Fragmented Spatial Maps: State Abstraction and Efficient Planning from Surprisal

Mirko Klukas[1*], Sugandha Sharma[1], YiLun Du[2], Tomas Lozano-Perez[2], Leslie Kaelbling[2], and Ila Fiete[1*]

[1]BCS & McGovern Institute, MIT, Cambridge MA 02139, USA
[2]EECS & CSAIL, MIT, Cambridge MA 02139, USA
[*]{mklukas, fiete}@mit.edu

When animals explore spatial environments, their representations often fragment into multiple maps. What determines these map fragmentations, and can we predict where they will occur with simple principles? We pose the problem of fragmentation of an environment as one of (online) spatial clustering. Taking inspiration from the notion of a *contiguous region* in robotics, we develop a theory in which fragmentation decisions are driven by surprisal. When this criterion is implemented with boundary, grid, and place cells in various environments, it produces map fragmentations from the first exploration of each space. Augmented with a long-term spatial memory and a rule similar to the distance-dependent Chinese Restaurant Process for selecting among relevant memories, the theory predicts the reuse of map fragments in environments with repeating substructures. Our model provides a simple rule for generating spatial state abstractions and predicts map fragmentations observed in electrophysiological recordings. It further predicts that there should be "fragmentation decision" or "fracture" cells, which in multicompartment environments could be called "doorway" cells. Finally, we show that the resulting abstractions can lead to large (orders of magnitude) improvements in the ability to plan and navigate through complex environments.

Introduction

Contextual reorientation [1] and reanchoring, in which behavior, state estimates, or meaning are suddenly reevaluated based on new contextual information from the world,

are universal phenomena in psychology. One famous set of examples is the parsing of garden-path sentences such as "Time flies like an arrow, fruit flies like a banana" or "The woman brought the sandwich from the kitchen tripped" [2]. In the latter there is a sudden reorientation upon hearing the word tripped, so that *the woman* becomes the person *who was* brought the sandwich rather than the person bringing the sandwich. Similarly, spatial reorientation and reanchoring can occur when entering a building lobby from the outside or entering a different looking room from another one. Such reanchoring or reorientation events may constitute the basis on which the brain segments the continuous stream of experience into episodes or chunks that it uses to structure experience and memory [3–5].

In the brain, grid cells construct continuous 2-dimensional Euclidean maps of small environments [6] by the integration of self-movement cues as the animal explores the space, Fig. 1a. The advantage of such velocity integration-based Euclidean representations is that they provide a consistent encoding of seen and unseen locations and independent of paths taken to get there, making it possible to compute novel shortcut paths between locations [7–11].

However, between different environments, place and grid cells "remap": Representations of these environments involve different (if overlapping) sets of place cells and the spatial relationships between place cells in one environment are not preserved in the other [12]. Grid remapping is more subtle: grid cells exhibit coherent module-wide shifts that are differential across modules in their firing phases [13]. Remapping can be driven by non-spatial changes in context (e.g. changes in olfactory or visual cues within the same space) or by large spatial changes where the subject cannot easily determine its spatial displacement (e.g. after a journey in a closed vehicle).

These jumps in spatial representation, typically studied by discontinuously transplanting subjects from one environment to another or by switching non-spatial cues, can also occur when subjects smoothly navigate themselves within a single unchanging environment, particularly if it has many compartments or subregions [14, 15], Fig. 1b-c. This phenomenon – referred to as map fragmentation – is also a form of remapping. However, it is a distinct version of remapping different from the traditional use of the word and concept [12] because the environment remains stationary as the animal moves continuously through it, while in typical remapping experiments the environment changes [16]. We hypothesize that map fragmentation is a solution to multiple problems: First, it solves the problem of the accumulation of path integration errors that prevent the formation of consistent maps over larger spaces, resulting in the formation of smaller but consistent Euclidean maps. Thus, map fragmentation enables spatial inference and shortcut behaviors within each submap. Second, each submap represents a state abstraction in which contiguous locations are clustered together, and combining these abstractions with links between them can permit efficient and hierarchical representation and planning. Third, submaps can combine more globally to form a "topometric" map, a representation with enough expressiveness for topologically non-trivial cognitive spaces beyond real space, that preserves the advantages of both local metric structure and global hierarchy and abstraction.

Here, we propose a simple online rule for map fragmentation that avoids the large

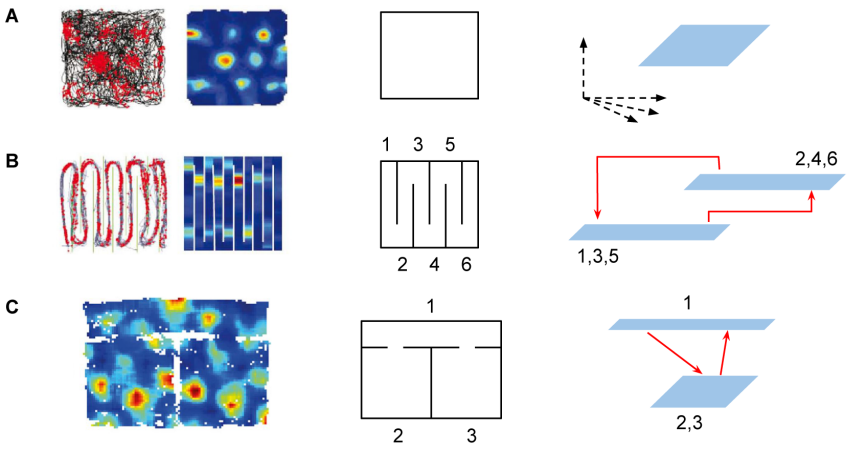

Figure 1: Map fragmentation in MEC. A-C: *Left,Middle.* Firing fields of grid cells in various environments, with environments illustrated schematically in the middle column (firing fields in A,B from [14] and in C from [15]). *Right.* Schematic map fragmentations of the environments. Blue regions are submaps, or regions with a continuous representation in the state space of the multi-module grid cell population. Solid arrows indicate discontinuous jumps in grid phase, which we interpret as transitions between submaps.

memory, time complexity and data-inefficiency of offline algorithms, and show that the resulting rule is a good potential model of map segmentations observed in grid and place cells. Finally, we demonstrate by implementing efficient random tree search algorithms that map fragmentation can facilitate efficient planning relative to using global maps, leading to a massive speed-up in complex and large environments without repeated substructures.

Results

Map fragmentation as clustering: an offline baseline

We propose that remapping across environments and fragmentation within environments can be considered to be a clustering problem: At each sampled location, the question is whether it should be categorized as a part of the most recently used map, or be assigned to a different one. A sensible answer would be that sufficiently "similar" locations should be assigned to the same map (cluster), while sufficiently different ones should be assigned to different maps (clusters).

We view a map as a (local) world model that enables the prediction of sensory inputs at any location within the map. Thus, we consider that a key metric for map fragmentation may be predictability or surprisal, Fig. 2. A similar metric has been used in robotics methods for simultaneous localization and mapping (SLAM) [17]. Specifically, sets of poses (locations and orientations) where the predictability of external observations remains high while moving between them ("contiguous regions") should be

clustered together into one map, Fig. 2. This view complements the use of other metrics that have been implemented in offline settings to construct spatial maps, including the graph Laplacian [18] and successor representation [19] methods, both of which use temporal proximity as their metric (indeed, under a random exploration policy, the successor representation is closely related to the graph Laplacian). Our primary focus here is on how biological and artificial agents might generate sensible maps in an online fashion. Secondarily, we use the metric of prediction or surprisal to generate these online fragmentations. In Discussion, we will consider how additional metrics can be used within the same online framework.

Define a model $\mathbb{P}(z' \mid x', x, z)$ that predicts the sensory input z' at pose x', based on the sensory input z at pose x (Fig. 2a,b; see Methods for details). The sensory observations and their predictions are given in terms of a range sensor centered on the agent, in the actual environment (Fig. 2a, right) or in a reconstructed map based on the observations z (Fig. 2a, middle), respectively. For each pose x, we delineate the surrounding region where predictability remains above threshold; this, by definition, is a contiguous region. We call the boundary of the region the prediction horizon for x. The radius of the prediction horizon varies depending on location within the environment, Fig. 2c. We can use the mutual surprise between poses, which we define as $-\frac{1}{2}(\log \mathbb{P}(z' \mid x', x, z) + \log \mathbb{P}(z \mid x, x', z'))$ (see Methods for details), as a measure of proximity that we illustrate with an Isomap embedding [20] of the environment (Fig. 2d). In this visualization, contiguous (high predictability) regions are compressed, while transition or bottleneck regions (low predictability) are stretched.

Finally, we define the average surprisal (see Methods) of a pose x by averaging over the mutual surprise of all nearby poses at a fixed Euclidean distance, and apply a clustering procedure similar to DBSCAN [21][1]. The procedure computes the connected components of all locations whose average surprise lies below a fixed threshold and decomposes the map into core fragments and transition regions, Fig. 2e. Additionally, in order to make an informed choice about the fragmentation threshold, we compute a contour tree (cf. [22]) of the surprisal values, which provides a visualization of how the connectivity ot space evolves with increasing thresholds, Fig. 2f. As we see, there are regions of the contour tree that are relative robust to the detailed threshold choice, providing similar connectivities over a range of threshold values.

