

The hippocampus as a predictive map

Kimberly L. Stachenfeld^{1,2,*}, Matthew M. Botvinick^{1,3}, and Samuel J. Gershman⁴

¹DeepMind, London, UK

²Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

³Gatsby Computational Neuroscience Unit, University College London, London, UK

⁴Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA, USA

*stachenfeld@google.com

ABSTRACT

A cognitive map has long been the dominant metaphor for hippocampal function, embracing the idea that place cells encode a geometric representation of space. However, evidence for predictive coding, reward sensitivity, and policy dependence in place cells suggests that the representation is not purely spatial. We approach this puzzle from a reinforcement learning perspective: what kind of spatial representation is most useful for maximizing future reward? We show that the answer takes the form of a predictive representation. This representation captures many aspects of place cell responses that fall outside the traditional view of a cognitive map. Furthermore, we argue that entorhinal grid cells encode a low-dimensional basis set for the predictive representation, useful for suppressing noise in predictions and extracting multiscale structure for hierarchical planning.

1 Introduction

Learning to predict long-term reward is fundamental to the survival of many animals. Some species may go days, weeks or even months before attaining primary reward, during which time aversive states must be endured. Evidence suggests that the brain has evolved multiple solutions to this reinforcement learning (RL) problem¹. One solution is to learn a model or “cognitive map” of the environment², which can then be used to generate long-term reward predictions through simulation of future states¹. However, this solution is computationally intensive, especially in real-world environments where the space of future possibilities is virtually infinite. An alternative “model-free” solution is to learn, from trial-and-error, a value function mapping states to long-term reward predictions³. However, dynamic environments can be problematic for this approach, because changes in the distribution of rewards necessitates complete relearning of the value function.

Here, we argue that the hippocampus supports a third solution: learning of a “predictive map” that represents each state in terms of its “successor states” (upcoming states)^{4,5}. This representation is sufficient for long-term reward prediction, is learnable using a simple, biologically plausible algorithm, and explains a wealth of data from studies of the hippocampus.

Our primary focus is on understanding the computational function of hippocampal place cells, which respond selectively when an animal occupies a particular location in space⁶. A classic and still influential view of place cells is that they collectively furnish an explicit map of space^{7,8}. This map can then be employed as the input to a model-based^{9–11} or model-free^{12,13} RL system for computing the value of the animal’s current state. In contrast, the predictive map theory views place cells as encoding predictions of future states, which can then be combined with reward predictions to compute values. This theory can account for why the firing of place cells is modulated by variables like obstacles, environment topology, and direction of travel. It also generalizes to hippocampal coding in non-spatial tasks. Beyond

24 the hippocampus, we argue that entorhinal grid cells¹⁴, which fire periodically over space, encode a
25 low-dimensional decomposition of the predictive map, useful for stabilizing the map and discovering
26 subgoals.

27 Results

28 The successor representation

29 An animal's optimal course of action will frequently depend on the location (or more generally, the "state")
30 that the animal is in. The hippocampus' purported role of representing location is therefore considered
31 to be a very important one. The traditional view of state representation in the hippocampus is that the
32 place cells index the current location by firing when the animal visits the encoded location, remaining
33 silent otherwise⁷. The main idea of the SR model, elaborated below, is that place cells do not encode
34 place *per se*, but rather a predictive representation of future states given the current state. Two states
35 that predict similar future states will have similar representations, and two physically adjacent states that
36 predict divergent future states will have dissimilar representations.

37 To motivate our use of the SR in the RL setting, we demonstrate that this representation emerges
38 naturally as a term M in the definition of value (V) often used in RL. We consider the problem of RL in a
39 Markov decision process consisting of the following elements¹⁵: a set of states (e.g., spatial locations), a
40 set of actions, a transition distribution $P(s'|s, a)$ specifying the probability of transitioning to state s' from
41 state s after taking action a , a reward function $R(s)$ specifying the expected immediate reward in state s ,
42 and a discount factor $\gamma \in [0, 1]$ that down-weights distal rewards. An agent chooses actions according to
43 a policy $\pi(a|s)$ and collects rewards as it moves through the state space. The value of a state is defined
44 formally as the expected discounted cumulative future reward under policy π :

$$V(s) = \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t) \mid s_0 = s \right], \quad (1)$$

45 where s_t is the state visited at time t . Our focus here is on policy evaluation (computing V). In our
46 simulations we feed the agent the optimal policy; in the Supplemental Methods we discuss algorithms
47 for policy improvement. To simplify notation, we assume implicit dependence on π and define the state
48 transition matrix T , where $T(s, s') = \sum_a \pi(a|s) P(s'|s, a)$.

49 The value function can be decomposed into the inner product of the reward function with a predictive
50 state representation known as the successor representation (SR)⁴, denoted by M :

$$V(s) = \sum_{s'} M(s, s') R(s'), \quad (2)$$

51 The SR encodes the expected discounted future occupancy of state s' along a trajectory initiated in state s :

$$M(s, s') = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathbb{I}(s_t = s') \mid s_0 = s \right], \quad (3)$$

52 where $\mathbb{I}(\cdot) = 1$ if its argument is true, and 0 otherwise.

53 An estimate of the SR (denoted \hat{M}) can be incrementally updated using a form of the temporal
54 difference learning algorithm^{4,16}. After observing a transition $s_t \rightarrow s_{t+1}$, the estimate is updated according
55 to:

$$\hat{M}_{t+1}(s, s') = \hat{M}_t(s, s') + \eta \left[\mathbb{I}(s_t = s') + \gamma \hat{M}_t(s_{t+1}, s') - \hat{M}_t(s, s') \right], \quad (4)$$

56 where η is a learning rate (unless specified otherwise, $\eta = 0.1$ in our simulations). The form of this
57 update is identical to the temporal difference learning rule for value functions¹⁵, except that in this case
58 the reward prediction error is replaced by a *successor prediction error* (the term in brackets). Note that
59 these prediction errors are distinct from state prediction errors used to update an estimate of the transition
60 function¹⁷; the SR predicts not just the next state but a superposition of future states over a possibly
61 infinite horizon. The transition and SR functions only coincide when $\gamma = 0$. We assume the SR matrix
62 M is initialized to the identity matrix, meaning $M(s, s') = 1$ if $s = s'$, and $M(s, s') = 0$ if $s \neq s'$. This
63 initialization can be understood to mean that each state will necessarily predict only itself.

64 The SR combines some of the advantages of model-free and model-based algorithms. Like model-
65 free algorithms, policy evaluation is computationally efficient with the SR. However, factoring the
66 value function into a state dynamics SR term and a reward term confers some of the flexibility usually
67 associated with model-based methods. Having separate terms for state dynamics and reward permits
68 rapid recomputation of new value functions when reward is changed independently of state dynamics,
69 as demonstrated in Fig. 1. The SR can be learned before any reward has been seen, so that at the first
70 introduction of reward, a value function can be computed immediately. When the reward function changes
71 – such as when the animal becomes satiated, or when food is redistributed about the environment – the
72 animal can immediately recompute a new value function based on its expected state transitions. A model-
73 free agent would have to relearn value estimates for each location in order to make value predictions, and
74 a model-based agent would need to aggregate the results of time-consuming searches through its model
75 before it could produce an updated value prediction^{1,4}. In Fig. S2, we demonstrate that while changing
76 the reward function completely disrupts model free learning of a value function in a 2-step tree maze, SR
77 learning can quickly adjust. Thus, the SR combines the efficiency of model-free control with some of the
78 flexibility of model-based control.

79 For an agent trying to optimize expected discounted future reward, two states that predict similar
80 successor states are necessarily similarly valuable, and can be safely grouped together¹⁸. This makes the
81 SR a good metric space for generalizing value. Since adjacent states will frequently lead to each other, the
82 SR will naturally represent adjacent states similarly and therefore be smooth over time and space in spatial
83 tasks. Since the SR is well defined for any Markov decision process, we can use the same architecture for
84 many kinds of tasks, not just spatial ones.

85 Hippocampal encoding of the successor representation

86 We now turn to our main theoretical claim: that the SR is encoded by the hippocampus. This hypothesis is
87 based on the central role of the hippocampus in representing space and context¹⁹, as well as its contribution
88 to sequential decision making^{20,21}. Although the SR can be applied to arbitrary state spaces, we focus on
89 spatial domains where states index locations.

90 Place cells in the hippocampus have traditionally been viewed as encoding an animal's current location.
91 In contrast, the predictive map theory views these cells as encoding an animal's *future* locations. Crucially,
92 an animal's future locations depend on its policy, which is constrained by a variety of factors such as the
93 environmental topology and the locations of rewards. We demonstrate that these factors shape place cell
94 receptive field properties in a manner consistent with a predictive map.

95 According to our model, the hippocampus represents the SR as a rate code across the population. Each
96 neuron represents some possible future state (e.g., spatial position) in the environment. At any current state
97 s , the population will encode a row of the SR matrix, $M(s, :)$. The firing rate of a single neuron encoding
98 state s' in the population is proportional to the discounted expected number of times it will be visited
99 under the present policy given the current position s . An SR place field refers to the firing rate of a single
100 SR-encoding neuron at each state in the task and corresponds to a column of the SR matrix, $M(:, s')$. This

101 vector contains the expected number of times a single encoded state s' will be visited under the current
102 policy, starting from any state s . In general, we will refer to place fields simulated under our model as
103 “SR receptive fields” or “SR place fields.” To summarize the relationship between the SR matrix M and
104 simulated hippocampal cells: The firing of *all* the neurons at *one* state s is modeled by a *row* $M(s, :)$ of
105 the SR matrix M , and the firing of *one* neuron encoding s' evaluated at *all* states is modeled by a *column*
106 $M(:, s')$. This is illustrated in Fig. 1.

107 We first try to build some intuition for this idea, and how it relates to a more traditional view of place
108 cells. In an open, 2D environment, the canonical place has a gradually decaying, roughly circular firing
109 field. These are often modeled as approximately Gaussian. In such an environment, the SR place fields
110 look essentially the same, with peaks of high firing surrounded by a radius of gradually reduced firing.
111 The SR model makes this prediction because under a random walk, the animal is likely to visit its current
112 location and nearby locations immediately, and more distant locations later. Thus, the states closer to the
113 encoded location of an SR place cell will predict a higher expected discounted number of visits to the
114 encoded location, and will elicit higher firing of the encoding cell.

