller, T.H., Whittington, J.C., Mark, S., Baram, A.B., Stachenfeld, K.L., and Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. Neuron *100*, 490–509.

162. Heckers, S., Zalesak, M., Weiss, A.P., Ditman, T., and Titone, D. (2004). Hippocampal activation during transitive inference in humans. Hippocampus *14*, 153–162.

163. Zeithamova, D., Schlichting, M., and Preston, A. (2012). The hippocampus and inferential reasoning: building memories to navigate future decisions. Frontiers in Human Neuroscience *6*. https://doi.org/10.3389/fnhum.2012.00070.

164. Barron, H.C., Reeve, H.M., Koolschijn, R.S., Perestenko, P.V., Shpektor, A., Nili, H., Rothaermel, R., Campo-Urriza, N., O'Reilly, J.X., Bannerman, D.M., et al. (2020). Neuronal Computation Underlying Inferential Reasoning in Humans and Mice. Cell *183*, 228-243.e21. https://doi.org/10.1016/j.cell.2020.08.035.

165. Cohen, N.J., and Eichenbaum, H. (1993). Memory, amnesia, and the hippocampal system (MIT press).

166. Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron *44*, 109–120.

167. Spiers, H.J. (2020). The Hippocampal Cognitive Map: One Space or Many? Trends in Cognitive Sciences *24*, 168–170. https://doi.org/10.1016/j.tics.2019.12.013.

168. Zeithamova, D., Dominick, A.L., and Preston, A.R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. Neuron *75*, 168–179. https://doi.org/10.1016/j.neuron.2012.05.010.

169. Wood, E.R., Dudchenko, P.A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. Nature *397*, 613–616.

170. Moita, M.A., Rosis, S., Zhou, Y., LeDoux, J.E., and Blair, H.T. (2003). Hippocampal place cells acquire location-specific responses to the conditioned stimulus during auditory fear conditioning. Neuron *37*, 485–497.

171. MacDonald, C.J., Carrow, S., Place, R., and Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. The Journal of neuroscience : the official journal of the Society for Neuroscience *33*, 14607–14616. https://doi.org/10.1523/JNEUROSCI.1537-13.2013.

172. McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., and Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. Neuron *83*, 202–215. https://doi.org/10.1016/j.neuron.2014.05.019.

173. Aronov, D., Nevers, R., and Tank, D.W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. Nature *543*, 719–722. https://doi.org/10.1038/nature21692.

174. Danjo, T., Toyoizumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. Science *359*, 213–218.

175. Duff, M.C., and Brown-Schmidt, S. (2012). The hippocampus and the flexible use and processing of language. Front Hum Neurosci *6*, 69. https://doi.org/10.3389/fnhum.2012.00069.

176. Duff, M.C., Kurczek, J., Rubin, R., Cohen, N.J., and Tranel, D. (2013). Hippocampal amnesia disrupts creative thinking. Hippocampus *23*, 1143–1149. https://doi.org/10.1002/hipo.22208.

177. Eichenbaum, H. (2013). What H.M. taught us. Journal of cognitive neuroscience *25*, 14–21. https://doi.org/10.1162/jocn_a_00285.

178. Wang, M.Z., and Hayden, B.Y. (2021). Latent learning, cognitive maps, and curiosity. Curr Opin Behav Sci *38*, 1–7. https://doi.org/10.1016/j.cobeha.2020.06.003.

179. Chu, J., and Schulz, L.E. (2020). Play, Curiosity, and Cognition. Annual Review of Developmental Psychology *2*, 317–343. https://doi.org/10.1146/annurev-devpsych-070120-014806.

180. Corballis, M.C. (2019). Mental time travel, language, and evolution. Neuropsychologia *134*, 107202. https://doi.org/10.1016/j.neuropsychologia.2019.107202.

181. Buzsáki, G., and Tingley, D. (2018). Space and Time: The Hippocampus as a Sequence Generator. Trends Cogn Sci *22*, 853–869. https://doi.org/10.1016/j.tics.2018.07.006.