The surprisal-based segmentations align well with both intuitive fragmentations and with neural data (cf. Fig. 1), suggesting that predictability may be a key and principled objective for map segmentation decisions.

However, the algorithm is offline, requiring full exploration of the space before it can generate the fragmented map. This is unlike in experiments, where animals generate map fragmentations in real-time as they explore an environment [15]; in non-spatial contexts too, there is evidence that event boundaries are defined in real-time [5, 23]. The algorithm also has high complexity, requiring fine spatial discretization and a large memory and computational buffer for the storage of and computation on the full predic-

[1]The density notion in DBSCAN is based on a count of neighbors. We use the average mutual surprise instead.

tivity matrix over all pairs of positions in the space. The same is true for Laplacian and successor matrix-based methods. Further, there is an additional gap between observed map fragmentations in biology and the latter two algorithms because while they provide multi-scale representations of the space (in the form of eigenvectors of some similarity matrix), they are not actually fragmentations of the environment, Fig. S3,S4.

Online fragmentation based on predictability: Our model

We next build a simple and biologically plausible online map fragmentation model based on surprisal, with the goal of generating fragmentations that are consistent with the principled offline clustering-based algorithm above. Our model is an agent that integrates its velocity as it explores an environment to update its pose estimate, and uses a short-term memory (STM) and a long-term memory (LTM) to make predictions about what it expects to see next.

The sensory observations for the online model consist of the activities of a population of cells that encode the presence of environmental boundaries at some distance, similar to boundary vector cells (BVCs) [24, 25] in entorhinal cortex or boundary-coding cells in the occipital place area [26]. These encode a binary, idiothetically-centered local view of the space[2] (with observation field-of-view angle ϕ), Fig. 3a,c. The velocity-based position estimates are represented by a population of idealized grid cells from multiple modules. For simplicity and to match the experimental setups in [14, 15] we assume that the pose angle is specified by a global orienting cue – effectively, the agent has access to its true head direction. The STM consists of an exponentially decaying moving average of recent observations, each shifted according to the internal velocity estimate of the agent, Fig. 3d. The STM is used to generate the prediction for the next observation (motivated by [17]), and a normalized dot product between the prediction and the current observation (BVC activity) yields our predictability signal (Fig. 3b,e). Due to its implementation as a moving average, STM activity slightly lags BVC activity. While high predictability is maintained along a trajectory, no fragmentation occurs. Once the predictability signal dips below a threshold, then at the first subsequent stabilization of spatial information, signaled by predictability returning to threshold, a fragmentation event is triggered (Fig. 3b). At this point, the agent must make a decision about which map to use, for which it uses its LTM.

The LTM consists of associations between the grid cell-encoded position representations and the sensory observations (filtered through the STM) encountered by the agent in the past (Fig. 3f). At a fragmentation event, the agent stochastically retrieves a previously visited state and corresponding location estimate in proportion to its match with the current obseration, or with some small constant probability (Fig. 3g), the agent selects (initializes) a new map, which corresponds to selecting a randomized new internal position representation by randomizing the set of phases across the grid modules.

The stochastic selection of an item from LTM based on overlap with the current observation serves two purposes simultaneously: first, an observation is likely to drive

[2]This is also known as a grid occupancy map in robotics [27].

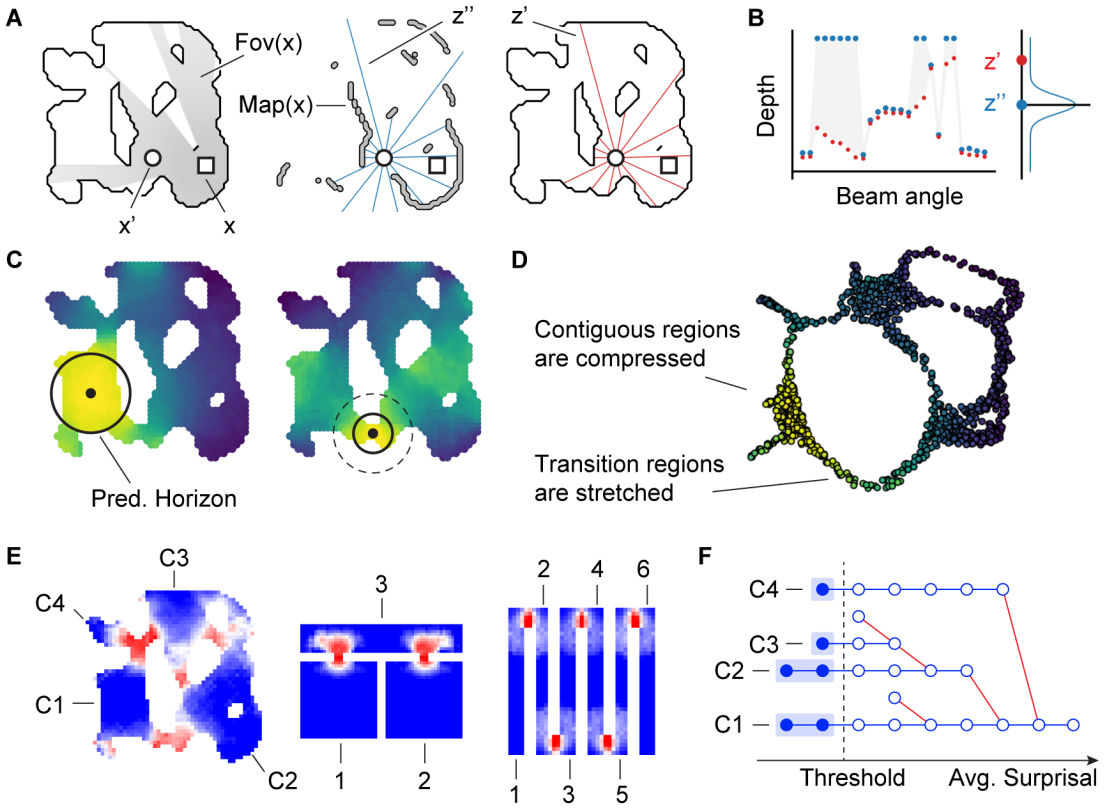

Figure 2: Fragmentations from predictability-based clustering. **A:** *Left.* Two agents (circle and square) in an organically shaped environment, with gray indicating the square agent's field of view. *Middle.* A set of range sensor observations (blue, z'') by the circular agent, shown on the map built from the square-agent's observations. *Right.* A set of range sensor observations (red, z') by the circular agent, shown on the actual environment. **B:** The blue observations in A constitute the square agent's prediction of the circular agent's measurements, with a prediction model given by a multivariate diagonal Gaussian with means given by the blue measurements, and which is evaluated at the vector of actual measurements in red. Greater vertical deviations between red and blue dots correspond to larger prediction errors. **C:** Each black dot represents a fixed reference location. All locations (pixels) in the map are colored by their predictability from the reference location. Solid black circle: the prediction horizon at that reference location: The horizon is large in open (contiguous) spaces, and small at bottleneck (transition) regions. **D:** An Isomap embedding of the environment based on mutual predictability gives rise to a warped embedding: distances are large when predictability is low. **E:** Locations in three environments, colored by their *average surprisal*. Numbered subregions correspond to connected components with average surprisal below a threshold level. **F:** A hierarchical clustering tree for the evolution of connected subregions for the first environment from E: Although the delineation of subregions depends on the choice of threshold, some subregions are relatively persistent and thus robust, maintaining their identity over a range of thresholds.

selection of a closely matching prior observation, and second, the retrieval of a previous observation is also proportional to the number of times that observation has been made

