Symposium

Learning to Navigate a Changing World: From Brain Systems to Novel Algorithms

J. Quinn Lee,[1,2,3] Ipshita Zutshi,[4,5] Andrew S. Alexander,[6,7,8] and Ann M. Hermundstad[9]

[1]Department of Psychology, University of Alberta, Edmonton, Alberta T6G 2E9, [2]Alberta Machine Intelligence Institute, Edmonton, Alberta T5J 1S5, [3]Canadian Institute for Advanced Research (CIFAR) AI Chair, Toronto, Ontario MG5 1M1, Canada, [4]Department of Neuroscience and Cell Biology, Rutgers Health, Rutgers, Piscataway, New Jersey 08854, [5]The State University of New Jersey, Piscataway, New Jersey 08854, [6]Interdepartmental Graduate Program in Dynamical Neuroscience, UC Santa Barbara, Santa Barbara, California 93106-9661, [7]Neuroscience Research Institute, UC Santa Barbara, Santa Barbara, California 93106-5060, [8]Department of Psychological and Brain Science, UC Santa Barbara, Santa Barbara, California 93106-9660, and [9]Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia 20147

Natural organisms navigate a world in constant flux and are frequently confronted with sensory ambiguity and shifting goals. While biological systems can adapt rapidly to navigate in novel and nonstationary environments, this challenge remains a significant hurdle for artificial intelligence. In this review, we examine recent progress in understanding neurobiological systems that allow for adaptive, efficient spatial and abstract navigation in dynamic environments and how these observations have informed new developments for the same behavior in machines. We highlight how predictive learning serves as a unifying framework to describe the organization of neural activity across distributed systems for navigation, including the hippocampal formation, parietal cortex, and higher-order associative structures. Specifically, we explore how such neurobiological systems encode expected future states, dynamic target trajectories, action plans, and internal behavioral strategies. Finally, we contrast the diffuse, inefficient exploration of canonical artificial agents with the structured, efficient control policies for goal-directed navigation observed in natural organisms. By grounding learning in the interplay between action, prediction, and natural behavioral statistics, navigation provides a roadmap for building resilient, neurobiologically inspired continual learning in artificial agents.

Key words: **artificial intelligence; continual learning; naturalistic behaviour; navigation; systems neuroscience**

Introduction

Natural organisms learn in a world that continues to change throughout their lifetime, and survival demands adaptation of knowledge and behavior in the face of such change. Sensory observations in natural environments are often ambiguous and shaped by unobserved variables that are in constant flux. Learning further requires organisms to decrease their uncertainty about the environment through interaction (Meister, 2022; Vanchurin et al., 2022) and to remember how to act from prior experience.

This challenge is often referred to as the continual learning problem, wherein an organism or agent is required to learn and remember in a changing or nonstationary world. While biological intelligence has evolved to survive and thrive in nonstationary environments, artificial systems struggle to learn, remember, and adapt to change, likely due to a combination of catastrophic interference and loss of plasticity (McCloskey and Cohen, 1989; McClelland et al., 1995; Dohare et al., 2024). While attention and progress with these challenges have grown, there remains no consensus on how to build artificial systems that continue to learn and remember in a changing environment. A number of groups have suggested that the gap between biological and artificial intelligence (AI) to learn continually could be closed through a deeper understanding of how animals learn and remember and the inextricable link between adaptive neural coding in the brain and behavior (Flesch et al., 2023; Zador et al., 2023; Mathis, 2026).

While any naturalistic behavior might demand continual learning, navigation offers an example critical for survival (e.g., locating food, shelter, or threats). Navigation tasks are learned rapidly across species, even when spaces are large and complex, and natural organisms can rapidly adapt their behavior to changing environments and goals, often with one- or few-shot learning (Wolf and Heisenberg, 1991; Ofstad et al., 2011; Meister, 2022; Dan et al., 2024; Mathis, 2026; AbdelRahman et al., 2026; Gong et al., 2026) and rapid adaptation of prior task knowledge (El-Gaby et al., 2024). Indeed, a wide range of tasks, involving abstract, structural knowledge, and nonspatial features, could be considered "navigation" problems (Whittington et al., 2022; Neupane et al., 2024; Qu et al., 2026), and decades of research have uncovered neural circuits and brain systems for learning and remembering how to navigate in both stationary and changing settings (Moser et al., 2017; Basu and Nagel, 2024).

Received May 27, 2026; revised Aug. 20, 2026; accepted Aug. 31, 2026.

Author contributions: J.Q.L., I.Z., A.S.A., and A.M.H. wrote the paper.

J.Q.L. was supported by the Canadian Institute for Advanced Research and the Natural Sciences and Engineering Research Council of Canada (RGPIN-2026-05763). I.Z. was supported by the Simons Foundation (SCGB TTI), and A.M.M. was supported by Howard Hughes Medical Institute. A.S.A. was supported by NIH R00 NS11966 and the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR) program as part of the DOE Collaborative Research in Computational Neuroscience (CRCNS) Program.

The authors declare no competing financial interests.

Correspondence should be addressed to J. Quinn Lee at quinn.lee@ualberta.ca.

https://doi.org/10.1523/JNEUROSCI.1331-26.2026

Copyright © 2026 the authors

In this review, we discuss recent findings on the neural systems and adaptive behavior in animals and agents learning to navigate in a changing world. Specifically, the ability of natural organisms to acquire new information during navigation while preserving prior knowledge is likely supported through interactions across brain systems for rapid acquisition of and transfer of structural knowledge, action planning, coordinate transformation, and flexible switching between behavioral policies. A common theme emerging in this research, which builds on theoretical foundations of predictive and efficient coding (Barlow, 1961; Tishby et al., 1999; Palmer et al., 2015; Chalk et al., 2018), is that animals learn to predict changes in the environment that relate their actions to sensory observations, and this prediction is reflected in neural coding schemes across systems known to have a causal role in navigation, learning, and memory. Based on these observations, we suggest the spatiotemporal structure of neural coding in higher-order brain systems reflects changing observations of the environment that organisms actively predict. The neural encoding of predictions is expressed in the local, predictive representation of environment structure that emerges during navigation and nonlocal trajectories through navigable space during action planning. From this perspective, behavior (action) takes a central role in shaping neural representation and highlights a present gap in our understanding of comparative behavior in artificial and biological agents in common task spaces. Namely, animals demonstrate structured, efficient exploration, learning, and knowledge transfer that canonical artificial agents often do not show or have yet to be examined in analogous tasks. In pursuing a deeper understanding of how biological intelligence addresses the need to navigate in a changing world, we consider how these developments might shape the future of intelligence research.

Learning Structured Representations of Changing Environments: The Power of Prediction

Predictive learning has long been used as a theoretical framework to explain and model intelligent behavior in natural organisms and artificial agents. Evidence for learning with temporal prediction errors, from rewards to high-dimensional sensory observations, has been observed in brain systems across diverse species. The differences between the role of distinct brain systems in predictive learning might largely be explained by what a system cares to predict: while distributions of reward prediction errors have been used to describe dopamine neuronal population activity (Schultz et al., 1997; Dabney et al., 2020; Masset et al., 2025; but see also Coddington et al., 2023; Grima et al., 2026), self-supervised predictions from visual observations can account for learning in the visual cortex (Bakhtiari et al., 2021; Nejad et al., 2025). At the single-neuron level, prediction has been suggested as a normative model of learning to minimize the metabolic cost of neuronal activity (Luczak et al., 2022). Importantly, the type of predictive learning observed across brain systems and tasks likely differs based on the computational problem to be addressed, from learning signals to the timescales of prediction. While a complete review of predictive learning and computation is beyond the scope of this article, we focus here on several models of predictive learning that account for neural activity observed in several brain systems in the navigation setting.