115 Fig. 3 illustrates the experimental conditions in which the predictions of the SR model (Fig. 3C)
116 depart from the predictions of two alternative models (Fig. 3A-B). As examples, we implement the three
117 models for a 2D room containing an obstacle and for a 1D track with an established preferred direction
118 of travel. The first alternative model is a Gaussian place field in which firing is related to the Euclidean
119 distance from the field center (Fig. 3A), usually invoked for modeling place field activity in open spatial
120 domains^{22,23}. The second alternative model is a topologically sensitive place field in which firing is related
121 to the average path length from the field center, where paths cannot pass through obstacles¹³ (Fig. 3A).
122 Like the topological place fields and unlike the Gaussian place fields, the SR place fields respect obstacles
123 in the 2D environment. Since states on opposite sides of a barrier cannot occur nearby in time, SR place
124 fields will tend to be active on only one side of a barrier.

125 On the 1D track, the SR place fields skew opposite the direction of travel. This backward skewing
126 arises because upcoming states can be reliably predicted further in advance when traveling repeatedly
127 in a particular direction. Neither of the control models provide ways for a directed behavioral policy to
128 interact with state representation, and therefore cannot show this effect. Evidence for predictive skewing
129 comes from experiments in which animals traveled repeatedly in a particular direction along a linear
130 track^{24,25}. The authors noted this as evidence for predictive coding in hippocampus^{24,26}. In Fig. 2, we
131 explain how a future-oriented representation evokes a forward-skewing representation in the population at
132 any given point in time but implies that receptive fields for any individual cell should skew backwards. In
133 order for a given cell to fire predictively, it must begin firing before its encoded state is visited, causing a
134 backward-skewed receptive field. Figure 4 compares the predicted and experimentally observed backward
135 skewing, demonstrating that the model captures the qualitative pattern of skewing observed when the
136 animal has a directional bias.

137 Consistent with the SR model, experiments have shown that place fields become distorted around
138 barriers^{27–29}. In Figure 5, we explore the effect of placing obstacles in a Tolman detour maze on the SR
139 place fields and compare to experimental results obtained by Alvernhe *et al.*²⁹. When a barrier is placed in
140 a maze such that the animal is forced to take a detour, the place fields engage in “local remapping.” Place
141 fields near the barrier change their firing fields significantly more than those further from the barrier (Fig.
142 5A-C). When barriers are inserted, SR place fields change their fields near the path blocked by the barrier
143 and less so at more distal locations where the optimal policy is unaffected (Fig. 5D-F). This locality is
144 imposed by the discount factor. The full set of place fields is included in the supplement (Fig. S3).

145 The SR model can be used to explain how hippocampal place fields depend on behaviorally relevant
146 features that alter an animal’s transition policy, such as reward. Using an annular watermaze, Hollup and

147 colleagues demonstrated that a hidden, stationary reward affects the distribution of place fields³⁰. Animals
148 were required to swim in some preferred direction around a ring-shaped maze filled with an opaque liquid
149 until they reached a hidden platform where they could rest. Hollup and colleagues found that the segment
150 containing the platform had more place fields centered within it than any other segment, and that the
151 preceding segment consistently had the second-largest number of place fields centered within it (Fig. 6A).

152 We simulated this task using a sequence of states connected in a ring. The transition policy was such
153 that the animal lingered longer near the rewarded location and had a preferred direction of travel (right,
154 or counterclockwise, in this case), matching behavioral predictions recorded by the authors³⁰. We set
155 the probability of transitioning left to 0 to illustrate the predictions of our model more clearly. As we
156 show in Figure 6A-B, the SR model predicts elevated firing near the rewarded location and backward
157 skewing of place fields. This creates an asymmetry, whereby the locations preceding the rewarded location
158 will experience slightly higher firing rates as well. Furthermore, this asymmetric backward skew makes
159 it likely that fields will overlap with the previous segment, not the upcoming segment. Figure 6C-D
160 demonstrates how this backward skewing can equate to a backward shift in cell peak in the presence of
161 noise or location uncertainty. This may explain the asymmetry found in the distribution of place field
162 peaks about the rewarded segment.

163 While Hollup and colleagues found an asymmetric distribution of place cells about the rewarded
164 segment, they also found that place fields were roughly the same size at reward locations as at other
165 locations. In contrast, the SR predicts that fields should get larger near reward locations (Fig. 6B), with
166 the magnitude of this effect modulated by the discount factor (Fig. S6). Thus, the SR is still an incomplete
167 account of reward-dependent place fields.

168 Note that the SR model does not predict that place fields would be immediately affected by the
169 introduction of a reward. Rather, the shape of the fields should change as the animal gradually adjusts its
170 policy and experiences multiple transitions consistent with that policy. The SR is affected by the presence
171 of the reward because rewards induce a change in the animal's policy, which determines the predictive
172 relationships between states.

173 Under a sufficiently large discount, the SR model predicts that firing fields centered near rewarded
174 locations will expand to include the surrounding locations and increase their firing rate under the optimal
175 policy. The animal is likely to spend time in the vicinity of the reward, meaning that states with or near
176 reward are likely to be common successors. SR place fields in and near the rewarded zone will cluster
177 because it is likely that states near the reward were anticipated by other states near the reward (Fig. S7).
178 For place fields centered further from the reward, the model predicts that fields will skew opposite the
179 direction of travel toward the reward, due to the effect illustrated in Fig. 2: a state will only be predicted
180 when the animal is approaching reward from some more distant state. Given a large potentially rewarded
181 zone or a noisy policy, these somewhat contradictory effects are sufficient to produce clustering of place
182 fields near the rewarded zone (Fig. S7). The punished locations will induce the opposite effect, causing
183 fields near the punished location to spread away from the rarely-visited punished locations (Fig. S5F). The
184 SR place fields for each of these environments are shown in Figure S5.

185 In addition to the influence of experimental factors, changes in parameters of the model will have
186 systematic effects on the structure of SR place fields. Motivated by data showing a gradient of increasing
187 field sizes along the hippocampal longitudinal axis^{31,32}, we explored the consequences of modifying the
188 discount factor γ in Figure S4 and Figure S6. Hosting a range of discount factors along the hippocampal
189 longitudinal axis provides a multi-timescale representation of space. It also circumvents the problem of
190 having to assume the same discount parameter for each problem or adaptively computing a new discount.
191 Another consequence is that larger place fields reflect the community structure of the environment. In
192 Figure S5, we show how the SR fields begin to expand their fields to cover all states with the same

193 compartment for a large enough discount. This overlap drives the clustering of states within the same
194 community. A gradient of discount factors might therefore be useful for decision making at multiple levels
195 of temporal abstraction^{18,33,34}.

196 An appealing property of the SR model is that it can be applied to non-spatial state spaces. Fig.
197 7A-D shows the SR embedding of an abstract state space used in a study by Schapiro and colleagues^{18,35}.
198 Human subjects viewed sequences of fractals drawn from random walks on the graph while brain activity
199 was measured using fMRI. We compared the similarity between SR vectors for pairs of states with pattern
200 similarity in the hippocampus. The key experimental finding was that hippocampal pattern similarity
201 mirrored the community structure of the graph: states with similar successors were represented similarly³⁵.
202 The SR model recapitulates these findings, since states in the same community tend to be visited nearby in
203 time, making them predictive of one another (Fig. 7E-G). A recent related fMRI result from Garvert and
204 colleagues provides further support that the hippocampus represents upcoming successors in a non-spatial,
205 relational task by showing that a successor model provided the best metric for explaining variance in
206 recorded hippocampal adaptation³⁶.

207 To demonstrate further how the SR model can integrate spatial and temporal coding in the hippocampus,
208 we simulated results from a recent study³⁷ in which subjects were asked to navigate among pairs of
209 locations to retrieve associated objects in a virtual city (8A). Since it was possible to “teleport” between
210 certain location pairs, while others were joined only by long, winding paths, spatial Euclidean distance
211 was decoupled from travel time. The authors found that objects associated with locations that were nearby
212 in either space or time increased their hippocampal pattern similarity (Fig. 8B). Both factors (spatial and
213 temporal distance) had a significant effect when the other was regressed out (Fig. 8C). The SR predicts
214 this integrated representation of spatiotemporal distance: when a short path is introduced between distant
215 states, such as by a teleportation hub, those states come predict one another.

216 Dimensionality reduction of the predictive map by entorhinal grid cells

217 Because the firing fields of entorhinal grid cells are spatially periodic, it was originally hypothesized that
218 grid cells might represent a Euclidean spatial metric to enable dead reckoning^{8,14}. Other theories have
219 suggested that these firing patterns might arise from a low-dimensional embedding of the hippocampal
220 map^{5,23,38}. Combining this idea with the SR hypothesis, we argue that grid fields reflect a low-dimensional
221 eigendecomposition of the SR. A key implication of this hypothesis is that grid cells will respond differently
222 in environments with different boundary conditions.

223 The boundary sensitivity of grid cells was recently highlighted by a study that manipulated boundary
224 geometry³⁹. In square environments, different grid modules had the same alignment of the grid relative
225 to the boundaries (modulo 60°, likely due to hexagonal symmetry in grid fields), whereas in a circular
226 environment grid field alignment was more variable, with a qualitatively different pattern of alignment
227 (Fig. 9A-C). Krupic *et al.* performed a “split-halves” analysis, in which they compared grid fields in
228 square versus trapezoidal mazes, to examine the effect of breaking an axis of symmetry in the environment
229 (Fig 9D,E). They found that moving the animal to a trapezoidal environment, in which the left and right
230 half of the environment had asymmetric boundaries, caused the grid parameters to be different on the
231 two sides of the environment³⁹. In particular, the spatial autocorrelegrams – which reveal the layout of
232 spatial displacement at which the grid field repeats itself – were relatively dissimilar over both halves of
233 the trapezoidal environment. The grid fields in the trapezoid could not be attributed to linearly warping
234 the square grid field into a trapezoid, raising the question of how else boundaries could interact with grid
235 fields.