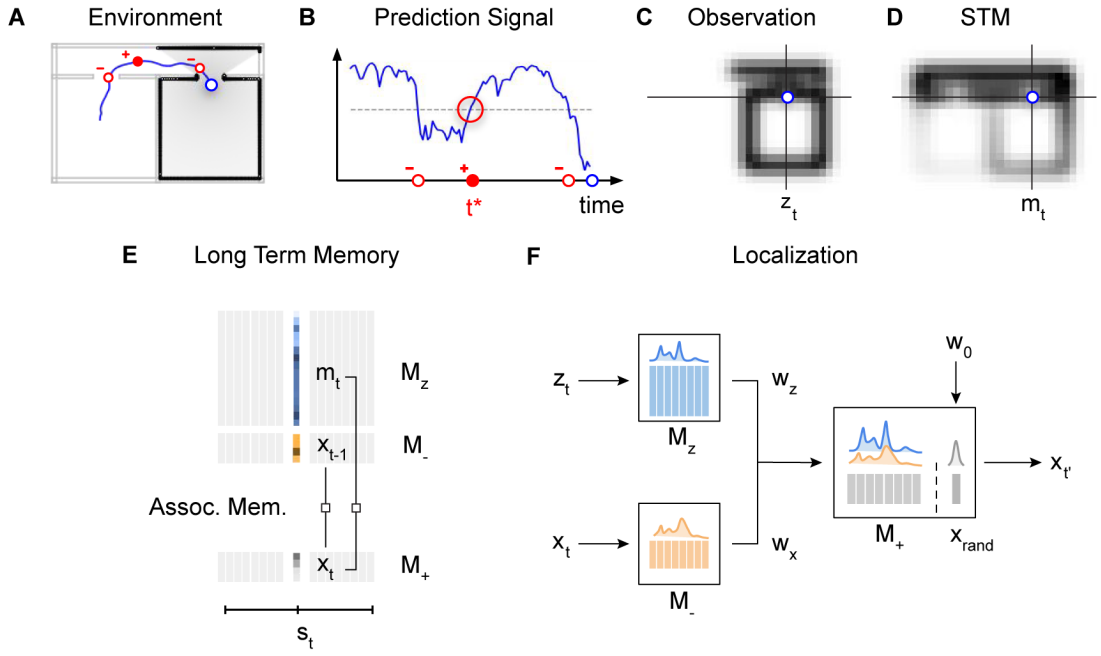

Figure 3: Online fragmentation model. **A:** Sample trajectory (blue) through a "2-room" environment highlighting three distinct events (red circles) that map to the events in B. Gray area: the agent's field of view at the agent's current location (blue circle). **B:** Prediction signal for the trajectory in A. Red dots indicate when the prediction crosses a threshold (dashed line), with the filled dot indicating a map fragmentation event (which occur at upward threshold crossings). **C** A snapshot of the current observation (z_t), which here consists of idiothetically centered BVC inputs to the agent (located at blue dot). Each pixel represents a BVC tuned to a location in space specified by the vector displacement between the agent and the pixel. Pixel intensity indicates the level of BVC activation. The image is cropped to exclude inactive BVCs. **D** Activity of the cell population encoding the STM (m_t) of recent observations for the trajectory from A. The STM slightly lags the current observation. Our prediction signal is derived from a normalized dot-product between the current observation and the STM. **E:** At each time t, the model stores associations between the STM m_t, the internal positional code x_t, and its predecessor x_{t-1} in a long-term memory (LTM). The slots s_t in the memory are chosen randomly. **F:** At a fragmentation event, the current observation z_t is compared to the LTM of observations M_z (blue columns), while the current position x_t is compared with the corresponding LTM of predecessor positions M_-. These comparisons result in weights w_z and w_x, respectively, and a memory is probabilistically selected for reuse from the LTM in proportion to its weight. Additionally, with a fixed probability, a random new map is initiated.

before, because stochastic selection from the set of past observations is a form of monte carlo volume estimation. In short, the selection of a submap after a fragmentation decision enables the reuse of existing submaps to represent new spaces when relevant based on similarity and frequency of past observations, while simultaneously permitting the creation of new maps. The frequency-dependence of this process together with the possibility to create new maps is similar to the Bayesian nonparametric Chinese

Restaurant Process (CRP) [16, 28, 29]; the observational similarity component makes it more akin to the distance dependent CRP (dd-CRP) [30, 31]. However, a key difference is that our observations are only implicitly clustered into submaps: each observation and location pair is stored independently of the rest in the LTM without a submap assignment, with submap boundaries only defined by the existence of a fragmentation decision and a jump in the grid-encoded spatial locations for the post-fragmentation observation relative to the immediate pre-fragmentation observation [3].

Further maintaining a "temporal" LTM which memorizes spatial transition probabilities, and using this information to bias the selection of a map at fragmentation events stabilizes how an environment is fragmented, though it is not critical (see SI, Fig. S6). Spatial transitions contain valuable information about the relationships between individual map fragments and are important for exploiting their hierarchical structure in route planning, as we illustrate later.

Fragmented maps in multiple environments

We explore the map fragmentations generated by our online model across organically shaped and previously experimentally tested structured multi-compartment environments, Fig. 4. The online model generates fragmentations at locations that correspond to observation bottlenecks, including at doorways or narrow openings and around the corner of sharp turns, Fig. 4a-c (top). Starting from the first trajectory through the space and across multiple trajectories, the remapping or fragmentation points and the selected maps are consistent, evidence of the robustness and reliability of online fragmentation (Fig. 4a-c, bottom). In the two-room and hallway environment, the model generates a fragmentation in which the two rooms are each represented by the same local map (rather than a single global map), and these maps are distinct from the map for the hallway. Moreover, the fragmentations generated by the online model are consistent with the fragments from the principled baseline method, Fig. 4d. This model can be used to generate fragmentations and predicted grid cell tuning curves for arbitrary enviromental geometries; we do so for model cells from different grid modules in two environments, Fig. 4e,h (fragmentations of more environments, including a square sprial maze and a simple linear track, shown in SI, Fig. S1a-c). If the angular field is of view ϕ is restricted rather than omnidirectional, the maps also acquire direction tuning, Fig. 4h and Fig. S1c.

[3] A temperature hyperparameter controls the degree of noise in the selection of a submap from LTM. This stochastic process allows us to not only use the degree of similarity but also the frequency of similar observations in selecting a submap: it performs a stochastic measurement of the volume of similar observations (submap occupancy), and then stochastically selects a map on that basis, without keeping an explicit count of how often each submap has been visited in the past. Thus, we may call this process a doubly-stochastic dd-CRP.