To know where she is and where she is going in the world, a navigator must integrate available sensory observations (vision, audition, and idiothetic cues) to infer her sense of space. To do so, the brain has been classically thought to build cognitive maps that reflect the spatiotemporal structure of the environment and one's location within it (O'Keefe and Nadel, 1978). From place to head direction and grid cells, the neurobiological instantiation of cognitive maps involves at least the hippocampal formation and connected thalamic and associative cortical networks (Moser et al., 2017). While a large body of research has characterized the formation and adaptation of cognitive maps in a number of model organisms, recent theoretical work has begun to unify a broad range of observations across these brain systems using predictive learning rules in artificial agents (Whittington et al., 2022).

While cognitive mapping in the hippocampus has been a central focus of modeling in navigation and memory since the discovery of place cells, earlier computational models used bespoke, hardwired networks to explain patterns observed in hippocampal place cells and entorhinal grid cells (Samsonovich and McNaughton, 1997; Hartley et al., 2000; Barry et al., 2006; Solstad et al., 2006), and thus such models did not learn. These influential models of cognitive mapping were often agnostic on the known role of the hippocampus in memory, while models of memory did not account for cognitive maps per se (McClelland et al., 1995; O'Reilly and Rudy, 2001; O'Reilly and Norman, 2002). Today, this theoretical division has begun to close with the application of learning algorithms that instead suggest the hippocampus encodes a predictive internal model of the world that links together observations of both spatial and nonspatial variables (Stachenfeld et al., 2017; George et al., 2021; Whittington et al., 2022; Carvalho et al., 2024; Levenstein et al., 2024; de Cothi et al., 2026).

A growing number of theoretical models have recently been proposed to account for neural coding observed in the hippocampal formation and connected brain systems (Stachenfeld et al., 2017; Whittington et al., 2020; George et al., 2021; Levenstein et al., 2024). At a high level, these models can learn to represent the structure of navigable task space, either physical or abstract, using predictive learning rules conditioned on observed state relationships in the environment. One popular example is the successor representation (SR), which was formalized in the field of computational reinforcement learning (RL; Dayan, 1993; Stachenfeld et al., 2017; de Cothi and Barry, 2020). In the RL framework (Dayan, 1993; Sutton and Barto, 2018), the SR encodes observed states in the environment (e.g., location) as the expected future states conditioned on the behavior of the agent or animal (Carvalho et al., 2024). The SR thus predicts future observations are likely to follow nearby in time based on the current state of an agent. The states need not be spatial but could be construed from any observations the agent or organism makes. In this type of predictive representation, each local, observed state contains predictive information about likely future states. Prima facie, state representations of the SR resemble place fields of principal neurons in the hippocampus (Stachenfeld et al., 2017). The SR also predicts a wide range of observations in nonspatial memory tasks, and its eigenvectors produce grid representations akin to those in the entorhinal cortex (Stachenfeld et al., 2017; Gershman, 2018). Furthermore, during goal-directed behavior in environments of changing shape, humans and rodents learn trajectories that are similar to artificial agents using an instantiation of the SR (Momennejad et al., 2017; de Cothi et al., 2022). Until recently, however, there were no direct comparisons of the SR to neural representations in the hippocampal formation, and direct empirical evaluations of its neural predictions have been lacking.

To address this gap, Lee et al. (2025) developed an approach with high-yield neural recording in freely behaving mice to directly compare neural representations in hippocampal subregion CA1 to theoretical model predictions of cognitive (or predictive) maps (see also Maisson et al., 2022). Using a geometrically changing environment, they first identified a reliable pattern of spatial mapping and remapping in CA1. To do so, Lee et al. (2025) leveraged a representational similarity-based approach to compare the structure of spatial mapping and remapping at the population-level in CA1 across animals and found that geometric changes drive a reliable pattern of remapping at the population-level in CA1. Next, they generated model predictions using true animal trajectories to learn successor features (SFs) from neurobiological inputs (de Cothi and Barry, 2020) and directly compared representations of SFs to those observed in CA1. Their results revealed that a specific range of model parameters accurately predicted the representations of CA1 populations in a geometrically changing environment.

Notably, there are several limitations of SR and SF models to explain neural representations; namely, they often assume preexisting localization (e.g., *x*–*y* state in an environment, direct place cell or boundary-vector cell inputs), discrete state spaces, or observability of environments not available to natural organisms. In recent work, Gomez et al. (2025) addressed these limitations using predictive learning in a neural network model from egocentric visual observations similar to those available to navigating animals. Specifically, this model learns to approximate eigenvectors of the graph Laplacian, which are equivalent to the eigenvectors of the SR and show grid-like representations akin to the entorhinal cortex, using an augmented Lagrangian Laplacian learning objective (ALLO; Machado et al., 2017; Gomez et al., 2024, 2025). More intuitively, the ALLO learning objective maps temporally adjacent sensory observations onto nearby states. In this sense, space is represented as a sequence of latent states, which is conceptually similar to recent alternative models for latent sequence representation (George et al., 2021; Raju et al., 2024). As described in Gomez et al. (2025), while the output of this model produces grid representations, the upstream representations in the neural network sequentially represent geometric and place-coding observed in the hippocampal system, matching the sequential representations suggested in earlier models of cognitive mapping (Barry et al., 2006; de Cothi and Barry, 2020). When they directly compared spatial representations from place-coding layers of the ALLO model to CA1 representations in Lee et al. (2025), they observed a close correspondence between these representations that could be learned directly from high-dimensional sensory observations in a geometrically changing environment. Importantly, this model could take any sensory observations and infer the predictive structure of the environment and thus offers a more generalized approach to representation learning based on the SR in both spatial and nonspatial settings.

These findings highlight the recent progress made in describing how the brain learns representations of changing environments through predictive learning of environmental structure (Gomez et al., 2025). When environmental structure changes, these representations adapt similarly to predictive neural representations of environmental structure observed in the hippocampus. Importantly, such flexible representations could significantly improve the ability of agents to transfer knowledge when environmental structure changes (Fig. 1*B*) and aid in the ability of animals to continue to learn in nonstationary settings wherein a model of the environmental structure does not have to be relearned but can instead be remapped similar to natural organisms. While a variety of conceptually related models that leverage different approaches for predictive learning of environmental structure have been described (Whittington et al., 2020; George et al., 2021; Levenstein et al., 2024), it remains unclear to what extent these models differ in the representations they learn in common task spaces, which should be explored in future work. However, these models do not presently address how the brain might relate actions to goals of the organism or agent. In natural environments, goal structure is often nonstationary