182. Sutton, R.S., and Barto, A.G. (1998). Reinforcement learning: an introduction (MIT Press).

183. Sosa, M., and Giocomo, L.M. (2021). Navigating for reward. Nat Rev Neurosci *22*, 472–487. https://doi.org/10.1038/s41583-021-00479-z.

184. Liu, Y., Mattar, M.G., Behrens, T.E.J., Daw, N.D., and Dolan, R.J. (2021). Experience replay is associated with efficient nonlocal learning. Science *372*. https://doi.org/10.1126/science.abf1357.

185. Krausz, T.A., Comrie, A.E., Frank, L.M., Daw, N.D., and Berke, J.D. Dual credit assignment processes underlie dopamine signals in a complex spatial environment. https://doi.org/10.1101/2023.02.15.528738.

186. Dennis, E.J., Hady, A.E., Michaiel, A., Clemens, A., Tervo, D.R.G., Voigts, J., and Datta, S.R. (2021). Systems Neuroscience of Natural Behaviors in Rodents. J. Neurosci. *41*, 911–919. https://doi.org/10.1523/JNEUROSCI.1877-20.2020.

187. Pearson, J.M., Watson, K.K., and Platt, M.L. (2014). Decision Making: The Neuroethological Turn. Neuron *82*, 950–965. https://doi.org/10.1016/j.neuron.2014.04.037.

188. Chettih, S.N., Mackevicius, E.L., Hale, S., and Aronov, D. (2024). Barcoding of episodic memories in the hippocampus of a food-caching bird. Cell *187*, 1922-1935.e20. https://doi.org/10.1016/j.cell.2024.02.032.

189. Rosenberg, M., Zhang, T., Perona, P., and Meister, M. (2021). Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration. eLife *10*, e66175. https://doi.org/10.7554/eLife.66175.

190. Ding, S.S., Fox, J.L., Gordus, A., Joshi, A., Liao, J.C., and Scholz, M. (2024). Fantastic beasts and how to study them: rethinking experimental animal behavior. Journal of Experimental Biology *227*, jeb247003. https://doi.org/10.1242/jeb.247003.

191. Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat Neurosci *8*, 1704–1711. https://doi.org/10.1038/nn1560.

192. Kim, S.M., and Frank, L.M. (2009). Hippocampal lesions impair rapid learning of a continuous spatial alternation task. PLoS ONE *4*, e5494.

193. Morris, R.G., Garrud, P., Rawlins, J.N., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. Nature *297*, 681–683.

194. Riedel, G., Micheau, J., Lam, A.G., Roloff, E.L., Martin, S.J., Bridge, H., de, H.L., Poeschel, B., McCulloch, J., and Morris, R.G. (1999). Reversible neural inactivation reveals hippocampal participation in several memory processes. Nat.Neurosci *2*, 898–905.

195. Packard, M.G., and McGaugh, J.L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. Neurobiol.Learn.Mem. *65*, 65–72.

196. Robinson, N.T.M., Descamps, L.A.L., Russell, L.E., Buchholz, M.O., Bicknell, B.A., Antonov, G.K., Lau, J.Y.N., Nutbrown, R., Schmidt-Hieber, C., and Häusser, M. (2020). Targeted Activation of Hippocampal Place Cells Drives Memory-Guided

Spatial Behavior. Cell *183*, 1586-1599.e10. https://doi.org/10.1016/j.cell.2020.09.061.

197. Robbe, D., and Buzsaki, G. (2009). Alteration of theta timescale dynamics of hippocampal place cells by a cannabinoid is associated with memory impairment. Journal of Neuroscience *29*, 12597–12605.

198. Siegle, J.H., and Wilson, M.A. (2014). Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. eLife *3*, e03061. https://doi.org/10.7554/eLife.03061.

199. Liu, C., Todorova, R., Tang, W., Oliva, A., and Fernandez-Ruiz, A. (2023). Associative and predictive hippocampal codes support memory-guided behaviors. Science *382*, eadi8237. https://doi.org/10.1126/science.adi8237.

200. Foster, D.J., and Wilson, M.A. (2007). Hippocampal theta sequences. Hippocampus *17*, 1093–1099. https://doi.org/10.1002/hipo.20345.