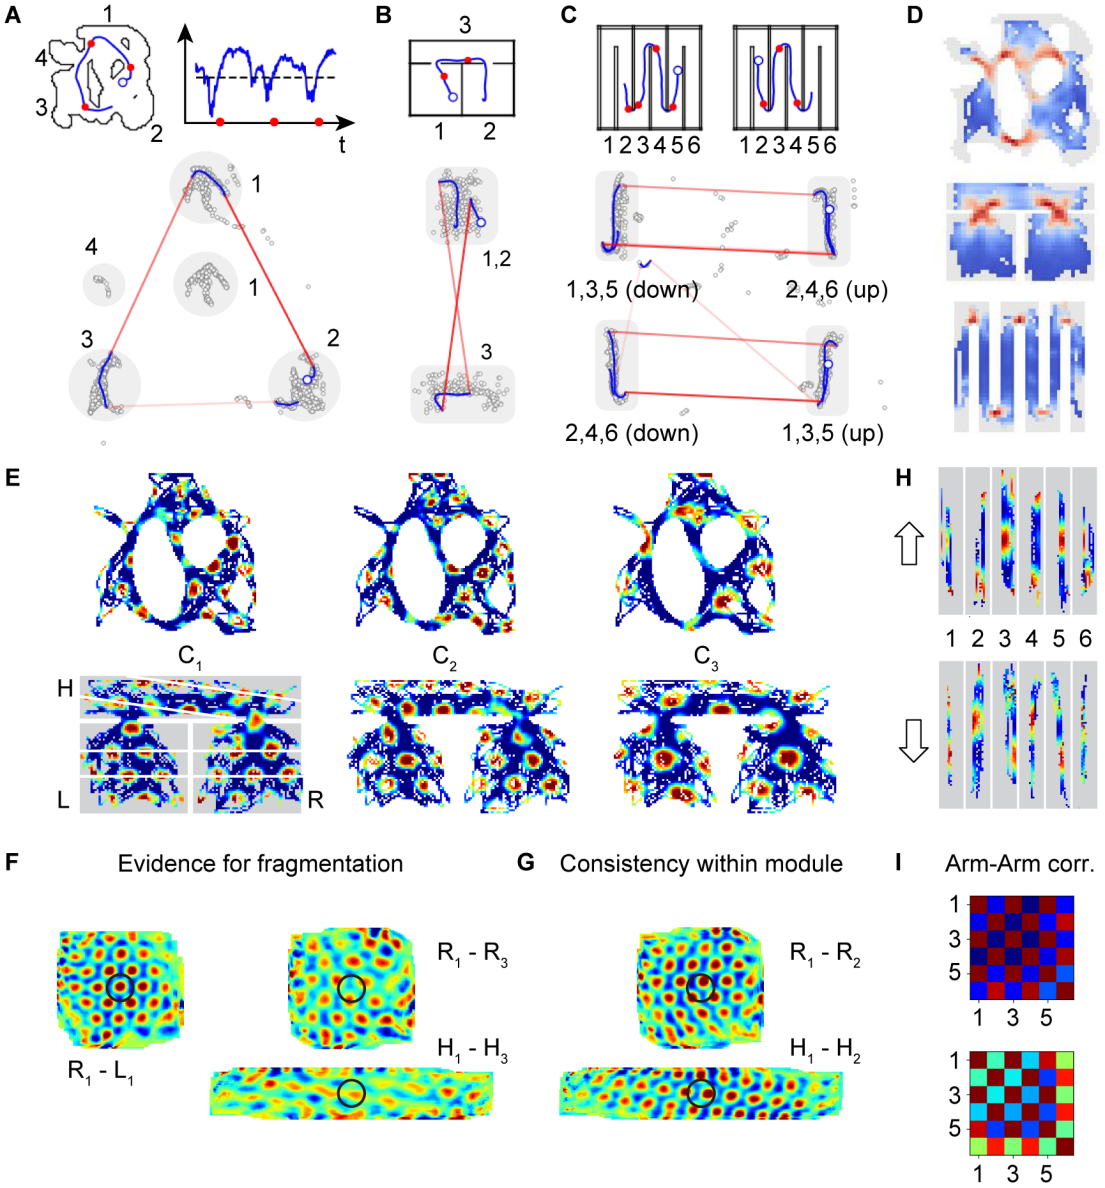

Figure 4: Online fragmentation results. A–C *Top:* Example trajectory snippets through three different environments (blue line), with current agent position indicated as an open blue circle. Fragmentation events indicated by red dots. The numbers indicate subregions identified in Figure 2E. A additionally shows the corresponding predictibility signal (blue) and threshold (dashed line). *Bottom:* The same trajectory snippets traversing the internal coding state space. The gray circular areas highlights the most-visited parts of state space for each environment, and the numbers correspond to the mapped area in the environment. Discontinuous jumps in state space, corresponding to transitions between submaps which occur at fragmentation events, are plotted in red. **D:** Heat maps indicating the density of online fragmentation events closely match the offline predictibility-based clustering fragmentations from Fig. Figure 2E. These maps can be interpreted as tuning curves of "fragmentation decision cells" or "doorway cells". **E:** Firing fields of three simulated grid cells in two distinct environments. C1 and C2 are from a common module; C3 is from a distinct module of larger scale. **F:** *Left.* Spatial cell-cell correlation of a single cell indicating that the tuning curves of cell 1 in both rooms (L and R) coincide and a common map is used for both rooms. *Right.* Spatial cell-cell correlations of cells from distinct modules. The cross correlation changes across different compartments showing that a relative phase shift must have occured (modules remap simultaneously and independently). **G:** Cross correlation for two cells from a common module in each of the compartments is presented, showing that the cross correlations are preserved across map fragments. **H,I:** Directionally tuned firing fields in the "Hairpin" maze (shown in C) of an idealized grid cell. The difference in firing fields for consequtive arms shows that the arms are mapped to different parts in mapping space depending on the direction they are traversed. The matrices show the correlation coefficients comparing signals of different arms.

Coherence of fragmentation across scales and maintenance of cell-cell relationships

Two key structural predictions of our model are, first, that the map fragmentations are consistent and coherent across scales (across grid modules), with all cells and modules remapping at the same spatial location in an environment. This is in contrast with eigenvector-based models [18, 19, 32], in which there is no specific or coherent remapping decision that is made across eigenvectors, Fig. S3,S4, Fig. 6.

Second, in our model all grid cells within each module maintain fixed cell-cell relationships across map fragments and environments. This too is in direct contrast with eigenvector-based models, Fig. 4g,S5, Fig. 6. Consistent with our model, grid cell data and analyses reveal that the pairwise relationships between co-modular grid cells remain stable across environments [33] and states [34, 35] even when place cells remap and their relationships change. These neural data are inconsistent more generally with models in which grid cell responses are derived from place cell responses [32, 36] because they would predict altered cell-cell relationships when place cells remap [34].

Efficient planning with fragmentations

Next, we quantify the functional utility of map fragmentation in a navigational planning problem. The fragmented maps, which represent a form of state abstraction, decompose the planning problem hierarchically, into a family of smaller and simpler sub-problems. Thus, they are expected to make planning more efficient. We perform computational

experiments to illustrate this point, comparing a bi-level navigation algorithm in the fragmented map with a simple baseline.

Consider a goal-directed problem in which agents, who have previously mapped the space, are tasked with finding a path to a cued goal location from a start location. For planning, we will assume that the LTM containing stored observation-location associations also includes an explicit submap identification (that is, all observations until a fragmentation event are assigned the same submap ID; at a fragmentation event, if the retrieved map has not yet been assigned a submap ID, a new submap ID is initiated and added to the LTM and associated with all subsequent observations until the next fragmentation event, and so on; all the observations between fragmentations are fused using local displacement information to form a submap for the whole fragment) and storing submap transitions. The environments are complex, but *without* repeating submap structure (Fig. 5a-b, d-e), because the fragmented representations generated by our simple agent do not distinguish between difference spaces with the same appearance (no global odometry assumed across submaps).

The baseline (global) agent is furnished with a global map, which includes ground-truth position informaton for all observations (Fig. 5a,d) and uses the Rapidly-exploring Random Tree (RRT) algorithm [37] to find a path through the space (Methods). The agent using a fragmented approach constructs a graph in which the nodes correspond to the submaps, and the edges correspond to observed transitions between submaps during exploration. It performs a depth-first search through the transition tree to find the sequence of submaps that lead to the node containing the target location (determined by querying the LTM with the target inputs). Within each submap, the agent uses the RRT algorithm to plan a path between the locations corresponding to the entry and exit edges. This agent possesses no global positional information.

In the environment of Fig. 5a-b, routes are found vastly more rapidly with fragmented maps than without: we see a $\sim$ 5-fold speedup. The relative advantage of planning with fragmented maps grows superlinearly with the complexity and size of the environment and separation between start and end locations within these spaces, Fig. 5c (right; steps are a proxy for the problem complexity).

Next, we simulate agents moving through 3D photorealistic virtual apartments in which observations are rich pixel images with range data, Fig. 5d, g. We apply convolutional visual recognition networks to the dense inputs to extract sparse landmarks and use these to generate online map segmentations (Methods). As before, the agent performs bi-level planning on the tree of transitions with submaps as nodes and RRT planning within submaps. Here we find a several orders of magnitude speedup in planning with map fragmentation, Fig. 5f.

Discussion

Relationship to existing work

Existing models of neural map fragmentation fall into two categories, Fig. 6: The first assumes that fragmentation is driven by large path integration errors that cause a mis-

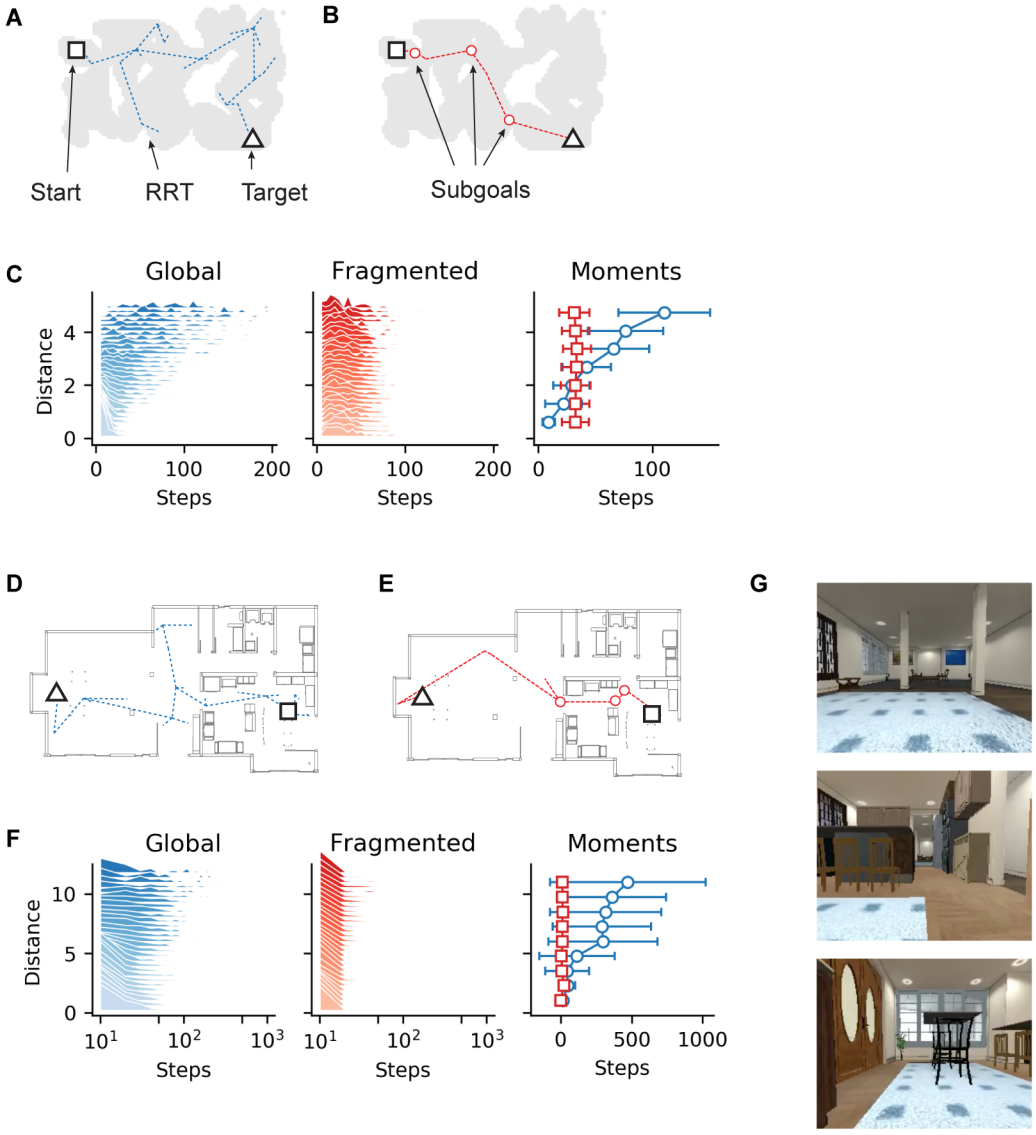

Figure 5: Efficient Hierarchical Planning with Map Fragmentations. A-C: Results of the planning algorithm applied to the global map of the environment (A) and to the hierarchical maps from our online fragmentation model (B). In C we plot the distributions of planning steps conditioned on the distance between start and target locations. **D-G:** Similar to A-C but using an alternative fragmentation algorithm based on semantic information extracted from visual inputs as shown in G.