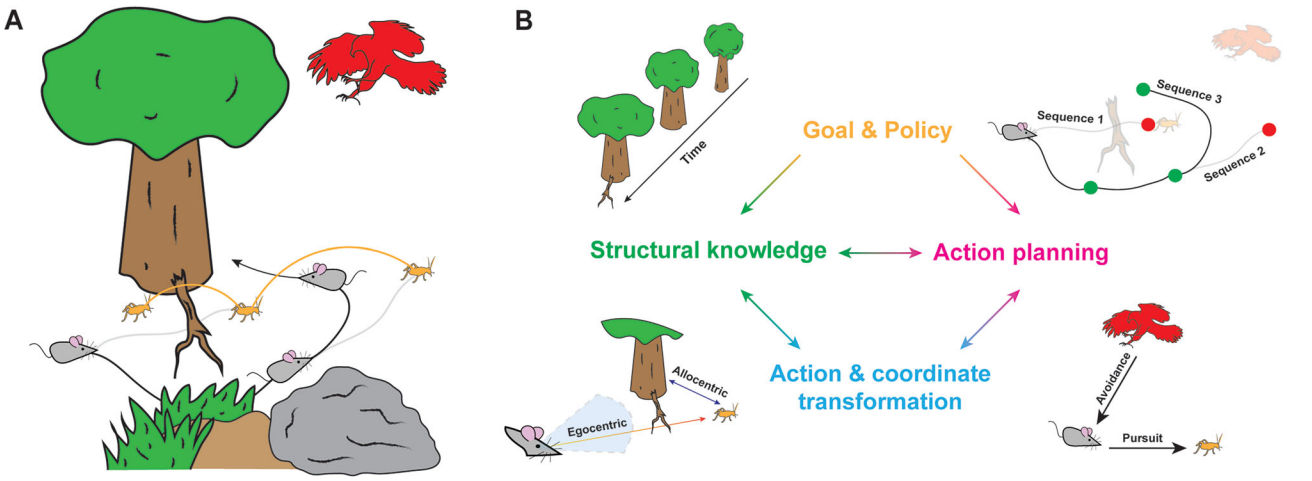


Figure 1. Navigating a changing world, with changing goals. ***A***, The diagram illustrates an example of a rodent navigating to pursue prey (an insect) in a natural environment and later avoiding predation (from an overhead bird of prey). ***B***, In such environments, features such as landmarks used for navigation are subject to change with time (e.g., a growing tree with a fallen branch), which can be represented through updates to structural knowledge of the environment. The goals of the organism are also highly dynamic and must switch from moment to moment in order to survive, such as between pursuit and avoidance, and require dynamic switching between the goals and policy of the navigator. To do so, the animal must infer their relationships between themselves and the environment (egocentric framing) and between features of the environment (allocentric framing) from available senses, necessitating action and coordinate transformation. An efficient strategy to adaptively navigate in this nonstationary setting is to construct sequences of actions between inferred states that can be rapidly updated according to the goals and policy of the navigator.

(Fig. 1*B*), such as during avoidance of threats, pursuit of prey or food caches, and adaptation to changing environmental conditions. To address these important gaps, new findings across several groups have begun to uncover how the neural representations in the brain enable organisms to pursue changing goals during navigation and how similar representations can be learned in machines.

Learning to Predict Dynamic Goals in Brains and Machines

As discussed, considerable evidence indicates that the hippocampus can learn a predictive map of the physical environment. Yet intelligent behavior also requires predicting features of the environment that change or move and using these predictions to guide behavior. How does the brain make predictions about dynamic features of the world? Doing so poses a spatial transformation problem (Fig. 1*B*): an animal must track a moving target relative to its own body (egocentrically) while anchoring its predictions in the stable coordinates of the world (allocentrically).

A long line of work has used predictive models to describe animal behavior during prey capture and escape, where many species adopt strategies consistent with internal models rather than purely reactive control (Ghose et al., 2006; Mischiati et al., 2015; Brighton et al., 2017; Evans et al., 2019). By tracking head movements in hunting dragonflies, Mischiati and colleagues demonstrated that insects infer the future position of a moving target rather than relying on a reactive strategy such as parallel navigation (Mischiati et al., 2015). Similarly, nonhuman primates playing a target-chasing video game employ a model-based predictive interception strategy that incorporates the physical dynamics of target movement (Yoo et al., 2020). In that study, neurons in the dorsal anterior cingulate cortex encoded multiple variables relevant to tracking the state of the target, including current and future target position, movement direction, speed, and acceleration. The insect and primate preparations, however, offer limited access to the underlying circuitry, and rodent models provide a complementary approach, offering the invasive recording and manipulation methods needed to study the ensemble activity that supports predictive interception.

Recent studies have therefore sought predictive behavior in mice and rats in two complementary roles: the rodent as prey and the rodent as pursuer (Fig. 1*B*). As prey, freely moving mice evading a predator robot in an enriched open field display "peeking" behavior around occluding objects, allowing them to scan for the predator's position while remaining under cover (Lai et al., 2024). Once the predator's location is determined, mice adjust their trajectory toward a static goal to avoid the robot. In the same task, mice also exhibit "baiting" behavior that appears to rely on stored knowledge that they can outrun the predator: they allow the robot to detect them, draw it back toward the arena's starting position, and then quickly run around it to reach their goal. In both cases, prediction supports route planning around both the moving predator and occlusions fixed within the environment. Route predictions based on the location of the predator and spatial affordances may rely on predictive representations observed in the hippocampal formation that reflect environmental structure (Stachenfeld et al., 2017; Danjo et al., 2018; Lee et al., 2025). These structural representations might further be used to generate nonlocal trajectories for planning in structured environments based on learned movement of targets and pursuers.

As pursuers, rats intercept moving visual targets for delayed reward (Alexander et al., 2022). Critically, targets moved either along pseudorandom paths or along "characteristic" paths that were stereotyped and anchored to the allocentric space of the environment. Rats took shortcuts on the characteristic paths, indicating that they learned to predict the target's trajectory through space. Because a shortcut requires anticipating where the target will be in world coordinates, this behavior requires a world-anchored, allocentric prediction, which again resembles SFs. This predictive interception emerged alongside multiplicative gain modulation of the self-motion tuning of neurons in the posterior parietal cortex (PPC). Relative to free exploration, this modulation lengthened the temporal window over which self-motion was encoded, improving the decoding of the animal's past, present, and future trajectory. Moreover, neurons whose self-motion tuning was gain-modulated tended to be sensitive to the egocentric position of the target during chasing, suggesting that multiplexing self-motion and spatial codes supports animals' ability to predict. Because predictions were observed on target paths anchored to the spatial environment, an interplay between egocentric representations in the neocortex and allocentric predictive representations in the hippocampus may be required (Bicanski and Burgess, 2018).

How do these changes in neural activity develop over learning, and how are they related to predictive behavior? A common approach for probing the computational principles of learning is to train recurrent neural networks (RNNs) on analogous tasks. In an early and influential example, Elman showed that RNNs trained to predict the next word in a sequence develop latent representations of grammar and lexical structure (Elman, 1990). When applied to spatial navigation, RNNs trained to path-integrate for self-localization develop latent representations resembling the place and grid cells of the hippocampal formation (Banino et al., 2018; Cueva and Wei, 2018). Just as those networks recover spatial codes from a self-localization objective, an RNN trained to pursue might be expected to recover the target-directed codes observed in PPC, providing a normative reason to treat such a network as a model of the biological circuit.

A recent study confirmed this intuition (Redman et al., 2026). A simple RNN received the target's velocity as input and generated its own velocity as output, with the objective of minimizing its distance to the target by the end of a fixed-length trial. The networks learned to pursue effectively and, notably, adopted predictive pursuit trajectories even though they were never explicitly trained to anticipate the target's future position, in contrast to the state-based learning signal used to predict environmental structure in previously discussed models. Latent representations resembling the egocentric target coding observed in PPC emerged in the networks, and ablating these units disrupted pursuit, offering a theoretical causal test that egocentric target position representations are required for prey capture. Predictive behavior emerged predominantly in high-rank networks, which also carried more information about the agent's position in allocentric space. Together, these results suggest that while low-dimensional, egocentric computations are sufficient for pursuit, predictive pursuit requires high-dimensional representations that also recover the allocentric frame.

The predictive models discussed thus far suggest that the temporal structure of a navigator's observations directly shapes neural representation. As a result, interactions between the navigator and their environment, which are determined both by the structure of navigable space (physical or abstract) and affordances of the organism or agent, should determine the structure of neural

representation. According to these models, the idiosyncrasies in actions of a navigator would have a central role in defining neural codes of higher-order brain structures that might implement predictive learning, which has been supported in recent experimental findings (Pecirno and Keinath, 2025). In changing environments, a useful neural code might not only predict environmental structure and proximity to goals but also how actions themselves relate to observed changes in the environment. In turn, we highlight new findings that identify action representation in high-order brain systems, including the hippocampus, and how such systems flexibly learn and adapt representations to dynamic environments and goals during navigation.