236 According to the SR eigenvector model, these effects arise because the underlying statistics of
237 the transition policy changes with the geometry. We simulated grid fields in a variety of geometric

238 environments used by Krupic and colleagues (Fig. 9F-H; Fig. 9A-S9). In agreement with the empirical
239 results, the orientation of eigenvectors in the circular environment tend to be highly variable, while those
240 recorded in square environments are almost always aligned to either the horizontal or vertical boundary
241 of the square (Fig. 9G,J). The variability in the circular environment arises because the eigenvectors are
242 subject to the rotational symmetry of the circular task space. SR eigenvectors also emulate the finding that
243 grids on either side of a square maze are more similar than those on either side of a trapezoid, because the
244 eigenvectors capture the effect of these irregular boundary conditions on transition dynamics.

245 Another main finding of Krupic et al.³⁹ was that when a square environment is rotated, grids remain
246 aligned to the boundaries as opposed to distal cues. SR eigenvectors inherently reproduce this effect, since
247 a core assumption of the theory is that grid firing is anchored to state in a transition structure, which is
248 itself constrained by boundaries. The complete set of the first 64 eigenvectors is shown in Figures S8A
249 and S9. While many fields conform to the canonical grid cell, others have skewed or otherwise irregular
250 waveforms. Our model predicts that such shapes would be included in the greater variety of firing fields
251 found in MEC that do not match the standard grid-like criterion.

252 A different manifestation of boundary effects is the fragmentation of grid fields in a hairpin maze⁴⁰.
253 Consistent with the empirical data, SR eigenvector fields tend to align with the arms of the maze, and
254 frequently repeat across alternating arms (Figure 10)⁴⁰. While patterns at many timescales can be found
255 in the eigenvector population, those at alternating intervals are most common and therefore replicate the
256 checkerboard pattern observed in the experimental data (Fig. S9).

257 To further explore how compartmentalized environments could affect grid fields, we simulated a recent
258 study⁴¹ that characterized how grid fields evolve over several days' exposure to a multi-compartment
259 environment (Fig. 11). While grid cells initially represented separate compartments with identical fields
260 (repeated grids), several days of exploration caused fields to converge on a more globally coherent grid
261 (Fig. 11D-F). With more experience, the grid regularity of the fields simultaneously decreased, as did the
262 similarity between the grid fields recorded in the two rooms (Fig. 11C). The authors conclude that grid
263 cells will tend to a regular, globally coherent grid to serve as a Euclidean metric over the full expanse of
264 the enclosure.

265 Our model suggests that the fields are tending not toward a globally *regular* grid, but to a predictive map
266 of the task structure, which is shaped in part by the global boundaries but also by the multi-compartment
267 structure. We simulated this experiment by initializing grid fields to a local eigenvector model, in which
268 the animal has not yet learned how the compartments fit together. After the SR eigenvectors have been
269 learned, we relax the constraint that representations be the same in both rooms and let eigenvectors and the
270 SR be learned for the full environment. As the learned eigenvectors converge, they increasingly resemble
271 a global grid and decreasingly match the predictions of the local fit (Fig. 11H-L; Fig. S10). As with the
272 recorded grid cells, the similarity of the fields in the two rooms drops to an average value near zero (Fig.
273 11I). They also have less regular grids compared to a single-compartment rectangular enclosure, explaining
274 the drop in grid regularity observed by Carpenter *et al.* as the grid fields became more “global”⁴¹. Since
275 separating barriers between compartments perturb the task topology from an uninterrupted 2D grid.

276 The eigenvectors of the SR are invariant to the discount factor of an SR matrix. This is because any
277 SR can be written as a weighted sum of transition policy matrices, as we explain in more detail in the
278 Supplemental Methods. The same eigenvectors will therefore support multiple SR matrices learned for the
279 same task but with different planning horizons. SR matrices with a large discount factor will place higher
280 eigenvalues on the eigenvectors with large spatial scales and low spatial frequency, whereas those with
281 smaller discounts and smaller place fields project more strongly onto higher spatial-frequency grid fields.
282 As discount is increased, the eigenvalues gradually shift their weight from the smaller scale to the larger
283 scale eigenvectors (Fig. S11). This mirrors data suggesting that hippocampal connections to and from

284 MEC vary gradually alongside place field spatial scale along the longitudinal axis^{31,32,42,43}. Grid fields,
285 in contrast, cluster in discrete modules⁴⁴. The SR eigenvectors are quantized as discrete modules as well,
286 as we show in Figure S12.

287 A normative motivation for invoking low-dimensional projections as a principle for grid cells is that
288 they can be used to smooth or “regularize” noisy updates of the SR. When the projection is based on an
289 eigendecomposition, this constitutes a form of *spectral regularization*⁴⁵. A smoothed version of the SR
290 can be obtained by reconstructing the SR from its eigendecomposition using only low-frequency (high
291 eigenvalue) components, thereby filtering out high-frequency noise (see Methods). This smoothing will fill
292 in the blanks in the successor representations, enabling faster convergence time and a better approximation
293 of the SR while it is still being learned. Spectral regularization has a long history of improving the
294 approximation of large, incomplete matrices in real-world domains, most commonly through matrix
295 factorization⁴⁵. The utility of a spectral basis for approximating value functions in spatial and other
296 environments has been demonstrated in the computational RL literature⁴⁶. In Figure S13A, we provide a
297 demonstration of how this kind of spectral regularization can allow the SR to be more accurately estimated
298 despite the presence of corrupting noise in a multi-compartment environment. In Figure S13B, we show
299 that spectral regularization provides a better reconstruction basis than a globally uniform Fourier basis,
300 because the former does not smooth over boundaries.

301 We also demonstrate how reweighting eigenvalues so that more weight is placed on the low-frequency
302 eigenvectors allows us to approximate the SR matrix for larger discounts with significantly less training
303 time (Fig. S13C). TD learning can take a long time to converge when the discount factor is large. Spectral
304 regularization can allow the SR to support planning over a longer timescale after significantly less training.

305 We include our own modest demonstration of how spectral regularization can improve SR-based
306 value function approximation in a noisy, multicompartment spatial task. Importantly, the regularization is
307 topologically sensitive, meaning that smoothing respects boundaries of the environment. Regularization
308 using a Fourier decomposition does not share this property, and will smooth over boundaries (Fig. S13).
309 The regularization hypothesis is consistent with data suggesting that although grid cell input is not required
310 for the emergence of place fields, place field stability and organization depends crucially on input from
311 grid cells^{47–49}. These eigenvectors also provide a useful partitioning of the task space, as discussed in the
312 following section.

313 **Subgoal discovery using grid fields**

314 In structured environments, planning can be made more efficient by decomposing the task into subgoals,
315 but the discovery of good subgoals is an open problem. The SR eigenvectors can be used for subgoal
316 discovery by identifying “bottleneck states” that bridge large, relatively isolated clusters of states, and
317 group together states that fall on opposite sides of the bottlenecks^{50,51}. Since these bottleneck states
318 are likely to be traversed along many optimal trajectories, they are frequently convenient waypoints to
319 visit. Navigational strategies that exploit bottleneck states as subgoals have been observed in human
320 navigation⁵². It is also worth noting that accompanying the neural results displayed in Fig. 7, the authors
321 found that when subjects were asked to parse sequences of stimuli into events, stimuli found at topological
322 bottlenecks were frequent breakpoints¹⁸.

323 The formal problem of identifying these bottlenecks is known as the k -way normalized min-cut
324 problem. An approximate solution can be obtained using spectral graph theory⁵³. First, the top $\log k$
325 eigenvectors of a matrix known as the graph Laplacian are thresholded such that negative elements of each
326 eigenvector go to zero and positive elements go to one. Edges that connect between these two labeled
327 groups of states are “cut” by the partition, and nodes adjacent to these edges are a kind of bottleneck
328 subgoal. The first subgoals that emerge will be the cut from the lowest-frequency eigenvector, and these

329 subgoals will approximately lie between the two largest, most separable clusters in the partition (see
330 Supplemental Methods for more detail). A prioritized sequence of subgoals is obtained by incorporating
331 increasingly higher frequency eigenvectors that produce partition points nearer to the agent.

332 The SR shares its eigenvectors with the graph Laplacian (see Supplemental Methods)⁵, making SR
333 eigenvectors equally suitable for this process of subgoal discovery. We show in Figure S14 that the
334 subgoals that emerge in a 2-step decision task and in a multi-compartment environment tend to fall near
335 doorways and decision points: natural subgoals for high-level planning. It is worth noting that SR matrices
336 parameterized by larger discount factors γ will project predominantly on the large-spatial-scale grid
337 components (Fig. S11). The relationship between more temporally diffuse, abstract SRs, in which states in
338 the same room are all encoded similarly (Fig. S4), and the subgoals that join those clusters can therefore
339 be captured by which eigenvalues are large enough to consider.

340 It has also been shown experimentally that entorhinal lesions impair performance on navigation tasks
341 and disrupt the temporal ordering of sequential activations in hippocampus while leaving performance on
342 location recognition tasks intact^{48,54}. This suggests a role of grid cells in spatial planning, and encourages
343 us to speculate about a more general role for grid cells in hierarchical planning.

344 Discussion

345 The hippocampus has long been thought to encode a cognitive map, but the precise nature of this map
346 is elusive. The traditional view that the map is essentially spatial^{7,8} is not sufficient to explain some of
347 the most striking aspects of hippocampal representation, such as the dependence of place fields on an
348 animal's behavioral policy and the environment's topology. We argue instead that the map is essentially
349 *predictive*, encoding expectations about an animal's future state. This view resonates with earlier ideas
350 about the predictive function of the hippocampus^{20,24,55-59}. Our main contribution is a formalization of
351 this predictive function in a reinforcement learning framework, offering a new perspective on how the
352 hippocampus supports adaptive behavior.

353 Our theory is connected to earlier work by Gustafson and Daw¹³ showing how topologically-sensitive
354 spatial representations recapitulate many aspects of place cells and grid cells that are difficult to rec-
355 oncile with a purely Euclidean representation of space. They also showed how encoding topological
356 structure greatly aids reinforcement learning in complex spatial environments. Earlier work by Foster
357 and colleagues¹² also used place cells as features for RL, although the spatial representation did not
358 explicitly encode topological structure. While these theoretical precedents highlight the importance of
359 spatial representation, they leave open the deeper question of why particular representations are better than
360 others. We showed that the SR naturally encodes topological structure in a format that enables efficient
361 RL.