match between estimated position and familiar observations [38]; in environments with little ambiguity in the external sensory cues, there would be no fragmentation. The second considers eigenvectors of different types of transition matrices, e.g. eigenvectors of the graph Laplacian of the adjacency matrix or of the successor representation [18,

19]. These models require a global buildup of the transition matrix between all pairs of locations in the environment, requiring in some sense a complete map of the environment before any fragmentations could happen. They also do not provide an explicit fragmentation of the environment, but rather a global map whose eigenvectors may be interpreted as the tuning curves of grid cells, and a subset of which appear visually fragmented (with different fragmentations by different grid cells and modules). By contrast, our model is fully online and provides explicit, robust fragmentations starting from the first trajectory through an environment, even in the absence of positional ambiguity; it also requires a smaller memory demand than transition matrix models. Finally, the model involves only simple, biologically plausible computational elements, with grid cells and BVCs and a short-term and a long-term memory, to explain a number of experimental results.

Our work, initially motivated by the empirical observations of fragmented maps in neuroscience, is closely related to work on segmented maps in the field of simultaneous localization and mapping (SLAM) in robotics [17, 39–41]. The main difference is that predictions in our model are based on a temporally limited window into the past, provided through the STM, whereas in [17] *all* observations are accumulated into a map that the prediction is based on. Further, our predictions are based on idiothetically-centered local views of the environment (BVC) – which are not assembled into a global allocentric map – and use an adapted moving average as a STM. For (re-)localization we use local views stored in a spatially indexed LTM.

As we have shown, spatial abstraction and spatial hierarchy in the form of map fragments can be of high utility in efficient search for solving goal-directed problems. State abstractions and hierarchical representations are broadly recognized to be important for more efficient reinforcement learning as well, and implemented in different forms including the classic options framework and more recent attempts [42, 43]. A key challenge for such approaches is to find rules that generate appropriate state abstractions, especially those capable of doing so in an online or streaming way. Our work is a contribution in this direction; related work includes the generation of temporal abstractions based on novelty rather than surprisal [44].

Our use of a surprisal signal is closely related to curiosity-based algorithms for reinforcement learning [45]. These algorithms use prediction error as an internal reward, to drive agents to explore unknown parts of the space. By contrast, we use prediction error as a way to generate state abstractions.

Model extensions: a broader set of metrics for fragmentation

The general principle of online state abstraction through online map fragmentation can use metrics in addition to surprisal for triggering a fragmentation event. Consider the case of two hallways with similar ideothetically-centered views, e.g. hallways 2 and 3 in Fig. 4c, that differ only in the permitted turn direction at the end. A natural extension of the model would be to incorporate a cell population encoding navigational affordances, to fragment and select maps based not only sensory surprisal but also on the set of actions that can be or are commonly taken. Other extensions include using the physical distance

between states [18, 19], the passage of time [46–48] with a dynamic (temporally decaying) threshold for fragmentation that makes fragmentation more likely as time elapses (also see [17]), the appearance unique or novel visual features including landmarks [44, 49–51], and sufficient mismatch in the estimates of state made from different cues or sensory modalities [38, 52], in addition to the metric of perceptual predictibility that we have used here and that the hippocampus has been shown to be sensitive to [53, 54]. The present model, which provides an online method for generating meaningful abstractions, may be applied with arbitrary combinations of these metrics to generate fragmentations influenced by multiple factors.

Merging of maps

In case of prolonged experience in the two compartment environment, map fragmentations tend to merge into a single, continuous representation that covers both compartments [15]. In our model some map fragments can, because of the stochastic nature of the fragmentation process, occasionally extend beyond an expected fragmentation boundary (see Fig. S1d). These events occur sparsely and are unlikely to be the source of the merging of maps observed in [15]. We expect the merging of maps to result from an improvement of the prediction signal with more experience, which can be modeled by allowing the prediction system to use not just recent observations from short term memory, but also past observations from long term memory. Exploring the dynamics of this process is an interesting potential extension of the present work.

Role of map fragmentation for general cognitive representation

Our model of online fragmentation of a continuous stream of experience enables the representation of a very general class of maps – including in spatial and non-spatial cognitive domains – in a way that exceeds the capabilities of a "pure" grid code. Grid cells generate Euclidean representations of Euclidean spaces [11]. Fragmented maps can each be viewed as separate local Euclidean "charts", mapped out by a multi-modular grid code, that are then associated to each other through transitions learned in the hippocampus according to the global layout of the charts. In other words, the combination of fragmented maps and the transitions between them can be viewed as a topological atlas [55] or topometric map [56], that can represent highly non-Euclidean structures while also permitting locally metric computations.

Thus, from a general perspective, map fragmentation and remapping (reanchoring) on cognitive representations can be viewed as faciliating the step from representing flat Euclidean space to representing richer manifolds. In combination with grid cells' ability to represent high-dimensional variables [11], such a coding scheme becomes highly expressive.

In contrast to the approach taken in [57, 58] there is no need to generate entirely new neural codes and representations to fit the local statistics of the explored space. Instead, we propose that the neural codes seen within submaps retain their native structure across spaces, in the form of a pre-formed and stable recurrent scaffold for memory through

grid cells. Even though grid cell representations in each module are 2-dimensional, theoretically the set of modules an represent even high-dimensional continuous spaces [11], while potential non-Euclidean aspects of cognitive varaibles can be captured by the between-submap transitions. This structuring of memory into continuous parts with preexisting scaffolds [59–63] together with occasional transitions between these continuous chunks simultaneously provides rapid learning and flexibility.

Episodic memory

Episodic memory, one example of a general cognitive representation, deserves special discussion because of the privileged role of the hippocampal system in its creation, storage, and use [64, 65]. Like spatial map fragmentations, episodic memory involves partitioning or clustering the continuous stream of temporal experience into chunks that involve similar perceptual, temporal, and contextual elements [3–5, 66]. Our proposal for surprisal-based spatial segmentations could be applied to study memory chunking. Interestingly, the memory for non-spatial items has also been shown to segment based on changes in spatial context, specifically by passage through doorways [67], as would be predicted by the present model.

The utility of applying our model first in the spatial domain is that it yields concrete predictions that are quantifiably consistent with observed map fragmentations. Applying it across cognitive domains will contribute to a unified computational model for how the hippocampal formation generates structured memories of spatial and non-spatial cognitive experience [57, 58, 65, 66, 68, 69].

Experimental Predictions

The decision to form a new map fragment in our model depends only on recent observations that are filtered through a STM, without requiring global information about the enironment. Thus, map fragmentations are predicted to occur in real time and on the very first pass through relevant regions of new environments, consistent with experimental results in the spatial and non-spatial domains [3–5, 70]. Further, in our model, all grid cells and grid modules undergo map fragmentation simultaneously, at the same time and location along a given trajectory, unlike in other models (Fig. S3,S3) [18, 19].

Fragmentations tend to occur at spatial bottlenecks that limit the prediction horizon, which correspond to "doorways" in the environment. The current evidence for cells firing at doorways is mixed [71, 72]. However, the necessity for a neural correlate that communicates the fragmentation decision and facilates across-module grid realignment under a fragmentation event predicts the existence of "fragmentation decision cells" or "doorway cells" whose tuning curves would resemble the heatmaps of Fig. 4d.

A common theme in MEC seems to be that cells with spatially structured tuning coexist with vector versions of themselves: i.e., cells that have similar tuning curves but are offset by a fixed vector (e.g. BVCs [24, 25] and landmark or object vector cells [73, 74]). In this light we might also expect "fragmentation vector cells" or "doorway vector cells" that fire if the rodent is at a fixed angle and distance from a fragmentation

location. These cells, which could be interpreted as encoding future action affordances or future map transitions, would faciliate planning.

Next, the model predicts that the stochastic process of generating map fragmentations can result in more than one map for the same region even when there is not an explicit manipulation of context or task. There are at least two implications of this result. First, it suggests that variations in the firing of grid and place cells on different visits to a location might be due not only to variable paths taken within a single map [75] but to the retrieval of entirely different maps. Second, these multiple, stochastically generated maps might subsequently be easy to harness for contextual differentiation, for instance like "splitter" cells [58, 76–78].