The Structure of Neural Activity Reflects Actions and Internal States

A central question arising from predictive models of navigation is what exactly neural activity is predicting. One possibility is that the brain simply predicts the sensory observations of a changing environment. Indeed, predictive coding has been extensively described in sensory systems, where neural activity reflects expected sensory inputs and prediction errors (Friston, 2010; Gavornik and Bear, 2014; Keller and Mrsic-Flogel, 2018). However, the hippocampus presents a more complicated picture. Hippocampal neurons are influenced by an enormous range of sensory and cognitive variables, including space, visual scenes, odors, sounds, as well as more abstract variables such as elapsed time, accumulated evidence, reward, and progress toward goals [see O'Keefe and Krupic (2021) for a detailed review]. Rather than encoding a single feature of the environment, hippocampal activity appears to flexibly incorporate whichever information is relevant for solving the current task. Yet this flexibility raises a fundamental question: what unifying principle determines which variables are predicted at any given moment?

An emerging view is that hippocampal activity is organized not around environmental variables themselves but around the behavioral significance of those variables. Neural tuning to sensory and internal cues develops only when those cues are behaviorally relevant or actively used by the animal (Zinyuk et al., 2000; Terrazas et al., 2005; Muzzio et al., 2009; Aronov et al., 2017; Nyberg et al., 2022). Moreover, fluctuations in attention and engagement profoundly alter the stability of hippocampal representations (Foster et al., 1989; Markus et al., 1995; Kentros et al., 2004; Fenton et al., 2010; Pettit et al., 2022). More recently, mice trained to switch between visual and odor cues to locate a hidden reward exhibited distinct hippocampal firing patterns according to the sensory modality that was behaviorally relevant, despite the sensory environment being identical across conditions (Radvansky et al., 2021). From this perspective, hippocampal activity reflects not the external world itself, but the animal's evolving interaction with it.

Direct evidence for this idea comes from Zutshi et al. (2025), who developed an auditory-guided navigation task in which an ascending auditory sweep instructed animals to navigate toward one of six reward locations. Hippocampal population activity progressively evolved toward the future goal before the animal reached the reward, closely following the accumulating auditory evidence. Critically, truncating the auditory sweep midway through the trial did not alter this predictive activity, despite eliminating the sensory cue that guided the decision. Thus, predictive hippocampal activity was more closely linked to the animal's impending action plan rather than to the sensory stimulus itself. This prediction was expressed at two complementary and interacting levels: hippocampal population activity evolved into distinct goal-specific states that identified the animal's selected destination, while transient nonlocal theta sequences briefly represented future locations ahead of the animal along the planned route.

Simple motor output alone, however, could not explain hippocampal activity. The goal-specific population state emerged as animals committed to a behavioral choice, whereas transient nonlocal sequences were particularly prominent during pauses and deliberation before commitment, consistent with previous observations of forward hippocampal sequences during vicarious trial-and-error (Johnson and Redish, 2007; Wikenheiser and Redish, 2015; Redish, 2016; Kay et al., 2020). This nonlocal predictive firing emerged not only during movement toward goals but also when animals deliberated or paused before committing to a choice, consistent with previous observations of forward hippocampal sequences during vicarious trial-and-error (Johnson and Redish, 2007; Wikenheiser and Redish, 2015; Redish, 2016; Kay et al., 2020). Moreover, nonlocal predictive firing on the auditory navigation task scaled with task complexity, weakening during habitual runs to known reward ports. Together, these findings suggest that hippocampal activity predicts upcoming behavior across multiple timescales, with stable population states reflecting the selected action plan within a specific environmental structure and transient theta sequences evaluating future trajectories within the current behavioral context. Thus, these two forms of predictive coding likely interact to support goal-directed behavior, wherein animals learn environmental structure through interaction, and can plan using nonlocal trajectories that predict upcoming action sequences.

This perspective also provides a unifying explanation for decades of observations on hippocampal remapping (Sanders et al., 2020; Fenton, 2024). Hippocampal spatial maps depend not only on the environment but also on how animals engage with it (Foster et al., 1989; Markus et al., 1995; Kentros et al., 2004; Fenton et al., 2010; Pettit et al., 2022). For example, place cells that are largely omnidirectional in open fields often become directionally selective during repeated traversals of a linear track (Markus et al., 1995; Battaglia et al., 2004) and are also predictive of upcoming turns on a T-maze (Frank et al., 2000; Wood et al., 2000; Ferbinteanu and Shapiro, 2003; Grieves et al., 2016). Consistent with this, Zutshi et al. (2025) found that hippocampal place cells remapped when animals switched from linear-track navigation and auditory-guided trials on the same track. Thus, behavioral context shaped the hippocampal map. Remarkably, returning to the original linear track after auditory trials did not reliably restore the linear-track map. Instead, map expression on the second block of linear-track trials depended on the animals' behavioral strategy on a trial-by-trial basis: trials in which mice ran directly to the end reward port reinstated the original linear-track map, whereas trials where they actively "searched" in intermediate ports reinstated the auditory-task map. Thus, hippocampal activity reflected the task as inferred and experienced by the animal rather than the task defined by the experimenter, indicating that hippocampal activity reflected the animal's internal beliefs about which task was currently being performed and therefore what future outcomes its actions were expected to produce. In approaching the problem of continual learning, these findings point to a potential solution whereby animals leverage distinct task representations to mitigate catastrophic forgetting through top–down switches between

behavioral policies based on the animal's internal belief about task context and current goals (Fig. 1*B*).

Together, these findings suggest that predictive hippocampal activity is organized not simply around future actions but around the behavioral strategies that generate those actions. The same movement may arise from exploitation, exploration, habits, or hesitation, and these distinct navigational modes produce different hippocampal dynamics despite similar sensory inputs and overt behavior. Preliminary observations from a dynamic patch-foraging task in the Zutshi laboratory further support this view, revealing distinct hippocampal population dynamics as animals transition between exploratory and exploitative strategies while navigating the same environment. Thus, rather than predicting sensory observations alone, hippocampal activity appears to predict future behavior within the context of the animal's current navigational strategy and predictions about the outcomes of her interactions with the environment. This view suggests that understanding neural firing requires understanding the behavioral algorithms that generate them—a principle that may also inform the design of artificial agents capable of continual learning.

Learning from Interactions in the Natural World

Many studies of adaptation and learning focus on understanding neural representations that arise after learning, using animals and agents that have been highly trained on particular tasks. As such, these studies often focus on the outcome of learning, and do not attempt to capture the dynamics of the learning process itself. As a result, the learning algorithms used to explain these data tend to be much slower than animals, requiring tens of thousands of reinforced action samples to converge on navigational strategies that animals learn with tens of reinforcements (Shamash et al., 2023). Perhaps equally striking is the observation that animals use highly structured trajectories to navigate space (Jiang et al., 2022), even during early learning (AbdelRahman et al., 2026; Gong et al., 2026), and they do not generate the random and diffuse trajectories produced by canonical learning agents. Taken together, this suggests that the evolving structure of animal behavior might provide important clues for building new algorithms that can explain the impressive learning efficiency of animals. Moreover, by capturing behavioral changes during learning, we might better constrain our search for the underlying neural mechanisms that mediate these changes.