362 The work is also related to work done by Dordek *et al.*²³, who demonstrated that gridlike activity
363 patterns from principal components of the population activity of simulated Gaussian place cells. As we
364 mentioned in the Results, one point of departure between empirically observed grid cell data and SR
365 eigenvector account is that in rectangular environments, SR eigenvector grid fields can have different
366 spatial scales aligned to the horizontal and vertical axis (see Fig. S8)¹⁴. In grid cells, the spatial scales
367 tend to be approximately constant in all directions unless the environment changes⁶⁰. The principal
368 components of Gaussian place fields are mathematically related to the SR eigenvectors, and naturally
369 also have grid fields that scale independently along the perpendicular boundaries of a rectangular room.
370 However, Dordek *et al.* found that when the components were constrained to have non-negative values
371 and the constraint that components be orthogonal was relaxed, the scaling became uniform in all directions
372 and the lattices became more hexagonal²³. This suggests that the difference between SR eigenvectors

373 and recorded grid cells is not fundamental to the idea that grid cells are doing spectral dimensionality
374 reduction. Rather, additional constraints such as non-negativity are required.

375 The SR can be viewed as occupying a middle ground between model-free and model-based learning.
376 Model-free learning requires storing a look-up table of cached values estimated from the reward history^{1, 61}.
377 Should the reward structure of the environment change, the entire look-up table must be re-estimated.
378 By decomposing the value function into a predictive representation and a reward representation, the SR
379 allows an agent to flexibly recompute values when rewards change, without sacrificing the computational
380 efficiency of model-free methods⁴. Model-based learning is robust to changes in the reward structure, but
381 requires inefficient algorithms like tree search to compute values^{1, 15}.

382 Certain behaviors often attributed to a model-based system can be explained by a model in which
383 predictions based on state dynamics and the reward function are learned separately. For instance, the
384 *context preexposure facilitation effect* refers to the finding that contextual fear conditioning is acquired
385 more rapidly if the animal has the chance to explore the environment for several minutes before the first
386 shock⁶². The facilitation effect is classically believed to arise from the development of a conjunctive
387 representation of the context in the hippocampus, though areas outside the hippocampus may also develop
388 a conjunctive representation in the absence of the hippocampus, albeit less efficiently⁶³. The SR provides a
389 somewhat different interpretation: over the course of preexposure, the hippocampus develops a *predictive*
390 representation of the context, such that subsequent learning is rapidly propagated across space. Figure S15
391 shows a simulation of this process and how it accounts for the facilitation effect.

392 Many models of prospective coding in the hippocampus have drawn inspiration from the well-
393 documented ordered temporal structure of firing in hippocampus relative to the theta phase^{20, 64, 65}, and
394 considered the many ways in which replaying hippocampal sweeps during sharp wave ripple events might
395 be used for planning^{66–71}. The firing of cells in hippocampus is aligned to theta such that cells encoding
396 more distant places fire later during a theta cycle than immediately upcoming states (a phenomenon
397 referred to as theta precession). States fire in a sequence ordered according to when they will next appear,
398 suggesting a likely mechanism for forward sequential planning^{65, 72}.

399 However, precession alone is probably not sufficient to enact backward expansion of place fields in
400 CA1, since NMDA antagonists that disrupt the persistent, backward expansion of place fields leave theta
401 precession intact⁷³. Furthermore, precession in CA1 likely originates outside of the hippocampus, as it
402 arises in MEC independently⁷⁴, and depends crucially on input from surrounding areas such as MEC and
403 CA3^{54, 75}. Thus, we think that it is worthwhile to consider the possible contributions of this backward
404 expansion to planning in addition to the contributions of the hippocampal temporal code examined by this
405 prior work.

406 The type of prospective coding implemented by theta precession and sharp wave ripple events is
407 reminiscent of model-based, sequential forward planning²⁰; many experiments and theoretical proposals
408 have looked at how replaying these sequences at decision points and at rest can underlie planning^{66–68, 70, 71}.
409 By integrating the reward reactivated at each state along a sweep through upcoming states, the value of a
410 specific upcoming trajectory can be predicted.

411 The SR is a different type of prospective code, with different tradeoffs. The SR marginalizes over all
412 possible sequences of actions, making predictions over an arbitrarily long timescale in constant time. This
413 results in a loss of flexibility relative to model-based planning, but greater computational efficiency. Thus,
414 the SR cannot replace the full functionality of model-based sweeps. However, it might provide a useful
415 adjunct to this functionality.

416 One way to combine the strengths of model-based planning with the SR would be to use the SR to
417 extend the range of forward sweeps. In Fig. S19, we illustrate how performing sweeps in the successor
418 representation space (Fig. S19F) or performing sweeps that terminate on a successor representation

419 of the terminal state (Fig. S19G) can extend the range of these predictions, making the hippocampal
420 representations a more powerful substrate for planning. This is tantamount to a “bootstrapped search”
421 algorithm, variants of which have been successful in a range of applications^{76,77}.

422 The SR model we describe is trained on the policy the animal has experienced. This means that when
423 the reward is changed, the new value function computed from the existing SR will initially be based on the
424 old policy. The new optimal policy is unlikely to be the same as the old one, which means that the new
425 value function is not correct, and must be refined as the animal optimizes its behavior. This problem is
426 encountered with all learning algorithms that learn cached statistics under the current policy dynamics.

427 In some cases, the old SR will be a reasonable initialization. In many environments, certain aspects of
428 the dynamics are not subject to the animal’s control, and the underlying adjacency structure is unlikely to
429 change. Furthermore, if rewards tend to be distributed in the same general area of a task, many policy
430 components will generalize. It is hard to make comprehensive claims about whether or not the space of
431 naturalistic tasks adheres to these properties in general. Recent computational work has demonstrated that
432 deep successor features (a more powerful generalization of the successor representation model) generalize
433 well across changing goals and environments in the domain of navigation⁷⁸.

434 To give an intuition of how the flexibility of the SR-based value computation depends on task hierarchy
435 and simulation parameters, we look at generalization using a simple tree-structured maze. Figure S16
436 illustrates how the quality of SR generalization depends on the policy stochasticity (parameterized by
437 β) and how similar the optimal paths are for the old and new rewarded location. When there is greater
438 stochasticity (closer to the random walk policy), the SR’s generalization to highly dissimilar locations is
439 less impaired, but there is also a reduced generalization advantage when the reward ends up nearby. The
440 random walk SR is used as a baseline. By diffusing value through the graph in accordance with the task’s
441 underlying adjacency structure, this representation always generalizes better than re-initializing to a state
442 index representation. The animal should maintain support for random actions until it is very certain of the
443 optimal path. Spectral regularization can promote this by smoothing the SR.

444 When the SR fails to support value computation given the new reward, there are other mechanisms
445 that can compensate. Models such as Dyna update cached statistics using sweeps through a model,
446 revising them flexibly⁷⁶. The original form of Dyna demonstrated how model-based and model-free
447 mechanisms could collaboratively update a value function. However, the value function can be replaced
448 with any statistic learnable through temporal differences, including the SR, as demonstrated by recent
449 work⁷⁹. Furthermore, there is evidence from humans that when reward is changed, revaluation occurs in a
450 policy-dependent manner, consistent with the kind of partial flexibility conferred by the SR⁸⁰.

451 Recent work has elucidated connections between models of episodic memory and the SR. Specifically,
452 Gershman *et al.* demonstrated that the SR is closely related to the Temporal Context Model (TCM) of
453 episodic memory^{16,19}. The core idea of TCM is that items are bound to their temporal context (a running
454 average of recently experienced items), and the currently active temporal context is used to cue retrieval of
455 other items, which in turn cause their temporal context to be retrieved. The SR can be seen as encoding a
456 set of item-context associations. The connection to episodic memory is especially interesting given the
457 crucial mnemonic role played by the hippocampus and entorhinal cortex in episodic memory. Howard
458 and colleagues⁸¹ have laid out a detailed mapping between TCM and the medial temporal lobe (including
459 entorhinal and hippocampal regions).

460 Spectral graph theory provides insight into the topological structure encoded by the SR. We showed
461 specifically that eigenvectors of the SR can be used to discover a hierarchical decomposition of the
462 environment for use in hierarchical RL. Spectral analysis has also frequently been invoked as a compu-
463 tational motivation for entorhinal grid cells (e.g.,⁸²). The fact that any function can be reconstructed by
464 sums of sinusoids suggests that the entorhinal cortex implements a kind of Fourier transform of space.

465 However, Fourier analysis is not the right mathematical tool when dealing with spatial representations in
466 a topologically structured environment, since we do not expect functions to be smooth over boundaries
467 in the environment. This is precisely the purpose of spectral graph theory: Instead of being maximally
468 smooth over Euclidean space, the eigenvectors of the graph Laplacian embed the smoothest approximation
469 of a function that respects the graph topology⁴⁶.

470 In conclusion, the SR provides a unifying framework for a wide range of observations about the
471 hippocampus and entorhinal cortex. The multifaceted functions of these brain regions can be understood
472 as serving a superordinate goal of prediction.

473 **Methods**

474 **Task simulation**

475 Environments were simulated by discretizing the plane into points, and connecting these points along a
476 triangular lattice (Fig. S1A). The adjacency matrix A was constructed such that $A(s, s') = 1$ wherever it is
477 possible to transition between states s and s' , and 0 otherwise.

478 The transition probability matrix T was defined such that $T(s, s')$ is the probability of transitioning
479 from state s to s' . Under a random walk policy, where the agent chooses randomly among all available
480 transitions, the transition probability distribution is uniform over allowable transitions. This amounts to
481 simply normalizing A so that each row of A sums to 1 to meet the constraint that the possible transition
482 from s must sum to 1. When reward or punishment was included as part of the simulated task, we
483 computed the optimal policy using value iteration and a softmax value function parameterized by β ¹⁵.

484 **SR computation**

485 The successor representation is a matrix, M where $M(s, s')$ is equal to the discounted expected number
486 of times the agent visits state s' starting from s (see Equation 3 for the mathematical definition and Fig.
487 S1B for an illustration). When the transition probability matrix is known, we can compute the SR as a
488 discounted sum over transition matrices raised to the exponent t . The matrix T^t is the t -step transition
489 matrix, where $T^t(s, s')$ is the probability of transitioning from s to s' in exactly t steps.

$$M = \sum_{t=0}^{\infty} \gamma^t T_{\pi}^t \quad (5)$$

490 This sum is a geometric matrix series, and for $\gamma < 1$, it converges to the following finite analytical solution:

$$M = \sum_{t=0}^{\infty} \gamma^t T_{\pi}^t = (I - \gamma T_{\pi})^{-1} \quad (6)$$

491 In most of our simulations, the SR was computed analytically from the transition matrix using this
492 expression.

493 The SR can be learned on-line using the temporal differences update rule of Equation 4⁴ (also see¹⁵
494 for background on TD learning) (Fig. 11, Fig. S1, Fig. S3). When noise was injected into the location
495 signal (Fig. S3). Noise was injected into the location signal by adding uniform random noise with mean 0
496 to the state indicator vector.