Finally, the large efficiencies in planning and goal-directed navigation afforded by the use of fragmented maps suggest that neural planning should exhibit hallmarks of the fragmentation process: If theta phase precession or waking neural replay events [79–85] correspond to planning [86–88], we should expect them to exhibit punctate trajectories with hierarchical dynamics between versus within fragments in multicompartment environments.

Methods

The source code will be made available online upon publication or request.

Offline fragmentations from predictability and surprise

We approximate the probabilistic observation model $P(z' \mid x', x, z)$ by

$$P(z' \mid x', x, z) = \int P(z' \mid x, m)\ P(m \mid x, z)\ \mathrm{d}m \approx P(z' \mid x', m^*_{x,z}).$$

Here $m^*_{x,z}$ is the maximum a posteriori estimation of an occupancy map given by an inverse sensor model as described in [27], and $P(z' \mid x, m)$ is the respective range sensor model. More precisely: Given a deterministic range sensor that takes measurements along a fixed number ($n = 1000, 1500$) of simulated beams, whose angles are chosen at equally spaced angles from the interval $[-\pi, \pi]$, we take three depth measurements z, z', and z''. The first two are taken in the actual environment at their resepective poses x and x', whereas the third is taken on a map $m^*_{x,z}$ built from the initial measurement z made at x. The observation model $P(z' \mid x, m)$ is then defined as a multivariate diagonal Gaussian with constant diagonal entries $\sigma = 1.0$ and mean z'', Fig. 2a,b.

The function underlying the distance matrix used for the Isomap embedding (cf. Fig. 2d) is given by the *mutual surprise* $s(x, x')$ between two poses x, x' which we define as

$$s(x, x') := -\tfrac{1}{2}\Big(\log P(z \mid x, x', z') + \log P(z' \mid x', x, z)\Big).$$

We refer to the negative mutual suprise as *mutual predictability*. With this in hand we define the *average surprise* $s(x) = s(x; r, \varepsilon)$ of a pose by averaging over the mutual

surprise about all poses at a fixed distance. To be more precise we define

$$s(x) := \frac{1}{|\Delta(x)|} \sum_{\Delta(x)} s(x, x'),$$

where $\Delta_x = \Delta_x(r, \varepsilon)$ is the set of all poses whose distance to x lies in the range $[r-\varepsilon, r+\varepsilon]$, for some previously fixed $r, \varepsilon > 0$; in our experiments we use $r \approx 0.4$ and $\varepsilon \approx 0.05$ depending of the minimal distances between poses. We sometimes refer to the negative avergage mutual suprise as *contiguity*. Informally, a high contiguity implies fewer suprises in direct proximity of the current pose and thus a low urge to remap.

To extract map fragmentations we uniformly sample poses from the environment and compute their avergage surprise, Fig. 2e. We then consider only those poses whose surprise lies below a previously fixed threshold (chosen acordingly for each environment). To make an informed choice about the threshold we compute a discrete contour tree [22] of the poses with respect to the average surprise visualizing the evolution of the connectivity with respect to increasing thresholds, Fig. 2f. The connected components of the subthreshold region yields a fragmentation into sub-maps, one for each connected region, and a suprathreshold transition region. We consider two poses to be connected if their Euclidean distance is below another previously fixed threshold that depends on the coverage of the environment by all the pose samples.

	Fragmentation w/o Sensor Ambiguity	Fragmentation on 1st Visit	Memory Demand	Coherent Fragm.	Cell–Cell corr. preserved
Spatial codes from BVC	**No**	**YES**	**n**	**YES**	**YES**
Laplacian & SR	**YES**	**No**	**n^2**	**No**	**No**
Our model	**YES**	**YES**	**n**	**YES**	**YES**

Figure 6: Comparison to other models. Our model improves on the potential shortcomings of other mapping approaches.

Online fragmentations from predictability

Observations and internal mapping locations

As before, our observation model is given by a range sensor that takes measurements along a fixed number of simulated beams. The beams' angles are chosen at equally spaced angles from the interval $[\theta_t - \phi/2, \theta_t + \phi/2]$. Here θ_t denotes the head direction at time t and $\phi = 360°, 270°$ defines the field of view of the agent; cf. Fig. 3a. We convert these range measurements into the activity z_t of a simulated population of boundary

vector cells by a binning process; cf. Fig. 3a,c. In our model we assume there is a $n \times n$ array of BVCs covering an area of $w \times w$, with $n = 91, 111$ and $w \approx 4$m.

We assume that internally locations are represented by a population of idealized grid cells of mulitple scales. For ease of computation, we interpret this multi-module grid code as a high capcity code for an unfolded 2-dimensional space [11]; cf. Fig. 4a-c (bottom). The Poisson rate maps f_c for an idealized grid cell c are then generated from superposing three cosinusoidal waves, each offset by an angle of 60°, over the unfolded 2-dimensional grid space, i.e.

$$f_c(x) := \tfrac{4\pi}{\sqrt{3}\lambda_c} \sum_{k=0}^{2} \Big[1 + \cos\big(\langle R_c(x - x_c), \mathrm{e}^{-ik\frac{\pi}{3}} \rangle\big)\Big].$$

Here λ_c and x_c encode the lattice scale and its offset, and R_c is a rotation matrix defining the orientation of the lattice.

Short term memory:

The short term memory (STM) is defined as an adapted exponential moving average of BVC activity:

$$m_t := \alpha \cdot \frac{\hat{z}_t}{\|\hat{z}_t\|} + (1 - \alpha) \cdot \frac{z_t}{\|z_t\|},$$

where the prediction

$$\hat{z}_t := \mathrm{shift}(m_{t-1}, -\lambda v_{t-1})$$

is a shifted version of the 2d-array m_{t-1} with respect to the scaled velocity of the agent. We found that a smoothing parameter $\alpha \approx 0.9$ works well. The scaling parameter $\lambda = \frac{n}{w}$ maps from the environment into pixel space. The shift of the BVC array results in a diffused version of the array caused by shifts with non-integer values. The extent of diffusion depends on the resolution (or number) of the BVCs.

Prediction model and fragmentation events

The prediction model is a normalized dot product of the current BVC observation z_t with the prediction $\hat{z}_t$ computed from the STM as described above:

$$\mathbb{P}(z_t \mid m_{t-1}, v_{t-1}) \;\propto\; \frac{\mathrm{vec}(z_t)}{\|\mathrm{vec}(z_t)\|} \cdot \frac{\mathrm{vec}(\hat{z}_t)}{\|\mathrm{vec}(\hat{z}_t)\|}^{\top}$$

where $\mathrm{vec}(z)$ is the unfolded version of a 2d-array z; cf. Fig. 3b-d. A fragmentation event is triggered after the prediction signal $\mathbb{P}(z_t \mid m_{t-1})$ recovers from falling below a fixed threshold θ ($\approx 0.9, 0.925$) and rises above again; cf. Fig. 3b. The normalization and the fact that both z_t and $\hat{z}_t$ are non-negative ensures that the prediction score lies within the intervall $[0, 1]$.

Long term memory and relocalization

The LTM is implemented as a matrix $M \in \mathbb{R}^{n \times S}$ whose s'th column is given by the concatenation of the internal position estimate x_{t_s}, its predecessor x_{t_s-1} and the state m_{t_s} of the STM at the time t_s the entry was written to memory, i.e. we have (cf. Fig. 3e)

$$M_{*,s} := \big(\mathrm{vec}(x_{t_s})\ \mathrm{vec}(x_{t_s-1})\ \mathrm{vec}(m_{t_s})\big)^\top.$$

We fill the memory as follows: At each time step t we choose a slot (column) s_t in the memory and replace the corresponding entry with the new one. Until we reach capacity, that is as long as $t \leq S$, we set $s_t = t$, after that the slots s_t are chosen uniformly at random – similar to the associative memory in [89]. Thus, the LTM consists of two associative memories: one storing assiciations between locations and observations, and the other storing state transitions. Alternatively, one could store associations with the actual observations z_{t_s} and not the filtered observations from the STM m_{t_s}, but we found that the associations with the STM work better and result in more stable fragmentations. The same is true when we restrict the capacity of the memory; cf. Fig. S6. Note that the LTM also maintains a temporal memory storing transitions (x_{t_s-1}, x_{t_s}) for each entry in the memory. We use a memory size S between 2000 and 6000.

To determine the new location during a fragmentation event we query the LTM and compute two distinct weight vectors w_1 and w_2. The first encodes how well a given observation z fits any of its entries and is given by

$$w_1(s, z; M) = \theta_M\big(\mathrm{vec}(m_{t_s}) \cdot \frac{\mathrm{vec}(z)}{\|\mathrm{vec}(z)\|}^\top\big).$$

With slight abuse of notation we denote by θ_M the function that sets all values below a certain threshold θ_M to $-\infty$. For ease of notation we set $\mathrm{e}^{-\infty} := 0$ – this becomes relevant in the probability computation below. We usually set this threshold to be equal to the fragmentation threshold $\theta \approx 0.93$. In order to allow for more flexibility during the above lookup we query the LTM not only with the actual observation z, but also with observations shifted by small pixel offsets δ, i.e. with $z_\delta = \mathrm{shift}(z, \delta)$ instead of just z, where $\delta \in \mathbb{Z}^2$ is chosen from a small region Δ around the origin. If a shifted observation fits a particular entry in the memory better, we replace the corresponding entry in the computed weight vector w_1. Then, if one of these adjusted entries, s say, is chosen during a remapping event we do not remap exactly to the associated position x_{t_s} but adjust it proportional to the respective offset δ and remap to $x_{t_s} + \frac{1}{\lambda}\delta$ (recall that λ translates from environment to pixel coordinates).