Indeed, there is a long and rich history of work in neuroethology and ecology that studies the structure of natural behavior and its relationship to an animal's ecological niche (Tinbergen, 1963; McNamara and Houston, 2009), and there has been renewed interest in using natural behavior as a lens to study neural computation (Dennis et al., 2021; Ulanovsky, 2025). In parallel, a set of conceptually related ideas has been pursued from an information-theoretic perspective (Attneave, 1954; Barlow, 1961; Simoncelli, 2003) to study the structure of natural stimuli and their relationship to neural coding in the brain. This line of work, termed efficient coding, hypothesizes that neural systems make efficient use of their limited resources to maximize information about incoming signals. This hypothesis has been generalized to derive coding schemes that preserve predictive information about future stimuli (Tishby, Pereira, and Bialek, 1999; Palmer et al., 2015; Chalk et al., 2018) or that can be dynamically adapted based on predictions about future environmental contexts (Rao and Ballard, 1999; Mlynarski and Hermundstad, 2018, 2021). Tests of this hypothesis have been most successful in the context of sensory and perceptual systems (Laughlin, 1981; Attic and Redlich, 1992; Olshausen and Field, 1997; Simoncelli and Olshausen, 2001; Lewicki, 2002; Smith and Lewicki, 2006; Hermundstad et al., 2014; Gervain and Geffen, 2019), where it is possible to explicitly measure the statistical structure of natural stimuli and derive the architectures that efficiently encode these stimuli. This suggests that a similar approach might be successful when applied to behavior, where it is possible to measure the statistical structure of natural behavior and derive architectures that efficiently generate this behavior.

This perspective was adopted in two recent studies of navigational learning in flies and mice (Dan et al., 2024; AbdelRahman et al., 2026). Both studies used the statistical structure of animal behavior to derive parameterized control policies for generating actions; these policies, in turn, were used to explain how animals rapidly modified their behavior in new settings. In AbdelRahman et al. (2026), mice learned to intercept hidden spatial targets for a delayed reward. The spatiotemporal structure of mouse trajectories provides important clues for understanding this rapid learning: naive animals explore the arena using smooth trajectory segments delineated by pauses. Despite substantial variability in their spatial coordinates, these trajectory segments were well captured by structured changes in speed and heading over time. This suggested that by shifting, scaling, and composing generative functions that control instantaneous speed and heading, mice could flexibly plan and execute a rich space of navigational trajectories. The authors then designed agents that came pre-equipped with these behavior-derived generative functions, rather than needing to learn them de novo. With such a control policy, learning could then be viewed as the process of selecting appropriate parameters of these functions, rather than learning a set of elemental actions (see also Yttri and Dudman, 2016; Jiang et al., 2022). By combining this control policy with inference and active sampling, the authors were able to explain the near-one-shot learning observed in the fastest mice while also capturing how mice rapidly infer hidden targets (Butler et al., 2019; Jiang et al., 2022) and evade obstacles (Shamash et al., 2023).

In Dan et al. (2024), flying flies in virtual reality learned within minutes to avoid visual patterns that were associated with heat punishment. As in AbdelRahman et al. (2026), the authors used the spatiotemporal structure of fly behavior to construct a generative control policy that was tethered to a flexible internal goal heading (see also Stone et al., 2017; Goulard et al., 2021; Matheson et al., 2022; Mussells Pires et al., 2024; Westeinde et al., 2024). Importantly, by integrating insights from connectomic (Hulse et al., 2021) and physiological data (Seelig and Jayaraman, 2015; Fisher et al., 2019; Green et al., 2019; Kim et al., 2019), the authors showed how the functional form of this control policy might be hardwired in the structure of neural circuits to exploit the predictable consequences of actions in 1D rotational space, whereas the parameters of this policy could be flexibly modified based on experience via Hebbian-like plasticity at localized synapses within those circuits. This combination of hardwired circuit architectures and localized sites of plasticity could thus enable animals to balance speed and flexibility in new settings. More broadly, because there were signatures of this dissociation—between hardwiring and flexibility—that were apparent in the evolving structure of behavior itself, this approach is likely to generate insights across a wider range of behavioral learning studies, even when physiological and connectomic data are not available.

When taken together, these studies highlight the importance of accurately capturing an animal's control policy in explaining

the evolving dynamics of learning. In both studies, the behavior-derived control policies implicitly encode information about the consequences of movement in physical space; as such, they serve as inductive biases that speed learning in navigational settings. Moreover, the algorithmic form of these control policies provides clues for deciphering the neural circuit architectures that might implement them. In the fly, these links can be made precise using the connectivity and dynamics of specific cell types in the fly central brain. In the mouse, the shared similarity in the underlying algorithm—which relies on the same types of vector computations that support goal learning in flies—suggests that similar circuit mechanisms might be at play. Finally, these lines of work point to an exciting open direction for future research: in both studies, rapid learning was supported by a structured exploration of space, which differs from the random exploration typically adopted in canonical learning algorithms. By studying the statistical structure of animal behavior in natural settings, we might thus hope to uncover new insights about the structure of exploration and its role in rapid behavioral flexibility.

Where Should We Explore from Here?

While natural organisms have evolved to learn in the face of continual change, both in their environments and goals, the same problem remains a significant challenge for learning in artificial systems. Here, we have considered navigation as a useful context and constraint to examine how biological neural systems continue to learn in changing environments and how such observations can inform algorithms for more rapid, efficient, and continual learning in machines.

One common thread that connects this recent progress is that predictive learning rules can yield structured representations in artificial systems closely aligned to those observed in behaving animals with dynamic environments and goals (Gomez et al., 2025; Lee et al., 2025; Redman et al., 2026). However, this raises the question—what do animals and (should) agents learn to predict? New findings discussed here reveal that higher-order neural systems in natural organisms, such as the hippocampus, encode predictions of action sequences based on the internal beliefs of an organism about the environment, its structure, and the outcomes of context-specific actions (Zutshi et al., 2025). Together, these findings highlight the central role of prediction in the structure of neural activity of brain systems necessary for organisms to learn and remember to navigate in changing environments. Furthermore, the organization of an organism's interaction with the environment has a central role in shaping the observed structure of neural activity in these brain systems. A consideration for future and continued work on this topic is the extent to which predictive environmental structure and the consequences of actions are preconfigured in specific brain systems, rather than being learned de novo, and can be rapidly adapted to navigate in new contexts. A deeper understanding of how organisms' prior knowledge can be mapped onto novel sequential experiences in new environments could significantly aid rapid learning in machines.

Based on recent work discussed here, we have suggested that prediction plays a central role in shaping neural codes for navigation and that different forms of predictive learning interact to solve distinct computational problems (Fig. 2*A*). Namely, animals learn about the structure of their environments and behavioral affordances through interaction, which could be supported through multitimescale predictions of state observations and action plans, and during pursuit of goals leverage predictions of targets to learn reference frame transformations (Fig. 2*B*). While recent models discussed here illustrate the emergence of neural coding aligned with brain systems using predictive learning rules and have made breakthroughs in the adaptability of neural representations for flexible behavior and

A

Computational problem	**Structural knowledge**	**Action planning**	**Action & coordinate transformation**
What is predicted?	State observations (multisensory, abstract) "What future states am I likely to observe?"	Goal-conditioned action-state observations "Which actions should I take to reach my goal?"	Action-conditioned state observations "How will the state change if I take this action?"
Timescale	Long (10^1 to 10^2 seconds)	Multiscale (10^{-2} to 10^1 seconds)	Short (10^{-2} to 10^{-1} seconds)
Candidate brain systems	Hippocampus, entorhinal cortex	Prefrontal cortex, dorsal striatum, hippocampus, central complex (flies)	Posterior parietal cortex, retrosplenial cortex, sensorimotor cortices, central complex (flies)
Candidate algorithms	George et al. (2021), Gomez et al. (2025); Levenstein et al. (2024), Stachenfeld et al. (2017), Whittington et al. (2020)	AbdelRahman et al. (2026), Dan et al. (2024), Levenstein et al. (2024), Whittington et al. (2025)	Redman et al. (2026), Bicanski & Burgess (2018), Mussells Pires et al. (2024), Westeinde et al. (2024)