201. Kudrimoti, H.S., Barnes, C.A., and McNaughton, B.L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. J Neurosci *19*, 4090–4101.

202. Csicsvari, J., O'Neill, J., Allen, K., and Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. Eur J Neurosci *26*, 704–716. https://doi.org/10.1111/j.1460-9568.2007.05684.x.

203. Diba, K., and Buzsaki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat Neurosci *10*, 1241–1242.

204. Kay, K., and Frank, L.M. (2018). Three brain states in the hippocampus and cortex. Hippocampus. https://doi.org/10.1002/hipo.22956.

205. Papale, A.E., Zielinski, M.C., Frank, L.M., Jadhav, S.P., and Redish, A.D. (2016). Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making. Neuron *92*, 975–982. https://doi.org/10.1016/j.neuron.2016.10.028.

206. Shadlen, M.N., and Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. Neuron *90*, 927–939. https://doi.org/10.1016/j.neuron.2016.04.036.

207. Jensen, O., and Lisman, J.E. (1996). Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. Learn.Mem. *3*, 279–287.

208. Jensen, O., and Lisman, J.E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. Trends Neurosci. *28*, 67–72.

209. Bi, G.Q., and Poo, M.M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. Journal of Neuroscience *18*, 10464–10472.

210. Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron *25*, 707--15.

211. Stephens, D.W., and Krebs, J.R. (1986). Foraging Theory (Princeton University Press).

212. Denovellis, E.L., Gillespie, A.K., Coulter, M.E., Sosa, M., Chung, J.E., Eden, U.T., and Frank, L.M. (2021). Hippocampal replay of experience at real-world speeds. Elife *10*. https://doi.org/10.7554/eLife.64505.

213. Deng, X., Liu, D.F., Kay, K., Frank, L.M., and Eden, U.T. (2015). Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter. Neural Computation *27*, 1438–1460. https://doi.org/10.1162/NECO_a_00744.

214. Johnson, A., Fenton, A.A., Kentros, C., and Redish, A.D. (2009). Looking for cognition in the structure within the noise. Trends Cogn Sci. *13*, 55–64.

215. Callaway, F., Rangel, A., and Griffiths, T.L. (2021). Fixation patterns in simple choice reflect optimal information sampling. PLOS Computational Biology *17*, e1008863. https://doi.org/10.1371/journal.pcbi.1008863.

216. Tolman, E.C. (1938). The determiners of behavior at a choice point. Psychol Rev *45*, 1–41.

217. Buzsaki, G. (1989). Two-stage model of memory trace formation: a role for “noisy” brain states. Neuroscience *31*, 551–570. https://doi.org/10.1016/0306-4522(89)90423-5.

218. Buzsaki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. Neuron *68*, 362–385. https://doi.org/10.1016/j.neuron.2010.09.023.

219. Daw, N.D. (2011). Model-based influences on humans' choices and striatal prediction errors. Neuron *69*, 1204–1215. https://doi.org/10.1016/j.neuron.2011.02.027.

220. Doll, B.B., Duncan, K.D., Simon, D.A., Shohamy, D., and Daw, N.D. (2015). Model-based choices involve prospective neural activity. Nat Neurosci *18*, 767–772. https://doi.org/10.1038/nn.3981.

221. Hasz, B.M., and Redish, A.D. (2018). Deliberation and Procedural Automation on a Two-Step Task for Rats. Front. Integr. Neurosci. *12*. https://doi.org/10.3389/fnint.2018.00030.

222. Voigts, J., Kanitscheider, I., Miller, N.J., Toloza, E.H.S., Newman, J.P., Fiete, I.R., and Harnett, M.T. (2022). Spatial reasoning via recurrent neural dynamics in mouse retrosplenial cortex. Preprint at bioRxiv, https://doi.org/10.1101/2022.04.12.488024 https://doi.org/10.1101/2022.04.12.488024.

223. Mattar, M.G., and Daw, N.D. (2018). Prioritized memory access explains planning and hippocampal replay. Nature Neuroscience *21*, 1609–1617. https://doi.org/10.1038/s41593-018-0232-z.