The second vector serves as a bias towards map transitions that have already been traversed and is given by

$$w_2(s, x; M) = \exp\big(-\frac{\|x - x_{t_s-1}\|^2}{0.25}\big).$$

Note that we use the Euclidean norm between two 2d vectors out of computational convenience, but we could have used the dot product of their corresponding multi modular

grid codes as well. Finally, when a fragmentation event is triggered we sample a new location from

$$P(x \mid z; M) \;\propto\; \mathrm{e}^{\tau \cdot w_0} \cdot P_0(x) \;+\; \sum_s \mathrm{e}^{\tau \cdot (w_1(s,z) + w_2(s,x))} \cdot \delta(x_{t_s}),$$

where $P_0(x)$ is a distribution over the space of possible locations, $w_0 = 1$ the concentration parameter, and $\tau = 1, 10, 20$ is the inverse temperature of the model; cf. Fig. 3g.

Trajectories

The trajectories are generated by choosing waypoints in the environment uniformly at random and navigating toward the next waypoint along a perturbed shortest path at a mean speed of 20 cm/sec. Time is discretized into steps of size $\Delta t = 0.1$ sec.

Hierarchical Planning

We apply Rapidly-exploring Random Trees (RRT) [37] to find a path between randomly chosen pairs of start and target positions, Fig. 5a. Next, we run our online segementation algorithm to get an environment-fragmentation into submaps and form a topological graph whose vertices and edges correspond to submaps and transitions respectively. For each map-fragment we superpose all its associated memories (STM-filtered BVC activity) and threshold this newly formed representation to form an occupancy grid map (in the sense of [27]) which we can apply the path planning algorithm to. We exploit the hierarchical structure by first finding a path of transitions in the topological graph, using a breadth-first search, and reduce the overall planning task into a family of sub-problems as follows: Each transition into- and out of a node defines a pair of local entry and exit postions on the submap associated with the traversed node defining a smaller planning problem that can be solved more effectively, Fig. 5b. In Fig. 5c we plot the distances between start and goal locations against the number of planning steps needed. The algorithm underlying the results in Fig. 2d–g is given as follows. Because the 3D environments involve dense observations of pixel-rich data, we add image processing and observation sparsification steps in the form of landmark identification. The agent receives RGB-D images as input, removes the floor plane, and segments the resulting point cloud. It retains as landmarks the large segments that are not vertical walls, which are generally large furniture items that are both relatively static and easy to recognize robustly from new viewpoints. As it moves through the environment, fragments are defined as follows: Starting at the initial location, the current fragment is defined a set of two visible landmarks, and the region of space from which both those landmarks remain in view constitutes the set of spatial locations assigned to that fragment. Whenever the agent moves into a part of the space where one or both of those landmarks are not visible, and if the current location does not correspond to any existing fragment, it starts a new fragment. Each fragment is connected topologically to the fragment it entered from.

Arm-arm correlation

The correlation matrices in Fig. 4i are computed as follows. For each arm in Fig. 4h we produce a 1-dimensional signal by averaging over the x-axis of the respective tuning curves in each arm. Each entry c_{ij} in the matrix is then given by the Pearson correlation coefficient of the 1-dimensional signals in arm i and j.

References

1. Julian, J. B., Keinath, A. T., Marchette, S. A. & Epstein, R. A. The Neurocognitive Basis of Spatial Reorientation. *Curr Biol* **28,** R1059–R1073 (Sept. 2018).
2. Levy, R., Reali, F. & Griffiths, T. L. *Modeling the Effects of Memory on Human Online Sentence Processing with Particle Filters* in *Proceedings of the 21st International Conference on Neural Information Processing Systems* (Curran Associates Inc., Vancouver, British Columbia, Canada, 2008), 937–944.
3. Swallow, K. M., Zacks, J. M. & Abrams, R. A. Event boundaries in perception affect memory encoding and updating. *J Exp Psychol Gen* **138,** 236–57 (May 2009).
4. Ezzyat, Y. & Davachi, L. What constitutes an episode in episodic memory? *Psychol Sci* **22,** 243–52 (Feb. 2011).
5. Baldassano, C. *et al.* Discovering Event Structure in Continuous Narrative Perception and Memory. *Neuron* **95,** 709–721.e5 (Aug. 2017).
6. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. Microstructure of a spatial map in the entorhinal cortex. eng. *Nature* **436,** 801–806 (2005).
7. Whishaw, I. Q., Hines, D. J. & Wallace, D. G. Dead reckoning (path integration) requires the hippocampal formation: evidence from spontaneous exploration and spatial learning tasks in light (allothetic) and dark (idiothetic) tests. eng. *Behav Brain Res* **127,** 49–69 (2001).
8. Blum, K. I. & Abbott, L. F. A model of spatial map formation in the hippocampus of the rat. eng. *Neural Comput* **8,** 85–93 (1996).
9. Burak, Y. & Fiete, I. R. Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol* **5,** e1000291 (Feb. 2009).
10. Banino, A. *et al.* Vector-based navigation using grid-like representations in artificial agents. *Nature* **557,** 429–433 (May 2018).
11. Klukas, M., Lewis, M. & Fiete, I. Efficient and flexible representation of higher-dimensional cognitive variables with grid cells. *PLoS Comput Biol* **16,** e1007796 (Apr. 2020).
12. Colgin, L. L., Moser, E. I. & Moser, M.-B. Understanding memory through hippocampal remapping. *Trends Neurosci* **31,** 469–77 (Sept. 2008).
13. Fyhn, M., Hafting, T., Treves, A., Moser, M.-B. & Moser, E. I. Hippocampal remapping and grid realignment in entorhinal cortex. eng. *Nature* **446,** 190–194 (2007).

14. Derdikman, D. *et al.* Fragmentation of grid cell maps in a multicompartment environment. *Nat Neurosci* **12,** 1325–1332 (Oct. 2009).

15. Carpenter, F., Manson, D., Jeffery, K., Burgess, N. & Barry, C. Grid cells form a global representation of connected environments. *Curr Biol* **25,** 1176–82 (May 2015).

16. Sanders, H., Wilson, M. A. & Gershman, S. J. Hippocampal remapping as hidden state inference. *Elife* **9** (June 2020).

17. Fairfield, N., Wettergreen, D. & Kantor, G. Segmented SLAM in three-dimensional environments. *Journal of Field Robotics* **27,** 85–103 (2010).

18. Machado, M. C., Bellemare, M. G. & Bowling, M. H. A Laplacian Framework for Option Discovery in Reinforcement Learning. *CoRR* **abs/1703.00956** (2017).

19. Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat Neurosci* **20,** 1643–1653 (Nov. 2017).

20. Tenenbaum, J. B., Silva, V. d. & Langford, J. C. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science* **290,** 2319–2323 (2000).

21. Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. *A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise* in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining* (AAAI Press, Portland, Oregon, 1996), 226–231.

22. Edelsbrunner, H. & Harer, J. *Computational Topology - an Introduction.* I–XII, 1–241 (American Mathematical Society, 2010).

23. Zacks, J. M. *et al.* Human brain activity time-locked to perceptual event boundaries. *Nat Neurosci* **4,** 651–5 (June 2001).

24. Barry, C. *et al.* The boundary vector cell model of place cell firing and spatial memory. eng. *Rev Neurosci* **17,** 71–97 (2006).

25. Lever, C., Burton, S., Jeewajee, A., O'Keefe, J. & Burgess, N. Boundary vector cells in the subiculum of the hippocampal formation. eng. *J Neurosci* **29,** 9771–9777 (2009).

26. Julian, J. B., Ryan, J., Hamilton, R. H. & Epstein, R. A. The Occipital Place Area Is Causally Involved in Representing Environmental Boundaries during Navigation. *Curr Biol* **26,** 1104–9 (Apr. 2016).

27. Thrun, S., Burgard, W. & Fox, D. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)* (The MIT Press, 2005).

28. Aldous, D. J. *Exchangeability and related topics* in *École d'Été de Probabilités de Saint-Flour XIII — 1983* (ed Hennequin, P. L.) (Springer Berlin Heidelberg, Berlin, Heidelberg, 1985), 1–198.

29. Schaeffer, R., Bordelon, B., Khona, M., Pan, W. & Fiete, I. Efficient online inference for nonparametric mixture models. *Proceedings of Machine Learning Research* **(to appear)** (2021).