B

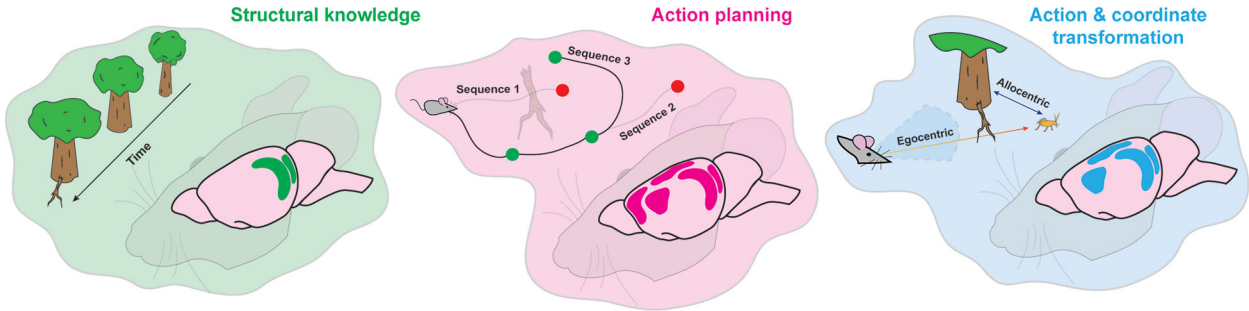


Figure 2. Candidate brain systems and novel algorithms for navigation in a changing world. ***A***, The table summarizes several core problems animals, and agents are confronted with during navigation in nonstationary environments, including learning about the structure of navigable space (physical or abstract), how to plan actions based on this structure, and how to transform actions and coordinate frames of reference with other, often moving objects and goals. The table proposes distinct forms of predictive objectives and timescales over which model results and experimental observations suggest each problem is resolved, along with candidate brain systems and algorithms that address these needs. ***B***, To illustrate the computational problems and respective brain systems that allow animals to address them, we highlight elements of the naturalistic navigation example from above and several (nonexhaustive) candidate brain systems from the table above.

rapid learning, future work should additionally focus on explicit tests of these model systems in continual learning problems. Continual learning specifically requires agents and animals to acquire new information while retaining (or even improving upon) prior task knowledge. While animals in the aforementioned tasks show impressive ability to flexibly switch between goal-directed behavior and neural representations, it remains unclear to what extent these new algorithms can achieve similar performance in changing navigational tasks when prior task knowledge is required. Recent work discussed here suggests that leveraging the statistics of animal navigation and the structure of neural representations as an inductive bias for such agents could aid in this effort.

An important gap in our knowledge that should be further explored is how natural organisms and artificial agents respectively interact with a changing world. Recent findings discussed here highlight that canonical agents do not explore their environment with the same efficient, structured behavior of natural organisms. Furthermore, leveraging the structure of actions of natural organisms as an inductive bias for navigation in artificial agents yields significant improvements in rapid learning during goal-directed navigation (Dan et al., 2024; AbdelRahman et al., 2026). Looking forward, an important area for future work to investigate further is how animals prioritize exploration of large and dynamic naturalistic environments, the role of such exploration in subsequent (goal-directed) learning and memory, and the structure of neural activity in relevant biological and artificial neural networks. Furthermore, how animals dynamically shift between behavioral contexts and neural representations supporting intelligent behavior is a large area of ongoing study. Uncovering the mechanisms natural organisms use to rapidly learn, update, and switch between relevant neural coding schemes and navigate changing environments could significantly inform the development of algorithms and architectures to better support similar behavior in machines, specifically in continual learning when earlier task knowledge should be maintained. In doing so, we suggest navigation poses a useful challenge and area of focus to better understand intelligent systems that continue to learn in the face of inevitable changes built into our natural world.

References

AbdelRahman NY, Jiang W-c, Coddington LT, Gong S, Dudman JT, Hermundstad AM (2026) Composing trajectories for rapid inference of navigational goals. Neuron 114:1–16.

Alexander AS, Tung JC, Chapman GW, Conner AM, Shelley LE, Hasselmo ME, Nitz DA (2022) Adaptive integration of self-motion and goals in posterior parietal cortex. Cell Rep 38:110504.

Aronov D, Nevers R, Tank DW (2017) Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. Nature 543:719–722.

Atick JJ, Redlich AN (1992) What does the retina know about natural scenes? Neural Comput 4:196–210.

Attneave F (1954) Some informational aspects of visual perception. Psychol Rev 61:183–193.

Bakhtiari S, Mineault P, Lillicrap T, Pack C, Richards B (2021) The functional specialization of visual cortex emerges from training parallel pathways with self-supervised predictive learning. Adv Neural Inf Process Syst 34: 25164–25178.

Banino A, et al. (2018) Vector-based navigation using grid-like representations in artificial agents. Nature 557:429–433.

Bányai M, Orbán G (2019) Noise correlations and perceptual inference. Curr Opin Neurobiol 58:209–217.

Barlow H (1961) Possible principles underlying the transformation of sensory messages. In: *Sensory communication* (Rosenblith WA, ed), pp 217–234. Cambridge, Massachusetts: MIT Press.

Barry C, Lever C, Hayman R, Hartley T, Burton S, O'Keefe J, Jeffery K, Burgess N (2006) The boundary vector cell model of place cell firing and spatial memory. Rev Neurosci 17:71–97.

Basu J, Nagel K (2024) Neural circuits for goal-directed navigation across species. Trends Neurosci 47:904–917.

Battaglia FP, Sutherland GR, McNaughton BL (2004) Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. J Neurosci 24:4541–4550.

Bicanski A, Burgess N (2018) A neural-level model of spatial memory and imagery. eLife 7:e33752.

Brighton CH, Thomas ALR, Taylor GK (2017) Terminal attack trajectories of peregrine falcons are described by the proportional navigation guidance law of missiles. Proc Natl Acad Sci U S A 114:13495–13500.

Butler WN, Hardcastle K, Giocomo LM (2019) Remembered reward locations restructure entorhinal spatial maps. Science 363:1447–1452.

Carvalho W, Tomov MS, de Cothi W, Barry C, Gershman SJ (2024) Predictive representations: building blocks of intelligence. Neural Comput 36:2225–2298.

Chalk M, Marre O, Tkačik G (2018) Toward a unified theory of efficient, predictive, and sparse coding. Proc Natl Acad Sci U S A 115:186–191.

Coddington LT, Lindo SE, Dudman JT (2023) Mesolimbic dopamine adapts the rate of learning from action. Nature 614:294–302.

Cueva CJ, Wei X-X (2018) Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In: *International Conference on Learning Representations*.

Dabney W, Kurth-Nelson Z, Uchida N, Starkweather CK, Hassabis D, Munos R, Botvinick M (2020) A distributional code for value in dopamine-based reinforcement learning. Nature 577:671–675.

Dan C, Hulse BK, Kappagantula R, Jayaraman V, Hermundstad AM (2024) A neural circuit architecture for rapid learning in goal-directed navigation. Neuron 112:2581–2599.e2523.

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

Dayan P (1993) Improving generalization for temporal difference learning: the successor representation. Neural Comput 5:613–624.

de Cothi W, et al. (2022) Predictive maps in rats and humans for spatial navigation. Curr Biol 32:3676–3689.e3675.

de Cothi W, Shipley S, Barry C (2026) Acetylcholine: a candidate substrate for hippocampal predictive learning? Nat Rev Neurosci 27:637–651.

de Cothi W, Barry C (2020) Neurobiological successor features for spatial navigation. Hippocampus 30:1347–1355.