224. Bhattarai, B., Lee, J.W., and Jung, M.W. (2020). Distinct effects of reward and navigation history on hippocampal forward and reverse replays. Proceedings of the

National Academy of Sciences *117*, 689–697. https://doi.org/10.1073/pnas.1912533117.

225. McNamara, C.G., Tejero-Cantero, A., Trouche, S., Campo-Urriza, N., and Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. Nat Neurosci *17*, 1658–1660. https://doi.org/10.1038/nn.3843.

226. Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., and Pennartz, C.M. (2009). Hippocampus leads ventral striatum in replay of place-reward information. PLoS Biol. *7*, e1000173. https://doi.org/10.1371/journal.pbio.1000173.

227. van der Meer, M.A., and Redish, A.D. (2011). Theta phase precession in rat ventral striatum links place and reward information. The Journal of neuroscience : the official journal of the Society for Neuroscience *31*, 2843–2854. https://doi.org/10.1523/JNEUROSCI.4869-10.2011.

228. Lee, H., Ghim, J.-W., Kim, H., Lee, D., and Jung, M. (2012). Hippocampal Neural Correlates for Values of Experienced Events. J. Neurosci. *32*, 15053–15065. https://doi.org/10.1523/JNEUROSCI.2806-12.2012.

229. Comrie, A.E., Frank, L.M., and Kay, K. (2022). Imagination as a fundamental function of the hippocampus. Philosophical Transactions of the Royal Society B: Biological Sciences *377*, 20210336. https://doi.org/10.1098/rstb.2021.0336.

230. McNamee, D.C. (2024). The generative neural microdynamics of cognitive processing. Current Opinion in Neurobiology *85*, 102855. https://doi.org/10.1016/j.conb.2024.102855.

231. Barron, H.C., Dolan, R.J., and Behrens, T.E. (2013). Online evaluation of novel choices by simultaneous representation of multiple memories. Nat. Neurosci. *16*, 1492–1498. https://doi.org/10.1038/nn.3515.

232. Reber, T.P., Luechinger, R., Boesiger, P., and Henke, K. (2012). Unconscious Relational Inference Recruits the Hippocampus. J. Neurosci. *32*, 6138–6148. https://doi.org/10.1523/JNEUROSCI.5639-11.2012.

233. Schwartenbeck, P., Baram, A., Liu, Y., Mark, S., Muller, T., Dolan, R., Botvinick, M., Kurth-Nelson, Z., and Behrens, T. (2023). Generative replay underlies compositional inference in the hippocampal-prefrontal circuit. Cell *186*, 4885-4897.e14. https://doi.org/10.1016/j.cell.2023.09.004.

234. Nour, M.M., Liu, Y., Arumuham, A., Kurth-Nelson, Z., and Dolan, R.J. (2021). Impaired neural replay of inferred relationships in schizophrenia. Cell *184*, 4315-4328.e17. https://doi.org/10.1016/j.cell.2021.06.012.

235. Wimmer, G.E. (2012). Preference by association: how memory mechanisms in the hippocampus bias decisions. Science *338*, 270–273. https://doi.org/10.1126/science.1223252.

236. Aru, J., Drüke, M., Pikamäe, J., and Larkum, M.E. (2023). Mental navigation and the neural mechanisms of insight. Trends in Neurosciences *46*, 100–109. https://doi.org/10.1016/j.tins.2022.11.002.

237. Kaplan, R., Schuck, N.W., and Doeller, C.F. (2017). The Role of Mental Maps in Decision-Making. Trends in Neurosciences *40*, 256–259. https://doi.org/10.1016/j.tins.2017.03.002.

238. Wallenstein, G.V., Eichenbaum, H., and Hasselmo, M.E. (1998). The hippocampus as an associator of discontiguous events. Trends.Neurosci. *21*, 317–323.

239. Knudsen, E.B., and Wallis, J.D. (2021). Hippocampal neurons construct a map of an abstract value space. Cell *184*, 4640-4650.e10. https://doi.org/10.1016/j.cell.2021.07.010.