30. Blei, D. M. & Frazier, P. I. Distance Dependent Chinese Restaurant Processes. *Journal of Machine Learning Research* **12,** 2461–2488 (2011).

31. Gershman, S. J., Frazier, P. I. & Blei, D. M. Distance Dependent Infinite Latent Feature Models. *IEEE Trans Pattern Anal Mach Intell* **37,** 334–45 (Feb. 2015).

32. Dordek, Y., Soudry, D., Meir, R. & Derdikman, D. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife* **5,** e10094 (Mar. 2016).

33. Yoon, K., Buice, M., Barry C.and Hayman, R., Burgess, N. & Fiete, I. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* **16,** 1077–84 (Aug. 2013).

34. Trettel, S., Trimper, J., Hwaun, E., Fiete, I. & Colgin, L. Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nat Neurosci* **22,** 609–617 (Apr. 2019).

35. Gardner, R. J., Lu, L., Wernle, T., Moser, M.-B. & Moser, E. I. Correlation structure of grid cells is preserved during sleep. *Nat Neurosci* **22,** 598–608 (Apr. 2019).

36. Kropff, E. & Treves, A. The emergence of grid cells: intelligent design or just adaptation? *Hippocampus* **18** (2008).

37. LaValle, S. Rapidly-exploring random trees : a new tool for path planning. *The annual research report* (1998).

38. Cheung, A. Probabilistic Learning by Rodent Grid Cells. *PLoS Comput Biol* **12** (2016).

39. Kuipers, B. The spatial semantic hierarchy. *Artificial intelligence* **119,** 191–233 (2000).

40. Bosse, M. *et al. An atlas framework for scalable mapping* in *2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422)* **2** (2003), 1899–1906.

41. Kuipers, B., Modayil, J., Beeson, P., MacMahon, M. & Savelli, F. *Local metrical and global topological maps in the hybrid spatial semantic hierarchy* in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004* **5** (2004), 4845–4851.

42. Sutton, R. S., Precup, D. & Singh, S. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* **112,** 181–211 (1999).

43. Bacon, P.-L., Harb, J. & Precup, D. *The Option-Critic Architecture* in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (AAAI Press, San Francisco, California, USA, 2017), 1726–1734.

44. Şimşek, Ö. & Barto, A. G. *Using relative novelty to identify useful temporal abstractions in reinforcement learning* in *Twenty-first international conference on Machine learning - ICML '04* (ACM Press, 2004).

45. Pathak, D., Agrawal, P., Efros, A. A. & Darrell, T. *Curiosity-Driven Exploration by Self-Supervised Prediction* in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (July 2017).

46. Sahay, A., Wilson, D. A. & Hen, R. Pattern separation: a common function for new neurons in hippocampus and olfactory bulb. *Neuron* **70,** 582–8 (May 2011).

47. Denny, C. A. *et al.* Hippocampal memory traces are differentially modulated by experience, time, and adult neurogenesis. *Neuron* **83,** 189–201 (July 2014).

48. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16,** 264–6 (Mar. 2013).

49. Knight, R. Contribution of human hippocampal region to novelty detection. *Nature* **383,** 256–9 (Sept. 1996).

50. Antunes, M. & Biala, G. The novel object recognition memory: neurobiology, test procedure, and its modifications. *Cogn Process* **13,** 93–110 (May 2012).

51. Du, Y., Fan, X., Lozano-Perez, T. & Kaelbling, L. Large Scale Robotic Navigation with Spatial Landmarks. *In Preperation* (2021).

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

53. Kumaran, D. & Maguire, E. A. An unexpected sequence of events: mismatch detection in the human hippocampus. *PLoS Biol* **4,** e424 (Nov. 2006).

54. Lee, I., Hunsaker, M. R. & Kesner, R. P. The role of hippocampal subregions in detecting spatial novelty. *Behav Neurosci* **119,** 145–53 (Feb. 2005).

55. Lee, J. M. *Introduction to Smooth Manifolds* (Springer New York, 2012).

56. Badino, H., Huber, D. & Kanade, T. *Visual topometric localization* in *2011 IEEE Intelligent Vehicles Symposium (IV)* (2011), 794–799.

57. Whittington, J. C. R. *et al.* The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell* **183,** 1249–1263.e23 (Nov. 2020).

58. George, D. *et al.* Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nat Commun* **12,** 2392 (Apr. 2021).

59. Farooq, U. & Dragoi, G. Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science* **363,** 168–173 (Jan. 2019).

60. Dragoi, G. & Tonegawa, S. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469,** 397–401 (Jan. 2011).

61. Dragoi, G. & Tonegawa, S. Selection of preconfigured cell assemblies for representation of novel spatial experiences. *Philos Trans R Soc Lond B Biol Sci* **369,** 20120522 (Feb. 2014).

62. Yim, M. Y., Sadun, L. A., Fiete, I. R. & Taillefumier, T. Place-cell capacity and volatility with grid-like inputs. *Elife* **10** (May 2021).

63. McKenzie, S. *et al.* Preexisting hippocampal network dynamics constrain optogenetically induced place fields. *Neuron* **109,** 1040–1054.e7 (Mar. 2021).

64. Milner, B. The medial temporal-lobe amnesic syndrome. *Psychiatr Clin North Am* **28,** 599–611, 609 (Sept. 2005).

65. Eichenbaum, H. The role of the hippocampus in navigation is memory. *J Neurophysiol* **117,** 1785–1796 (Apr. 2017).

66. Richmond, L. L. & Zacks, J. M. Constructing Experience: Event Models from Perception to Action. *Trends Cogn Sci* **21,** 962–980 (Dec. 2017).

67. Radvansky, G. A. & Copeland, D. E. Walking through doorways causes forgetting: situation models and experienced space. *Mem Cognit* **34,** 1150–6 (July 2006).

68. Eichenbaum, H. Is the rodent hippocampus just for 'place'? *Curr Opin Neurobiol* **6,** 187–95 (Apr. 1996).

69. Zacks, J. M., Kurby, C. A., Eisenberg, M. L. & Haroutunian, N. Prediction error associated with the perceptual segmentation of naturalistic events. *J Cogn Neurosci* **23,** 4057–66 (Dec. 2011).

70. Carpenter, F., Manson, D., Jeffery, K., Burgess, N. & Barry, C. Grid Cells Form a Global Representation of Connected Environments. *Current Biology* **25,** 1176–1182 (2015).

71. Spiers, H. J., Hayman, R. M. A., Jovalekic, A., Marozzi, E. & Jeffery, K. J. Place field repetition and purely local remapping in a multicompartment environment. *Cereb Cortex* **25,** 10–25 (Jan. 2015).

72. Duvelle, É. *et al.* Hippocampal place cells encode global location but not connectivity in a complex space. *Curr Biol* **31,** 1221–1233.e9 (Mar. 2021).

73. Deshmukh, S. S. & Knierim, J. J. Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus* **23,** 253–267 (Feb. 2013).

74. Høydal, Ø. A., Skytøen, E. R., Andersson, S. O., Moser, M.-B. & Moser, E. I. Object-vector coding in the medial entorhinal cortex. *Nature* **568,** 400–404 (Apr. 2019).

75. Low, R. J., Lewallen, S., Aronov, D., Nevers, R. & Tank, D. W. Probing variability in a cognitive map using manifold inference from neural dynamics. *bioRxiv* (2018).

76. Dudchenko, P. A. & Wood, E. R. in *Space, Time and Memory in the Hippocampal Formation* 253–272 (Springer Vienna, 2014).

77. Kinsky, N. R. *et al.* Trajectory-modulated hippocampal neurons persist throughout memory-guided navigation. *Nature Communications* **11** (May 2020).

78. Levy, S. J., Kinsky, N. R., Mau, W., Sullivan, D. W. & Hasselmo, M. E. Hippocampal spatial memory representations in mice are heterogeneously stable. *Hippocampus* **31,** 244–260 (Oct. 2020).

79. Kudrimoti, H., Barnes, C. & McNaughton, B. Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J Neurosci* **19,** 4090–101 (May 1999).

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

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

82. Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the hippocampus. *Nature Neuroscience* **12,** 913–918 (June 2009).

83. Buzsáki, G. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25,** 1073–188 (Oct. 2015).

84. Roux, L., Hu, B., Eichler, R., Stark, E. & Buzsáki, G. Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nat Neurosci* **20,** 845–853 (June 2017).

85. Joo, H. R. & Frank, L. M. The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nat Rev Neurosci* **19,** 744–757 (Dec. 2018).

86. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497,** 74–9 (May 2013).

87. Ólafsdóttir, H. F., Bush, D. & Barry, C. The Role of Hippocampal Replay in Memory and Planning. *Curr Biol* **28,** R37–R50 (Jan. 2018).

88. Erdem, U. M. & Hasselmo, M. A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *Eur J Neurosci* **35,** 916–31 (Mar. 2012).

89. Uria, B. *et al.* The Spatial Memory Pipeline: a model of egocentric to allocentric understanding in mammalian brains. *bioRxiv* (2020).

90. Yoon, K.-J. *et al.* Evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* (In review, 2012).