497 **Eigenvector computation and Spectral Regularization**

498 In generating the grid cells shown, we assume random walk policy, which is the maximum entropy prior
499 for policies (see⁸³ for why maximum entropy priors can be good priors for regularization). However, since

500 the learned eigenvectors are sensitive to the sampling statistics, our model predicts that regions of the
501 task space more frequently visited would come to be over-represented in the grid space (see Figure S8 for
502 examples). For most figures, we compute the eigenvectors of the SR using the built-in MATLAB `eig`
503 function (Fig. S1C). We then thresholded the eigenvectors at 0 so that firing rates are not negative (Fig.
504 S1D).

505 For Figure 11, eigenvectors were computed incrementally using a Candid Covariance-free Incremental
506 PCA (CCIPCA), an algorithm that efficiently implements stochastic gradient descent to compute principal
507 components⁸⁴ (eigenvectors and principal components are equivalent in this and many domains). Spectral
508 regularization was implemented by reconstructing the SR from the truncated eigendecomposition (Fig.
509 S13). Spectral reconstruction for Figure S13 was implemented by shifting the eigenvalues so that more
510 weight was placed on low-frequency eigenvectors, rather than imposing a hard cutoff on high-frequency
511 eigenvectors, and reconstructing an SR that corresponded to a larger discount factor. This allowed larger-
512 discount SRs to be more exactly approximated. The reconstructed SR matrices M_{recon} were compared to
513 the ground truth matrix M_{gt} by taking the correlation between M_{recon} and M_{gt} (Fig. S13). This measure
514 indicates whether policies based on SR-based value functions for different reward functions will tend
515 send the animal in the right direction. Details can be found in the Supplemental Methods.

516 Plotting receptive fields

517 To visualize place fields under the SR model, we showed heat maps of how active each SR-encoding
518 neuron would be at each state in the environment (Fig. S1E-F). This shows the discounted expected
519 number of times the neuron's encoded state s will be visited from each other state in the environment,
520 and corresponds to taking a column $M(s, :)$ from the SR matrix and reshaping it so that each element
521 appears at the x, y location of its corresponding state. We use the same reshaping and plotting procedure to
522 visualize eigenvector grid cells, using the columns of the thresholded eigenvector matrix U in place of M .

523 Quantifying place and grid fields

524 To quantify place field clustering, center of mass (CoM) of SR place fields was computed by summing the
525 locations of firing, weighted by the firing rate at that location (normalized so that the total firing summed
526 to 1):

$$\text{CoM}(s) = \frac{\sum_{s'} M(s, s') \mathbf{p}(s')}{\sum_{s'} M(s, s')}, \quad (7)$$

527 where $\mathbf{p}(s')$ is the (X, Y) coordinate of the place field centered at state s' .

528 In Fig. 5, spatial similarity was computed by taking the Fisher z transform of spatial correlation between
529 fields. Statistics were computed in this z space.

530 Grid field quantifications paralleled the analyses of Krupic *et al.*³⁹: an ellipse was fit to the 6 peaks
531 closest to the central peak, “orientation” refers to the orientation of the main axes (a, b) . “Correlation”
532 always refers to the Pearson correlation, “spatial correlation” refers to the Pearson correlation computed
533 over points in space (as opposed to points in a vector), and spatial autocorrelation refers to the 2D
534 auto-convolution.

535 To measure similarity between halves of the environment in Figure 9, we 1) computed the spatial
536 autocorrelation for each half, 2) selected a circular window in the center of the autocorrelation, and 3)
537 computed the correlation between autocorrelations of the two halves in the window. This paralleled the
538 analysis taken by Krupic *et al.*³⁹ and provides a measure of grid similarity across halves of the environment.
539 The circular window is used to control for the fact that the boundaries of the square and trapezoid in
540 the two halves of the respective environments differ. The mean similarity was *not* computed in Fisher

541 z -transformed space, as one would normally do, but rather in correlation space. This was because the
542 similarity for many of the square eigenvectors and at least one trapezoidal eigenvector was exactly 1, for
543 which $z = \infty$. A dot plot is superimposed over this plot so the statistics of the distribution can be visualized.

544 In evaluating our simulations of the grid fields reported by Carpenter *et al.*⁴¹ (Fig. 11), the local model
545 consisted of the set of 2D Fourier components bounded by the size of the compartment and the global
546 model consisted of the set of 2D Fourier components bounded by the size of the environment. “Model fit”
547 was measured for each eigenvector by finding maximum correlation over all model components between
548 the eigenvector and model component.

549 Code availability

550 These results were generated using code written in MATLAB. If you are interested in accessing the code,
551 you can email the corresponding author and we will be happy to make it available.

552 References

- 553 1. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty based competition between prefrontal and dorsolateral
554 striatal systems for behavioral control. *Nature Neuroscience* **8**, 1704–1711 (2005).
- 555 2. Tolman, E. C. Cognitive maps in rats and men. *Psychological Review* **55**, 189–208 (1948).
- 556 3. Schultz, W., Dayan, P. & Montague, P. A neural substrate of prediction and reward. *Science* **275**,
557 1593–1599 (1997).
- 558 4. Dayan, P. Improving generalization for temporal difference learning: The successor representation.
559 *Neural Computation* **5**, 613–624 (1993).
- 560 5. Stachenfeld, K. L., Botvinick, M. & Gershman, S. J. Design principles of the hippocampal cognitive
561 map. In *Advances in Neural Information Processing Systems 27*, 2528–2536 (MIT Press, 2014).
- 562 6. O’Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. preliminary evidence from unit
563 activity in the freely-moving rat. *Brain Research* **34**, 171–175 (1971).
- 564 7. O’Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Oxford: Clarendon Press, 1978).
- 565 8. McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M. B. Path integration and the
566 neural basis of the ‘cognitive map’. *Nature Reviews Neuroscience* **7**, 663–678 (2006).
- 567 9. Muller, R. U., Stead, M. & Pach, J. The hippocampus as a cognitive graph. *The Journal of General*
568 *Physiology* **107**, 663–694 (1996).
- 569 10. Penny, W., Zeidman, P. & Burgess, N. Forward and backward inference in spatial cognition. *PLoS*
570 *Computational Biology* **9**, e1003383 (2013).
- 571 11. Rueckert, E., Kappel, D., Tanneberg, D., Pecevski, D. & Peters, J. Recurrent spiking networks solve
572 planning tasks. *Scientific reports* **6** (2016).
- 573 12. Foster, D., Morris, R. & Dayan, P. A model of hippocampally dependent navigation, using the
574 temporal difference learning rule. *Hippocampus* **10**, 1–16 (2000).
- 575 13. Gustafson, N. J. & Daw, N. D. Grid cells, place cells, and geodesic generalization for spatial
576 reinforcement learning. *PLoS Computational Biology* **7**, e1002235 (2011).
- 577 14. Hafting, T., Fyhn, M., Molden, S., Moser, M. B. & Moser, E. I. Microstructure of a spatial map in the
578 entorhinal cortex. *Nature* **436**, 801–806 (2005).
- 579 15. Sutton, R. & Barto, A. *Reinforcement Learning: An Introduction* (MIT Press, 1998).

- 580 **16.** Gershman, S., Moore, C., Todd, M., Norman, K. & Sederberg, P. The successor representation and
581 temporal context. *Neural Computation* **24**, 1553–1568 (2012).
- 582 **17.** Gläscher, J., Daw, N., Dayan, P. & O’Doherty, J. States versus rewards: dissociable neural prediction
583 error signals underlying model-based and model-free reinforcement learning. *Neuron* **66**, 585–595
584 (2010).
- 585 **18.** Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M. M. Neural
586 representations of events arise from temporal community structure. *Nature neuroscience* **16**, 486–492
587 (2013).
- 588 **19.** Howard, M. & Kahana, M. A distributed representation of temporal context. *Journal of Mathematical*
589 *Psychology* **46**, 269–299 (2002).
- 590 **20.** Lisman, J. & Redish, A. D. Prediction, sequences and the hippocampus. *Philosophical Transactions*
591 *of the Royal Society of London B: Biological Sciences* **364**, 1193–1201 (2009).
- 592 **21.** Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered
593 goals. *Nature* **497**, 74–79 (2013).
- 594 **22.** Burgess, N., Recce, M. & O’Keefe, J. A model of hippocampal function. *Neural Networks* **7**,
595 1065–1081 (1994).
- 596 **23.** Dordek, Y., Meir, R. & Derdikman, D. Extracting grid characteristics from spatially distributed place
597 cell inputs using non-negative PCA. *eLife* (2015).
- 598 **24.** Mehta, M. R., Barnes, C. A. & McNaughton, B. L. Experience-dependent, asymmetric expansion of
599 hippocampal place fields. *Proceedings of the National Academy of Sciences* **94**, 8918–8921 (1997).
- 600 **25.** Mehta, M., Quirk, M. & Wilson, M. Experience-dependent asymmetric shape of hippocampal
601 receptive fields. *Neuron* **25**, 707–715 (2000).
- 602 **26.** Mehta, M. & McNaughton, B. ’expansion and shift of hippocampal place fields: Evidence for synaptic
603 potentiation during behavior. In *Computational Neuroscience: Trends in Research*, 741–745 (Plenum
604 Press, 1997).
- 605 **27.** Muller, R. U. & Kubie, J. L. The effects of changes in the environment on the spatial firing of
606 hippocampal complex-spike cells. *The Journal of Neuroscience* **7**, 1951–1968 (1987).
- 607 **28.** Skaggs, W. & McNaughton, B. Spatial firing properties of hippocampal CA1 populations in an
608 environment containing two visually identical regions. *The Journal of Neuroscience* **18**, 8455–8466
609 (1998).
- 610 **29.** Alvernhe, A., Save, E. & Poucet, B. Local remapping of place cell firing in the Tolman detour task.
611 *European Journal of Neuroscience* **33**, 1696–1705 (2011).
- 612 **30.** Hollup, S., Molden, S., Donnett, J., Moser, M. & Moser, E. Accumulation of hippocampal place fields
613 at the goal location in an annular watermaze task. *Journal of Neuroscience* **21**, 1635–1644 (2001).
- 614 **31.** Kjelstrup, K. *et al.* Finite scale of spatial representation in the hippocampus. *Science* **321**, 140–143
615 (2008).
- 616 **32.** Strange, B., Witter, M., Lein, E. & Moser, E. Functional organization of the hippocampal longitudinal
617 axis. *Nature Reviews Neuroscience* **15**, 655–669 (2014).
- 618 **33.** Sutton, R. Td models: Modeling the world at a mixture of time scales. In *Proceedings of the 12th*
619 *International Conference on Machine Learning* (1995).