240. Whittington, J.C.R., Muller, T.H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T.E.J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. Cell *183*, 1249-1263.e23. https://doi.org/10.1016/j.cell.2020.10.024.

241. Davidson, J.D., and Hady, A.E. (2019). Foraging as an evidence accumulation process. PLOS Computational Biology *15*, e1007060. https://doi.org/10.1371/journal.pcbi.1007060.

242. Brown, L.S., Cho, J.R., Bolkan, S.S., Nieh, E.H., Schottdorf, M., Tank, D.W., Brody, C.D., Witten, I.B., and Goldman, M.S. (2024). Neural circuit models for evidence

accumulation through choice-selective sequences. Preprint at bioRxiv, https://doi.org/10.1101/2023.09.01.555612.

243. Harlow, H.F. (1949). The formation of learning sets. Psychological review *56*, 51.

244. Kolling, N., Behrens, T.E., Mars, R.B., and Rushworth, M.F. (2012). Neural Mechanisms of Foraging. Science *336*, 95–98. https://doi.org/10.1126/science.1216930.

245. Kennerley, S.W., Dahmubed, A.F., Lara, A.H., and Wallis, J.D. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. Journal of cognitive neuroscience *21*, 1162–1178. https://doi.org/10.1162/jocn.2009.21100.

246. Rushworth M.F., N.M.P., Boorman E.D.,. Walton M.E.,. Behrens T.E. (2011). Frontal cortex and reward-guided learning and decision-making. Neuron *70*, 1054–1069.

247. Mohebi, A., Pettibone, J.R., Hamid, A.A., Wong, J.T., Vinson, L.T., Patriarchi, T., Tian, L., Kennedy, R.T., and Berke, J.D. (2019). Dissociable dopamine dynamics for learning and motivation. Nature *570*, 65–70. https://doi.org/10.1038/s41586-019-1235-y.

248. Karlsson, M.P., Tervo, D.G., and Karpova, A.Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. Science *338*, 135–139. https://doi.org/10.1126/science.1226518.

249. Wang, J.X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J.Z., Hassabis, D., and Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. Nat Neurosci *21*, 860–868. https://doi.org/10.1038/s41593-018-0147-8.

250. Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. Nat Neurosci *14*, 933–939. https://doi.org/10.1038/nn.2856.

251. Vertechi, P., Lottem, E., Sarra, D., Godinho, B., Treves, I., Quendera, T., Lohuis, M.N.O., and Mainen, Z.F. (2020). Inference-Based Decisions in a Hidden State Foraging Task: Differential Contributions of Prefrontal Cortical Areas. Neuron *106*, 166-176.e6. https://doi.org/10.1016/j.neuron.2020.01.017.

252. Shin, J.D., and Jadhav, S.P. (2016). Multiple modes of hippocampal-prefrontal interactions in memory-guided behavior. Current opinion in neurobiology *40*, 161–169. https://doi.org/10.1016/j.conb.2016.07.015.

253. Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., and Wiener, S.I. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. Neuron *66*, 921–936. https://doi.org/10.1016/j.neuron.2010.05.013.

254. Joshi, A., Denovellis, E.L., Mankili, A., Meneksedag, Y., Davidson, T.J., Gillespie, A.K., Guidera, J.A., Roumis, D., and Frank, L.M. (2023). Dynamic synchronization

between hippocampal representations and stepping. Nature *617*, 125–131. https://doi.org/10.1038/s41586-023-05928-6.

255. Chung, J.E., Joo, H.R., Fan, J.L., Liu, D.F., Barnett, A.H., Chen, S., Geaghan-Breiner, C., Karlsson, M.P., Karlsson, M., Lee, K.Y., et al. (2019). High-Density, Long-Lasting, and Multi-region Electrophysiological Recordings Using Polymer Electrode Arrays. Neuron *101*, 21-31.e5. https://doi.org/10.1016/j.neuron.2018.11.002.