Dennis EJ, El Hady A, Michaiel A, Clemens A, Tervo DRG, Voigts J, Datta SR (2021) Systems neuroscience of natural behaviors in rodents. J Neurosci 41:911–919.

Dohare S, Hernandez-Garcia JF, Lan Q, Rahman P, Mahmood AR, Sutton RS (2024) Loss of plasticity in deep continual learning. Nature 632:768–774.

El-Gaby M, Harris AL, Whittington JCR, Dorrell W, Bhomick A, Walton ME, Akam T, Behrens TEJ (2024) A cellular basis for mapping behavioural structure. Nature 636:671–680.

Elman J (1990) Finding structure in time. Cogn Sci 14:179–211.

Evans DA, Stempel AV, Vale R, Branco T (2019) Cognitive control of escape behaviour. Trends Cogn Sci (Regul Ed) 23:334–348.

Fenton AA, Lytton WW, Barry JM, Lenck-Santini P-P, Zinyuk LE, Kubik S, Bures J, Poucet B, Muller RU, Olypher AV (2010) Attention-like modulation of hippocampus place cell discharge. J Neurosci 30:4613–4625.

Fenton AA (2024) Remapping revisited: how the hippocampus represents different spaces. Nat Rev Neurosci 25:428–448.

Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. Neuron 40:1227–1239.

Fisher YE, Lu J, D'Alessandro I, Wilson RI (2019) Sensorimotor experience remaps visual input to a heading-direction network. Nature 576:121–125.

Flesch T, Saxe A, Summerfield C (2023) Continual task learning in natural and artificial agents. Trends Neurosci 46:199–210.

Foster TC, Castro CA, McNaughton BL (1989) Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. Science 244:1580–1582.

Frank LM, Brown EN, Wilson M (2000) Trajectory encoding in the hippocampus and entorhinal cortex. Neuron 27:169–178.

Friston K (2010) The free-energy principle: a unified brain theory? Nat Rev Neurosci 11:127–138.

Gavornik JP, Bear MF (2014) Learned spatiotemporal sequence recognition and prediction in primary visual cortex. Nat Neurosci 17:732–737.

George D, Rikhye RV, Gothoskar N, Guntupalli JS, Dedieu A, Lázaro-Gredilla M (2021) Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. Nat Commun 12:2392.
Gershman SJ (2018) The successor representation: its computational logic and neural substrates. J Neurosci 38:7193.
Ghose K, Horiuchi TK, Krishnaprasad PS, Moss CF (2006) Echolocating bats use a nearly time-optimal strategy to intercept prey. PLoS Biol 4:e108.
Gomez D, Bowling M, Machado MC (2024) Proper Laplacian Representation Learning. In: International conference on learning Representations. Available at: https://proceedings.iclr.cc/paper_files/paper/2024/hash/7c88ed03238db64cef6f23bd0d6b5d11-Abstract-Conference.html
Gomez D, Bowling M, Lee JQ, Machado MC (2025) Representations in the hippocampal-entorhinal system emerge from learning sensory predictions. *bioRxiv*, 2025.2010.2003.680189.
Gong S, Martell A, Dudman JT, Coddington LT (2026) Reward magnitude determines reinforcement learning efficiency. Science 392:eaeb0813.
Goulard R, Buehlmann C, Niven JE, Graham P, Webb B, Ayers J (2021) A unified mechanism for innate and learned visual landmark guidance in the insect central complex. PLoS Comput Biol 17:e1009383.
Green J, Vijayan V, Mussells Pires P, Adachi A, Maimon G (2019) A neural heading estimate is compared with an internal goal to guide oriented navigation. Nat Neurosci 22:1460–1468.
Grieves RM, Wood ER, Dudchenko PA (2016) Place cells on a maze encode routes rather than destinations. eLife 5:e15986.
Grima LL, Guo Y, Narayan L, Hermundstad AM, Dudman JT (2026) A global dopaminergic learning rate enables adaptive foraging across many options. Neuron 114:1–16.
Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J (2000) Modeling place fields in terms of the cortical inputs to the hippocampus. Hippocampus 10:369–379.
Hermundstad AM, Briguglio JJ, Conte MM, Victor JD, Balasubramanian V, Tkačik G (2014) Variance predicts salience in central sensory processing. eLife 3:e03722.
Hulse BK, et al. (2021) A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. eLife 10:e66039.
Jiang WC, Xu S, Dudman JT (2022) Hippocampal representations of foraging trajectories depend upon spatial context. Nat Neurosci 25:1693–1705.
Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J Neurosci 27:12176–12189.
Kay K, Chung JE, Sosa M, Schor JS, Karlsson MP, Larkin MC, Liu DF, Frank LM (2020) Constant sub-second cycling between representations of possible futures in the hippocampus. Cell 180:552–567.e25.
Keller GB, Mrsic-Flogel TD (2018) Predictive processing: a canonical cortical computation. Neuron 100:424–435.
Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER (2004) Increased attention to spatial context increases both place field stability and spatial memory. Neuron 42:283–295.
Kim SS, Hermundstad AM, Romani S, Abbott LF, Jayaraman V (2019) Generation of stable heading representations in diverse visual scenes. Nature 576:126–131.
Lai AT, Espinosa G, Wink GE, Angeloni CF, Dombeck DA, MacIver MA (2024) A robot-rodent interaction arena with adjustable spatial complexity for ethologically relevant behavioral studies. Cell Rep 43:113671.
Laughlin S (1981) A simple coding procedure enhances a neuron's information capacity. Z Naturforsch C Biosci C 36:910–912.
Lee JQ, Keinath AT, Cianfarano E, Brandon MP (2025) Identifying representational structure in CA1 to benchmark theoretical models of cognitive mapping. Neuron 113:307–320.
Levenstein D, Efremov A, Eyono RH, Peyrache A, Richards B (2024) Sequential predictive learning is a unifying theory for hippocampal representation and replay. *bioRxiv*, 2024.2004.2028.591528.
Lewicki MS (2002) Efficient coding of natural sounds. Nat Neurosci 5:356–363.
Luczak A, McNaughton BL, Kubo Y (2022) Neurons learn by predicting future activity. Nat Mach Intell 4:62–72.
Machado MC, Rosenbaum C, Guo X, Liu M, Tesauro G, Campbell M (2017) Eigenoption discovery through the deep successor representation. arXiv preprint arXiv:1710.11089.
Maisson DJ, Wikenheiser A, Noel JG, Keinath AT (2022) Making sense of the multiplicity and dynamics of navigational codes in the brain. J Neurosci 42:8450–8459.
Markus EJ, Qin Y-L, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J Neurosci 15:7079–7094.
Masset P, Tano P, Kim HR, Malik AN, Pouget A, Uchida N (2025) Multi-timescale reinforcement learning in the brain. Nature 642:682–690.
Matheson AMM, Lanz AJ., Medina AM, Licata AM, Currier TA, Syed MH, Nagel KI (2022) A neural circuit for wind-guided olfactory navigation. Nat Commun 13:4613.
Mathis MW (2026) Leveraging insights from neuroscience to build adaptive artificial intelligence. Nat Neurosci 29:13–24.
Mcclelland JL, Mcnaughton BL, Reilly RCO (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.
McCloskey M, Cohen NJ (1989) Catastrophic interference in connectionist networks: the sequential learning problem. In: *Psychology of learning and motivation*, Vol. 24 (Bower GH, ed), pp 109–165. San Diego: Academic Press.
McNamara JM, Houston AI (2009) Integrating function and mechanism. Trends Ecol Evol 24:670–675.
Meister M (2022) Learning, fast and slow. Curr Opin Neurobiol 75:102555.
Mischiati M, Lin H, Herold P, Imler E, Olberg R, Leonardo A (2015) Internal models direct dragonfly interception steering. Nature 517:333–338.
Młynarski WF, Hermundstad AM (2018) Adaptive coding for dynamic sensory inference. eLife 7:e32055.
Młynarski WF, Hermundstad AM (2021) Efficient and adaptive sensory codes. Nat Neurosci 24:998–1009.
Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ (2017) The successor representation in human reinforcement learning. Nat Hum Behav 1:680–692.
Moser EI, Moser M-B, McNaughton BL (2017) Spatial representation in the hippocampal formation: a history. Nat Neurosci 20:1448–1464.
Mussells Pires P, Zhang L, Parache V, Abbott LF, Maimon G (2024) Converting an allocentric goal into an egocentric steering signal. Nature 626:808–818.
Muzzio IA, Levita L, Kulkarni J, Monaco J, Kentros C, Stead M, Abbott LF, Kandel ER (2009) Attention enhances the retrieval and stability of visuospatial and olfactory representations in the dorsal hippocampus. PLoS Biol 7:e1000140.
Nejad KK, Anastasiades P, Hertäg L, Costa RP (2025) Self-supervised predictive learning accounts for cortical layer-specificity. Nat Commun 16:6178.
Neupane S, Fiete I, Jazayeri M (2024) Mental navigation in the primate entorhinal cortex. Nature 630:704–711.
Nyberg N, Duvelle E, Barry C, Spiers HJ (2022) Spatial goal coding in the hippocampal formation. Neuron 110:394–422.
Ofstad TA, Zuker CS, Reiser MB (2011) Visual place learning in *Drosophila melanogaster*. Nature 474:204–207.
O'Keefe J, Krupic J (2021) Do hippocampal pyramidal cells respond to nonspatial stimuli? Physiol Rev 101:1427–1456.
O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
Olshausen BA, Field DJ (1997) Sparse coding with an overcomplete basis set: a strategy employed by V1? Vision Res 37:3311–3325.
O'Reilly RC, Norman KA (2002) Hippocampal and neocortical contributions to memory: advances in the complementary learning systems framework. Trends Cogn Sci (Regul Ed) 6:505–510.
O'Reilly RC, Rudy JW (2001) Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychol Rev 108:311–345.
Palmer SE, Marre O, Berry MJ, Bialek W (2015) Predictive information in a sensory population. Proc Natl Acad Sci U S A 112:6908–6913.
Pecirno SA, Keinath AT (2025) Environmental representations in mouse hippocampal CA1 reflect the predictive structure of navigation. Curr Biol 35:6100–6113.
Pettit NL, Yuan XC, Harvey CD (2022) Hippocampal place codes are gated by behavioral engagement. Nat Neurosci 25:561–566.
Qu Y, Ou J, Pang L, Wu S, Luo Y, Behrens T, Liu Y (2026) Development of non-spatial grid-like neural codes tracks inference and intelligence. Cell 189:3091–3106.e3015.