- 620 **34.** Modayil, J., White, A. & Sutton, R. Multi-timescale nexting in a reinforcement learning robot.
621 *arXiv:1112.1133 [cs]* (2011).
- 622 **35.** Schapiro, A. C., Turk-Browne, N., Norman, K. & Botvinick, M. Statistical learning of temporal
623 community structure in the hippocampus. *Hippocampus* **26**, 3–8 (2016).
- 624 **36.** Garvert, M. M., Dolan, R. J. & Behrens, T. E. A map of abstract relational knowledge in the human
625 hippocampal–entorhinal cortex. *eLife* e17086 (2017).
- 626 **37.** Deuker, L., Bellmund, J., Schröder, T. & Doeller, C. An event map of memory space in the
627 hippocampus. *eLife* **5**, e16534 (2016).
- 628 **38.** Franzius, M., Sprekeler, H. & Wiskott, L. Slowness and sparseness lead to place, head-direction, and
629 spatial-view cells. *PLoS Computational Biology* **3**, 3287–3302 (2007).
- 630 **39.** Krupic, J., Bauza, M., Burton, S., Barry, C. & O’Keefe, J. Grid cell symmetry is shaped by
631 environmental geometry. *Nature* **518**, 232–235 (2015).
- 632 **40.** Derdikman, D. *et al.* Fragmentation of grid cell maps in a multicompartiment environment. *Nature*
633 *Neuroscience* **12**, 1325–1332 (2009).
- 634 **41.** Carpenter, F., Manson, D., Jeffery, K., Burgess, N. & Barry, C. Grid cells form a global representation
635 of connected environments. *Current Biology* **25**, 1176–1182 (2015).
- 636 **42.** Witter, M. P., Wouterlood, F. G., Naber, P. A. & Van Haeften, T. Anatomical organization of the
637 parahippocampal-hippocampal network. *Ann N Y Acad Sci* **911**, 1–24 (2000).
- 638 **43.** Dolorfo, C. L. & Amaral, D. G. Entorhinal cortex of the rat: topographic organization of the cells of
639 origin of the perforant path projection to the dentate gyrus. *Journal of Computational Neurology* **398**,
640 25–48 (1998).
- 641 **44.** Stensola, H. *et al.* The entorhinal grid map is discretized. *Nature* **492**, 72 – 78 (2012).
- 642 **45.** Mazumder, R., Hastie, T. & Tibshirani, R. Spectral regularization algorithms for learning large
643 incomplete matrices. *Journal of Machine Learning Research* **11**, 2287–2322 (2010).
- 644 **46.** Mahadevan, S. & Maggioni, M. Proto-value functions: A Laplacian framework for learning represen-
645 tation and control in markov decision processes. *Journal of Machine Learning Research* **8**, 2169–2231
646 (2007).
- 647 **47.** Bonnevie, T. *et al.* Grid cells require excitatory drive from the hippocampus. *Nature Neuroscience* **16**,
648 309–317 (2013).
- 649 **48.** Hales, J. *et al.* Medial entorhinal cortex lesions only partially disrupt hippocampal place cells and
650 hippocampus-dependent place memory. *Cell Reports* **9**, 893–901 (2014).
- 651 **49.** Muessig, L., Hauser, J., Wills, T. & Cacucci, F. A developmental switch in place cell accuracy
652 coincides with grid cell maturation. *Neuron* **86**, 1167–1173 (2015).
- 653 **50.** Şimşek, Ö., Wolfe, A. & Barto, A. Identifying useful subgoals in reinforcement learning by local
654 graph partitioning. In *Proceedings of the 22nd International Conference on Machine Learning*,
655 816–823 (ACM, 2005).
- 656 **51.** Solway, A. *et al.* Optimal behavioral hierarchy. *PLoS Computational Biology* **559** (2014).
- 657 **52.** Ribas-Fernandes, J. *et al.* A neural signature of hierarchical reinforcement learning. *Neuron* **71**,
658 370–379 (2011).

- 659 **53.** Shi, J. & Malik, J. Normalized cuts and image segmentation. In *Pattern Analysis and Machine*
660 *Intelligence, IEEE Transactions on*, vol. 22, 888–905 (IEEE, 2000).
- 661 **54.** Schlesiger, M. *et al.* The medial entorhinal cortex is necessary for temporal organization of hippocam-
662 pal neuronal activity. *Nature Neuroscience* **18**, 1123–1132 (2015).
- 663 **55.** Blum, K. & Abbott, L. A model of spatial map formation in the hippocampus of the rat. *Neural*
664 *Computation* **8**, 85–93 (1996).
- 665 **56.** Mehta, M. Neuronal dynamics of predictive coding. *Neuroscientist* **7**, 490–495 (2001).
- 666 **57.** Levy, W. B., Hocking, A. B. & Wu, X. Interpreting hippocampal function as recoding and forecasting.
667 *Neural Networks* **18**, 1242–1264 (2005).
- 668 **58.** Hassabis, D. & Maguire, E. A. The construction system of the brain. *Philosophical Transactions of*
669 *the Royal Society B: Biological Sciences* **364**, 1263–1271 (2009).
- 670 **59.** Buckner, R. L. The role of the hippocampus in prediction and imagination. *Annual Review of*
671 *Psychology* **61**, 27–48 (2010).
- 672 **60.** Barry, C., Hayman, R., Burgess, N. & Jeffery, K. Experience-dependent rescaling of entorhinal grids.
673 *Nature Neuroscience* **10**, 682–684 (2007).
- 674 **61.** Dolan, R. J. & Dayan, P. Goals and habits in the brain. *Neuron* **80**, 312–25 (2013).
- 675 **62.** Fanselow, M. From contextual fear to a dynamic view of memory systems. *Trends in Cognitive*
676 *Sciences* **14**, 7–15 (2010).
- 677 **63.** Wiltgen, B. J., Sanders, M. J., Anagnostaras, S., Sage, J. & Fanselow, M. S. Context fear learning in
678 the absence of the hippocampus. *The Journal of Neuroscience* **26**, 5484–5491 (2006).
- 679 **64.** Maurer, A. P. & McNaughton, B. L. Network and intrinsic cellular mechanisms underlying theta
680 phase precession of hippocampal neurons. *Trends in Neuroscience* **30**, 325–333 (2007).
- 681 **65.** Hasselmo, M. E. & Stern, C. E. Theta rhythm and the encoding and retrieval of space and time.
682 *NeuroImage* **85**, 656–666 (2014).
- 683 **66.** Johnson, A. & Redish, A. Neural ensembles in ca3 transiently encode paths forward of the animal at
684 a decision point. *Journal of Neuroscience* **27**, 12176–12189 (2007).
- 685 **67.** Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a
686 simple function of experience. *Neuron* **65**, 695–705 (2010).
- 687 **68.** Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S. & Redish, A. D. Segmentation of spatial
688 experience by hippocampal theta sequences. *Nature Neuroscience* **15**, 1032–1039 (2012).
- 689 **69.** P., M. A., N., B. S., Lipa, P., Skaggs, W. E. & Barnes, C. A. Greater running speeds result in altered
690 hippocampal phase sequence dynamics. *Hippocampus* **22**, 737–747 (2012).
- 691 **70.** van der Meer, M. A. A., Johnson, A., Schmitzer-Torbert, N. C. & Redish, A. D. Triple dissociation
692 of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial
693 decision task. *Neuron* **67**, 25–32 (2010).
- 694 **71.** Pezzulo, G., van der Meer, M. A., Lansink, C. S. & Pennartz, C. M. Internally generated sequences in
695 learning and executing goal-directed behavior. *Trends in Cognitive Sciences* **18**, 647–657 (2014).
- 696 **72.** Sanders, H., Rennó-Costa, C., Idiart, M. & Lisman, J. Grid cells and place cells: An integrated view
697 of their navigational and memory function. *Trends in Neurosciences* **38**, 763–775 (2015).

- 698 **73.** Ekstrom, A., Meltzer, J., McNaughton, B. & Barnes, C. Nmda receptor antagonism blocks experience-
699 dependent expansion of hippocampal “place fields”. *Neuron* **31**, 631–638 (2001).
- 700 **74.** Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.-B. & Moser, E. I. Hippocampus-independent phase
701 precession in entorhinal grid cells. *Nature* **453**, 1248–1252 (2008).
- 702 **75.** Middleton, S. J. & McHugh, T. J. Silencing ca3 disrupts temporal coding in the ca1 ensemble. *Nat*
703 *Neurosci* **19**, 945–951 (2016).
- 704 **76.** Sutton, R. S. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART*
705 *Bulletin* **2**, 160–163 (1991).
- 706 **77.** Silver, D. *et al.* Mastering the game of go with deep neural networks and tree search. *Nature* **529**,
707 484–489 (2016).
- 708 **78.** Zhang, J., Springenberg, J. T., Boedecker, J. & Burgard, W. Deep reinforcement learning with
709 successor features for navigation across similar environments. *CoRR* **abs/1612.05533** (2016).
- 710 **79.** Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J. & Daw, N. D. Predictive
711 representations can link model-based reinforcement learning to model-free mechanisms. *bioRxiv*
712 (2017).
- 713 **80.** Momennejad, I. *et al.* The successor representation in human reinforcement learning. *bioRxiv* (2017).
- 714 **81.** Howard, M., Fotedar, M., Datey, A. & Hasselmo, M. The temporal context model in spatial navigation
715 and relational learning: toward a common explanation of medial temporal lobe function across
716 domains. *Psychological Review* **112**, 75–116 (2005).
- 717 **82.** Krupic, J., Burgess, N. & O’Keefe, J. Neural representations of location composed of spatially
718 periodic bands. *Science* **337**, 853–857 (2012).
- 719 **83.** Bialek, W. *Biophysics: Searching for Principles* (Princeton University Press, 2012).
- 720 **84.** Weng, J., Zhang, Y. & Hwang, W. Candid covariance-free incremental principal component analysis.
721 *IEEE Trans. Pattern Anal. Mach. Intell.* **25**, 1034–1040 (2003).