256. Guidera, J.A., Gramling, D.P., Comrie, A.E., Joshi, A., Denovellis, E.L., Lee, K.H., Zhou, J., Thompson, P., Hernandez, J., Yorita, A., et al. (2024). Regional specialization manifests in the reliability of neural population codes. Preprint at bioRxiv, https://doi.org/10.1101/2024.01.25.576941 https://doi.org/10.1101/2024.01.25.576941.

257. Lee, K.H., Denovellis, E., Ly, R., Magland, J., Soules, J., Comrie, A.E., Gramling, D.P., Guidera, J.A., Nevers, R., Adenekan, P., et al. (2024). Spyglass: a data analysis framework for reproducible and shareable neuroscience research (Neuroscience) https://doi.org/10.1101/2024.01.25.577295.

258. Sutton, R.S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. SIGART Bull. *2*, 160–163. https://doi.org/10.1145/122344.122377.

259. Huys, Q.J.M., Cools, R., Gölzer, M., Friedel, E., Heinz, A., Dolan, R.J., and Dayan, P. (2011). Disentangling the Roles of Approach, Activation and Valence in

Instrumental and Pavlovian Responding. PLOS Computational Biology *7*, e1002028. https://doi.org/10.1371/journal.pcbi.1002028.

260. Denovellis, E.L., Frank, L.M., and Eden, U.T. (2019). Characterizing hippocampal replay using hybrid point process state space models. In 2019 53rd Asilomar Conference on Signals, Systems, and Computers (IEEE), pp. 245–249. https://doi.org/10.1109/IEEECONF44664.2019.9048688.

261. Kloosterman, F., Layton, S.P., Chen, Z., and Wilson, M.A. (2013). Bayesian Decoding using Unsorted Spikes in the Rat Hippocampus. Journal of neurophysiology. https://doi.org/10.1152/jn.01046.2012.

262. Gillespie, A.K., Astudillo Maya, D.A., Denovellis, E.L., Desse, S., and Frank, L.M. (2022). Neurofeedback training can modulate task-relevant memory replay rate in rats. https://doi.org/10.1101/2022.10.13.512183.

263. Dijkstra, E.W. (2022). A Note on Two Problems in Connexion with Graphs. In Edsger Wybe Dijkstra: His Life,Work, and Legacy (Association for Computing Machinery), pp. 287–290.

264. Skaggs, W.E. (1995). Relations between the theta rhythm and activity patterns of hippocampal neurons.

265. Mountcastle, V.B. (1998). Perceptual neuroscience : the cerebral cortex (Harvard University Press).

266. Sherrington, C.S. (1906). The integrative action of the nervous system (Yale University Press) https://doi.org/10.1037/13798-000.

267. Campagner, D., Vale, R., Tan, Y.L., Iordanidou, P., Pavón Arocas, O., Claudi, F., Stempel, A.V., Keshavarzi, S., Petersen, R.S., Margrie, T.W., et al. (2023). A cortico-collicular circuit for orienting to shelter during escape. Nature *613*, 111–119. https://doi.org/10.1038/s41586-022-05553-9.

268. Yang, W., Sun, C., Huszár, R., Hainmueller, T., and Buzsáki, G. (2023). Selection of experience for memory by hippocampal sharp wave ripples. bioRxiv, 2023.11.07.565935. https://doi.org/10.1101/2023.11.07.565935.

269. Berners-Lee, A., Feng, T., Silva, D., Wu, X., Ambrose, E.R., Pfeiffer, B.E., and Foster, D.J. (2022). Hippocampal replays appear after a single experience and incorporate greater detail with more experience. Neuron *110*, 1829-1842.e5. https://doi.org/10.1016/j.neuron.2022.03.010.

270. Sutton, R., and Barto, A. (1998). Reinforcement Learning (MIT Press).

271. Drieu, C., Todorova, R., and Zugaro, M. (2018). Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. Science *362*, 675–679.

272. Eichenbaum, H. (2015). Does the hippocampus preplay memories? Nat. Neurosci. *18*, 1701–1702. https://doi.org/10.1038/nn.4180.

273. Pratt, W.E., and Mizumori, S.J. (2001). Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. Behav Brain Res *123*, 165–183. https://doi.org/10.1016/s0166-4328(01)00204-2.