Radvansky BA, Oh JY, Climer JR, Dombeck DA (2021) Behavior determines the hippocampal spatial mapping of a multisensory environment. Cell Rep 36:109444.
Raju RV, Guntupalli JS, Zhou G, Wendelken C, Lázaro-Gredilla M, George D (2024) Space is a latent sequence: a theory of the hippocampus. Sci Adv 10:eadm8470.
Rao RPN, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci 2:79–87.
Redish AD (2016) Vicarious trial and error. Nat Rev Neurosci 17:147–159.
Redman WT, Dinc F, Lin X, Chan MG, Alexander AS (2026) Predictive pursuit emerges in high-dimensional recurrent neural networks. *bioRxiv*, 2026.04.23.720457.
Samsonovich A, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. J Neurosci 17:5900.
Sanders H, Wilson MA, Gershman SJ (2020) Hippocampal remapping as hidden state inference. eLife 9:e51140.
Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599.
Seelig JD, Jayaraman V (2015) Neural dynamics for landmark orientation and angular path integration. Nature 521:186–191.
Shamash P, Lee S, Saxe AM, Branco T (2023) Mice identify subgoal locations through an action-driven mapping process. Neuron 111:1966–1978.e8.
Simoncelli EP (2003) Vision and the statistics of the visual environment. Curr Opin Neurobiol 13:144–149.
Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. Annu Rev Neurosci 24:1193–1216.
Smith EC, Lewicki MS (2006) Efficient auditory coding. Nature 439:978–982.
Solstad T, Moser EI, Einevoll GT (2006) From grid cells to place cells: a mathematical model. Hippocampus 16:1026–1031.
Stachenfeld KL, Botvinick MM, Gershman SJ (2017) The hippocampus as a predictive map. Nat Neurosci 20:1643–1653.
Stone T, Webb B, Adden A, Weddig NB, Honkanen A, Templin R, Wcislo W, Scimeca L, Warrant E, Heinze S (2017) An anatomically constrained model for path integration in the bee brain. Curr Biol 27:3069–3085.e11.
Sutton RS, Barto AG (2018) *Reinforcement learning: an introduction*, Ed. 2. Cambridge, Massachusetts: MIT Press.
Terrazas A, Krause M, Lipa P, Gothard KM, Barnes CA, McNaughton BL (2005) Self-motion and the hippocampal spatial metric. J Neurosci 25:8085–8096.
Tinbergen N (1963) On aims and methods of ethology. Zeitschrift für Tierpsychologie 20:410–433.
Tishby N, Pereira FC, Bialek W (1999) The information bottleneck method. In: Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing.
Ulanovsky N (2025) Natural Neuroscience: toward a systems neuroscience of natural behaviors. Cambridge, MA: MIT Press.
Vanchurin V, Wolf YI, Koonin EV, Katsnelson MI (2022) Thermodynamics of evolution and the origin of life. Proc Natl Acad Sci U S A 119: e2120042119.
Westeinde EA, Kellogg E, Dawson PM, Lu J, Hamburg L, Midler B, Druckmann S, Wilson RI (2024) Transforming a head direction signal into a goal-oriented steering command. Nature 626:819–826.
Whittington JCR, Muller TH, Mark S, Chen G, Barry C, Burgess N, Behrens TEJ (2020) The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. Cell 183:1249–1263.e1223.
Whittington JCR, McCaffary D, Bakermans JJW, Behrens TEJ (2022) How to build a cognitive map. Nat Neurosci 25:1257–1272.
Wikenheiser AM, Redish AD (2015) Hippocampal theta sequences reflect current goals. Nat Neurosci 18:289–294.
Wolf R, Heisenberg M (1991) Basic organization of operant behavior as revealed in Drosophila flight orientation. J Comp Physiol A 169:699–705.
Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27:623–633.
Yoo SBM, Tu JC, Piantadosi ST, Hayden BY (2020) The neural basis of predictive pursuit. Nat Neurosci 23:252–259.
Yttri EA, Dudman JT (2016) Opponent and bidirectional control of movement velocity in the basal ganglia. Nature 533:402–406.
Zador A, et al. (2023) Catalyzing next-generation artificial intelligence through NeuroAI. Nat Commun 14:1597.
Zinyuk L, Kubik S, Kaminsky Y, Fenton AA, Bures J (2000) Understanding hippocampal activity by using purposeful behavior: place navigation induces place cell discharge in both task-relevant and task-irrelevant spatial reference frames. Proc Natl Acad Sci U S A 97:3771–3776.
Zutshi I, Apostolelli A, Yang W, Zheng ZS, Dohi T, Balzani E, Williams AH, Savin C, Buzsaki G (2025) Hippocampal neuronal activity is aligned with action plans. Nature 639:153–161.