722 **Acknowledgments**

723 We are grateful to Tim Behrens, Ida Mommenejad, and Kevin Miller for helpful discussions, and to
724 Alexander Mathis and Honi Sanders for comments on an earlier draft of the paper. This research was
725 supported by the NSF Collaborative Research in Computational Neuroscience (CRCNS) Program Grant
726 IIS-120 7833 and The John Templeton Foundation. The opinions expressed in this publication are those of
727 the authors and do not necessarily reflect the views of the funding agencies.

728 **Author contributions statement**

729 All authors conceived the model and wrote the manuscript. Simulations were carried out by K.S.

730 **Additional information**

731 The authors declare no competing financial interests.

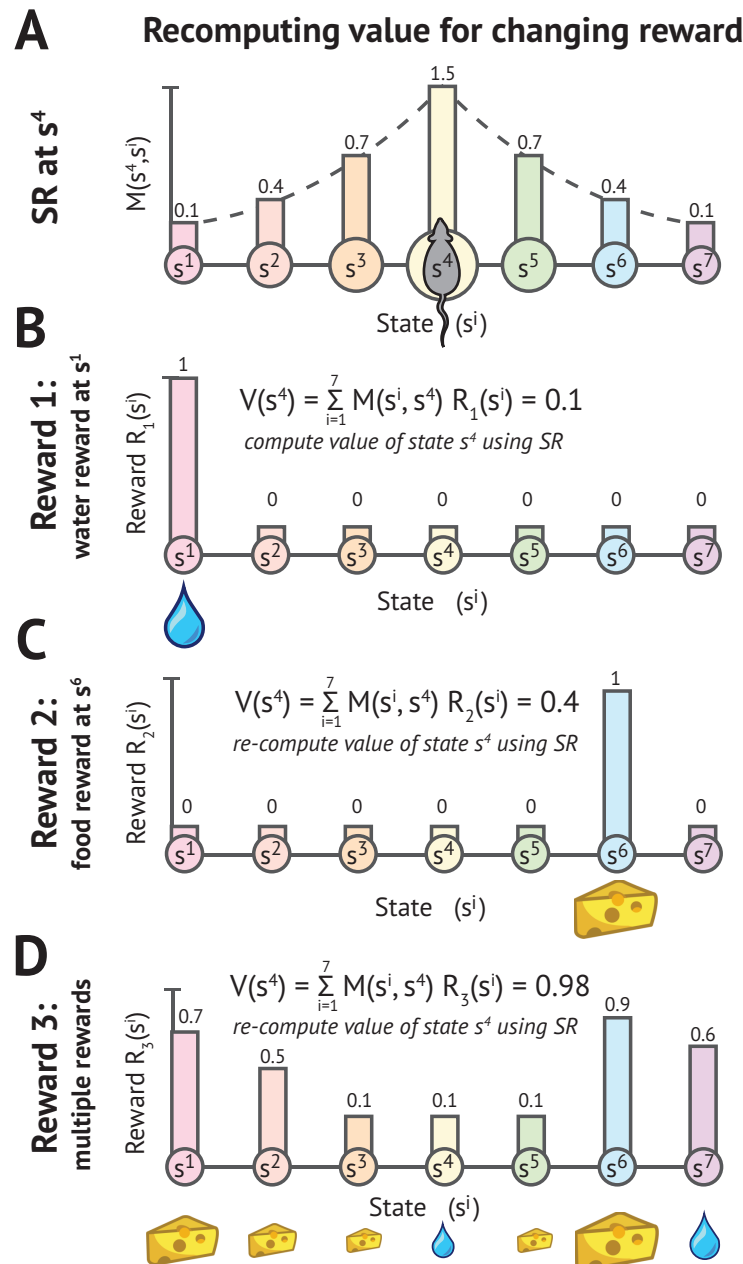


Figure 1. Updating value with the SR following change in reward. Since the representations of state and reward are decoupled, value functions can be rapidly recomputed for new reward functions without changing the SR. As formally defined in Equation 3, $M(s, s')$ gives the expected number of visits to state s' given a current location of s . Panel A shows the successor representation of state s^4 , which corresponds to a row $M(s^4, :)$ of the SR matrix. Panels B-D show how the value of s^4 changes under different reward functions.

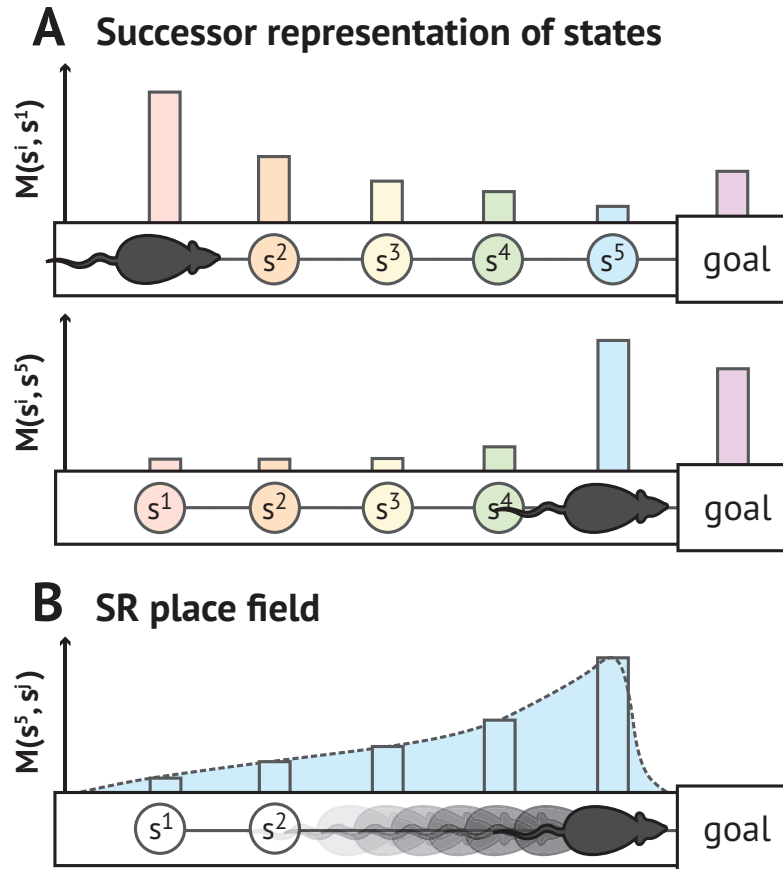


Figure 2. *Illustration of SR encoding population and individual SR place fields.* Under a prospective representation such as the SR, the population vector will be asymmetrically expanded in the direction of travel toward upcoming states. The place fields for individual cells will skew backwards. (A) A neural population encodes a prospective representation such that the firing rate of each cell is proportional to the discounted expected number of times its preferred state will be visited in the future. This population code is skewed toward upcoming states. Each colored bump represents the firing rate of a different place field located along the track. The value $M(s, s')$ is formally defined in Equation 3 as the expected number of visits to state s' given a current location of s , and the population vectors $M(s, :)$ illustrated here correspond to rows of the SR matrix. (B) The place field for a single SR-encoding cell skews backward toward past states that predict the cell's preferred state. When the blue state s^5 is visited, it becomes associated with all past states that predicted it. This automatically assigns credit for upcoming reward to preceding states. The receptive field $M(:, s')$ illustrated here corresponds to a column of the SR matrix.