274. Wallis, J.D., and Kennerley, S.W. (2010). Heterogeneous reward signals in prefrontal cortex. Current opinion in neurobiology *20*, 191–198. https://doi.org/10.1016/j.conb.2010.02.009.

275. Horst NK, L.M. (2012). Working with memory: evidence for a role for the medial prefrontal cortex in performance monitoring during spatial delayed alternation. J Neurophysiol *108*, 3276–3288.

276. Rushworth, M.F.S., and Behrens, T.E.J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. Nature neuroscience *11*, 389–397.

277. Funamizu A, I.M., Doya K, Kanzaki R, Takahashi H. (2015). Condition interference in rats performing a choice task with switched variable and fixed reward conditions. Front Neurosci.

278. Otis, J.M., Namboodiri, V.M., Matan, A.M., Voets, E.S., Mohorn, E.P., Kosyk, O., McHenry, J.A., Robinson, J.E., Resendez, S.L., Rossi, M.A., et al. (2017). Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. Nature *543*, 103–107. https://doi.org/10.1038/nature21376.

279. Siapas, A.G., Lubenov, E.V., and Wilson, M.A. (2005). Prefrontal phase locking to hippocampal theta oscillations. Neuron *46*, 141–151. https://doi.org/10.1016/j.neuron.2005.02.028.

280. Jones, M.W., and Wilson, M.A. (2005). Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. Hippocampus *15*, 867–873. https://doi.org/10.1002/hipo.20119.

281. Jones, M.W., and Wilson, M.A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. PLoS Biol *3*, e402. https://doi.org/10.1371/journal.pbio.0030402.

282. Hyman, J.M., Zilli, E.A., Paley, A.M., and Hasselmo, M.E. (2005). Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. Hippocampus *15*, 739–749. https://doi.org/10.1002/hipo.20106.

283. Swanson, L.W. (1981). A direct projection from Ammon's horn to prefrontal cortex in the rat. Brain Res *217*, 150–154.

284. Thierry, A.M., Gioanni, Y., Degenetais, E., and Glowinski, J. (2000). Hippocampo-prefrontal cortex pathway: anatomical and electrophysiological characteristics. Hippocampus *10*, 411--9.

285. Hoover, W.B., and Vertes, R.P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. Brain Struct Funct *212*, 149–179. https://doi.org/10.1007/s00429-007-0150-4.

286. Sun, X., Comrie, A.E., Monroe, E., Kahn, A., Guidera, J.A., Tong, L., Denovellis, E., Krausz, T.A., Shin, D., Berke, J., et al. Dynamic value codes in the medial prefrontal cortex to inform decisionmaking. Cosyne 2024 Abstracts.

287. Foster, D.J., and Wilson, M.A. (2007). Hippocampal theta sequences. Hippocampus *17*, 1093–1099. https://doi.org/10.1002/hipo.20345.

288. Chung, J.E., Joo, H.R., Smyth, C.N., Fan, J.L., Geaghan-Breiner, C., Liang, H., Liu, D.F., Roumis, D., Chen, S., Lee, K.Y., et al. (2019). Chronic Implantation of Multiple Flexible Polymer Electrode Arrays. J Vis Exp. https://doi.org/10.3791/59957.

289. Aydin, C., Okun, M., and Gardner, R.J. (2020). Neuropixels 2 . 0 : A miniaturized high-density probe for stable , long-term brain recordings.

290. Buzsáki, G. (2019). The brain from inside out (Oxford University Press).

291. Etter, G., van der Veldt, S., Choi, J., and Williams, S. (2023). Optogenetic frequency scrambling of hippocampal theta oscillations dissociates working memory retrieval from hippocampal spatiotemporal codes. Nat Commun *14*, 410. https://doi.org/10.1038/s41467-023-35825-5.

292. Quirk, C.R., Zutshi, I., Srikanth, S., Fu, M.L., Marciano, N.D., Wright, M.K., Parsey, D.F., Liu, S., Siretskiy, R.E., Huynh, T.L., et al. (2021). Precisely Timed Theta Oscillations are Selectively Required During the Encoding Phase of Memory. Nat Neurosci *24*, 1614–1627. https://doi.org/10.1038/s41593-021-00919-0.

293. Joshi, A., Comrie, A.E., Mankili, A., Guidera, J.A., Bray, S., Nevers, R., Sun, X., Monroe, E., Kharazia, V., Astudillo Maya, D.A., et al. (2024). Specific ablation of Hippocampal Theta Activity During Locomotion Impairs Learning. In FENS 2024.

294. Coulter, M.E., Gillespie, A.K., Chu, J., Denovellis, E.L., Nguyen, T.T.K., Liu, D.F., Wadhwani, K., Sharma, B., Wang, K., Deng, X., et al. (2024). Closed-loop modulation of remote hippocampal representations with neurofeedback. bioRxiv, 2024.05.08.593085. https://doi.org/10.1101/2024.05.08.593085.

295. Chu, J.P., Coulter, M.E., Denovellis, E.L., Nguyen, T.T.K., Liu, D.F., Deng, X., Eden, U.T., Kemere, C.T., and Frank, L.M. (2024). RealtimeDecoder: A fast software module for online clusterless decoding. Preprint, https://doi.org/10.1101/2024.05.03.592417.

296. Lai, C., Tanaka, S., Harris, T.D., and Lee, A.K. (2023). Volitional activation of remote place representations with a hippocampal brain–machine interface. Science *382*, 566–573. https://doi.org/10.1126/science.adh5206.

297. Gillespie, A.K., Maya, D.A.A., Denovellis, E.L., Desse, S., and Frank, L.M. (2024). Neurofeedback training can modulate task-relevant memory replay rate in rats. eLife *12*. https://doi.org/10.7554/eLife.90944.2.

298. Wijnen, K., Genzel, L., and van der Meij, J. (2024). Rodent maze studies: from following simple rules to complex map learning. Brain Struct Funct *229*, 823–841. https://doi.org/10.1007/s00429-024-02771-x.

299. Dayan, P., and Daw, N.D. (2008). Decision theory, reinforcement learning, and the brain. Cognitive, Affective, & Behavioral Neuroscience *8*, 429–453. https://doi.org/10.3758/CABN.8.4.429.

300. Gupta, A.S., van der Meer, M.A., Touretzky, D.S., and Redish, A.D. (2012). Segmentation of spatial experience by hippocampal theta sequences. Nature neuroscience *15*, 1032–1039.

301. Mathis, M.W., and Mathis, A. (2020). Deep learning tools for the measurement of animal behavior in neuroscience. Current Opinion in Neurobiology *60*, 1–11. https://doi.org/10.1016/j.conb.2019.10.008.

302. Chung, J.E., Magland, J.F., Barnett, A.H., Tolosa, V.M., Tooker, A.C., Lee, K.Y., Shah, K.G., Felix, S.H., Frank, L.M., and Greengard, L.F. (2017). A Fully Automated Approach to Spike Sorting. Neuron *95*, 1381-1394.e6. https://doi.org/10.1016/j.neuron.2017.08.030.

303. Chung, J.E., Joo, H.R., Fan, J.L., Liu, D.F., Barnett, A.H., Chen, S., Geaghan-Breiner, C., Karlsson, M.P., Karlsson, M., Lee, K.Y., et al. (2019). High-Density, Long-Lasting, and Multi-region Electrophysiological Recordings Using Polymer Electrode Arrays. Neuron *101*, 21-31.e5. https://doi.org/10.1016/j.neuron.2018.11.002.

304. Topalovic, U., Barclay, S., Ling, C., Alzuhair, A., Yu, W., Hokhikyan, V., Chandrakumar, H., Rozgic, D., Jiang, W., Basir-Kazeruni, S., et al. (2023). A wearable platform for closed-loop stimulation and recording of single-neuron and

local field potential activity in freely moving humans. Nat Neurosci *26*, 517–527. https://doi.org/10.1038/s41593-023-01260-4.

Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

Signed by:
Alison E. Comne
923BA6FAEB70497...
Author Signature

9/29/2024
Date