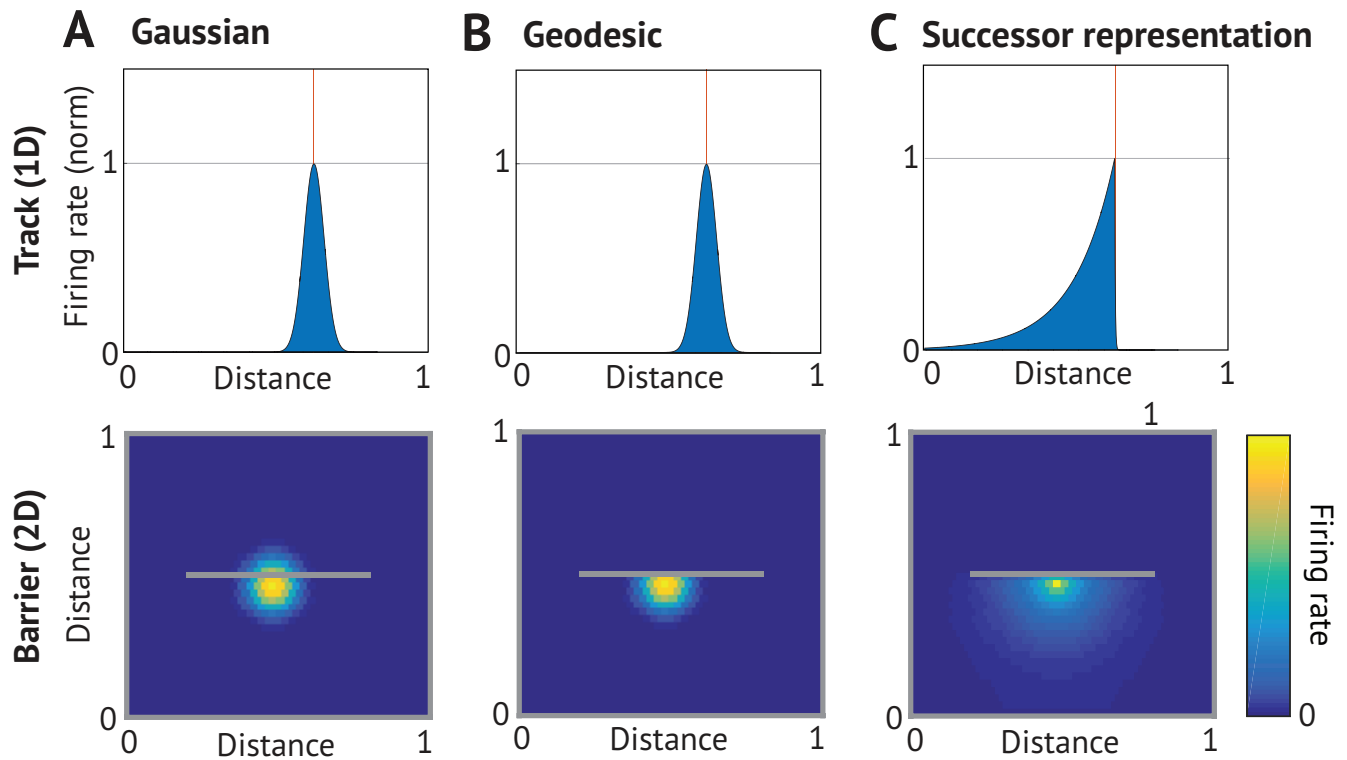


Figure 3. *Comparison of place cell models.* (Top) 1D track with left-to-right preferred direction of travel, red line marking field center; (bottom) 2D environment with a barrier indicated by gray line. (A) *Gaussian place field.* Firing of place cells decays with Euclidean distance from the center of the field regardless of experience and environmental topology. (B) *Topological place field.* Firing rate decays with geodesic distance from the center of the field, which respects boundaries in the environment but is invariant to the direction of travel¹³. (C) *SR place field.* Firing rate is proportional to the discounted expected number of visits to other states under the current policy. On the directed track, fields will skew opposite the direction of motion to anticipate the upcoming successor state. Since the policy will not permit traversing walls, successor fields warp around obstacles.

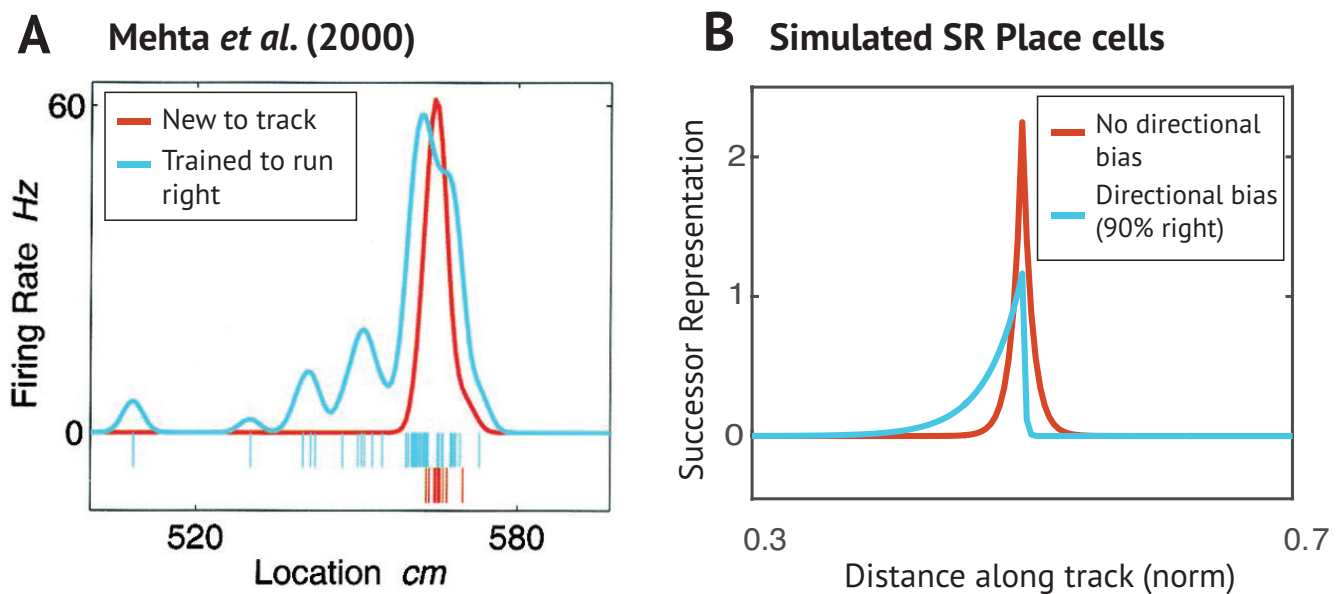


Figure 4. *Predictive skewing of place fields.* (A) As a rat is trained to run repeatedly in a preferred direction along a narrow track, initially symmetric place cells (red) begin to skew (blue) opposite the direction of travel²⁵. (B) When transitions in either direction are equally probable, SR place fields are symmetric (red). Under a policy in which transitions to the right are more probable than to the left, simulated SR place fields skew opposite the direction of travel toward states predicting the preferred state (blue).

Alvernhe *et al.* (2011) recordings from Tolman detour maze




Figure 5. *Place fields near detour.* (A) Maze used by Alvernhe and colleagues²⁹ for studying how place cell firing is affected by the insertion of barriers in a Tolman detour maze. Reward is delivered at location B. “Near” and “Far” zones are defined. In “early” and “late” detour conditions, a clear barrier blocks the shortest path, forcing the animal to take the short detour to the left or the longer detour to the right. (B) Example CA1 place fields recorded from a rat navigating the maze. (C) Over the population, place fields near the barrier changed their shape, while the rest remained unperturbed. This is shown by computing the Fisher z transformed spatial correlation between place field activity maps with and without barriers present. (D) The environment used to simulate the experimental results. (E) Example SR place fields near to and far from the barrier, before and after barrier insertion. More fields are shown in Fig. S6. (F) When barriers are inserted, SR place fields change their fields near the path blocked by the barrier and less so at more distal locations where policy is unaffected. The effect is more pronounced in the early detour condition because the detour appears closer to the start.

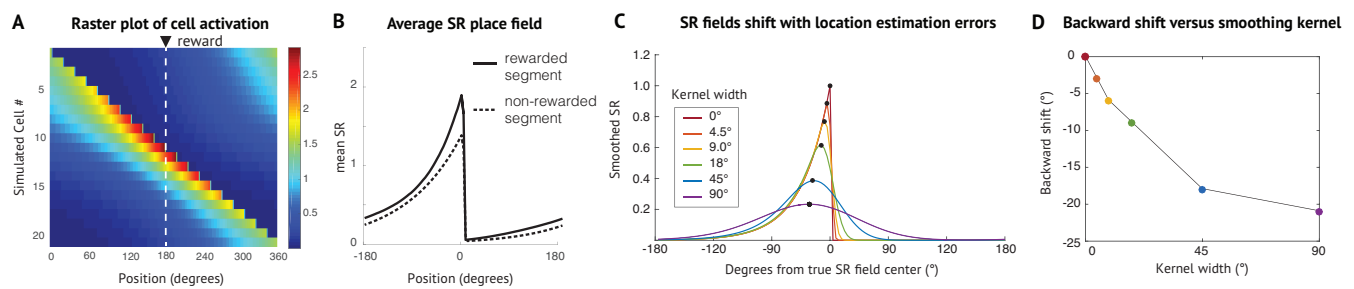
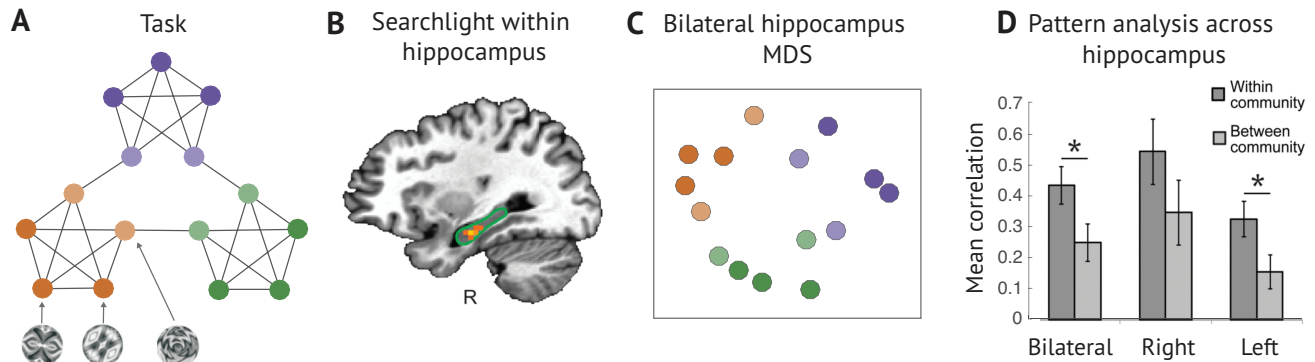


Figure 6. *Distribution of place fields in annular maze with reward.* (A) Simulated SR raster for annular watermaze. The transition model assumes that the animal spends more time near the rewarded platform and that the animal must move counter-clockwise (shown above as right-to-left) to get the reward. For this simulation, the probability of moving clockwise is 0. (B) The average SR place field in the rewarded and unrewarded segments. The states near the reward are visited more, so the SR model predicts more firing near these rewarded locations and the states that precede them. This difference is smaller when the discount factor is smaller. (C) When location is uncertain, the SR becomes smoother and the peak shifts toward the center of mass. For this reason, an asymmetric firing field may be accompanied by a backward migration of the firing field. (D) The magnitude of the shifts become more pronounced as the uncertainty distribution over possible locations of the animal becomes wider. For a given discount, the magnitude of the shift is bounded by distance between the SR field's center of mass and the encoded state.

Schapiro *et al.* (2015)



SR Simulations




Figure 7. *Hippocampal representations in non-spatial task.* (A) Schapiro *et al.*³⁵ showed subjects sequences of fractal stimuli drawn from the task graph shown, which has clusters of interconnected nodes (or “communities”). Nodes of the same color fall within the same community, with the lighter colored nodes connecting to adjacent communities. (B) A searchlight within hippocampus showed a stronger within-community similarity effect in anterior hippocampus. (C, D) States within the same cluster had a higher degree of representational similarity in hippocampus, and multidimensional scaling (MDS) of the hippocampal BOLD dissimilarity matrix captured the community structure of the task graph³⁵. (E) The SR matrix learned on the task. The block diagonal structure means that states in the same cluster predict each other with higher probability. (F) Multidimensional scaling of dissimilarity between rows of the SR matrix reveals the community structure of the task graph. (G) Consistent with this, the average within-community SR state similarity is consistently higher than the average between-community SR state similarity.

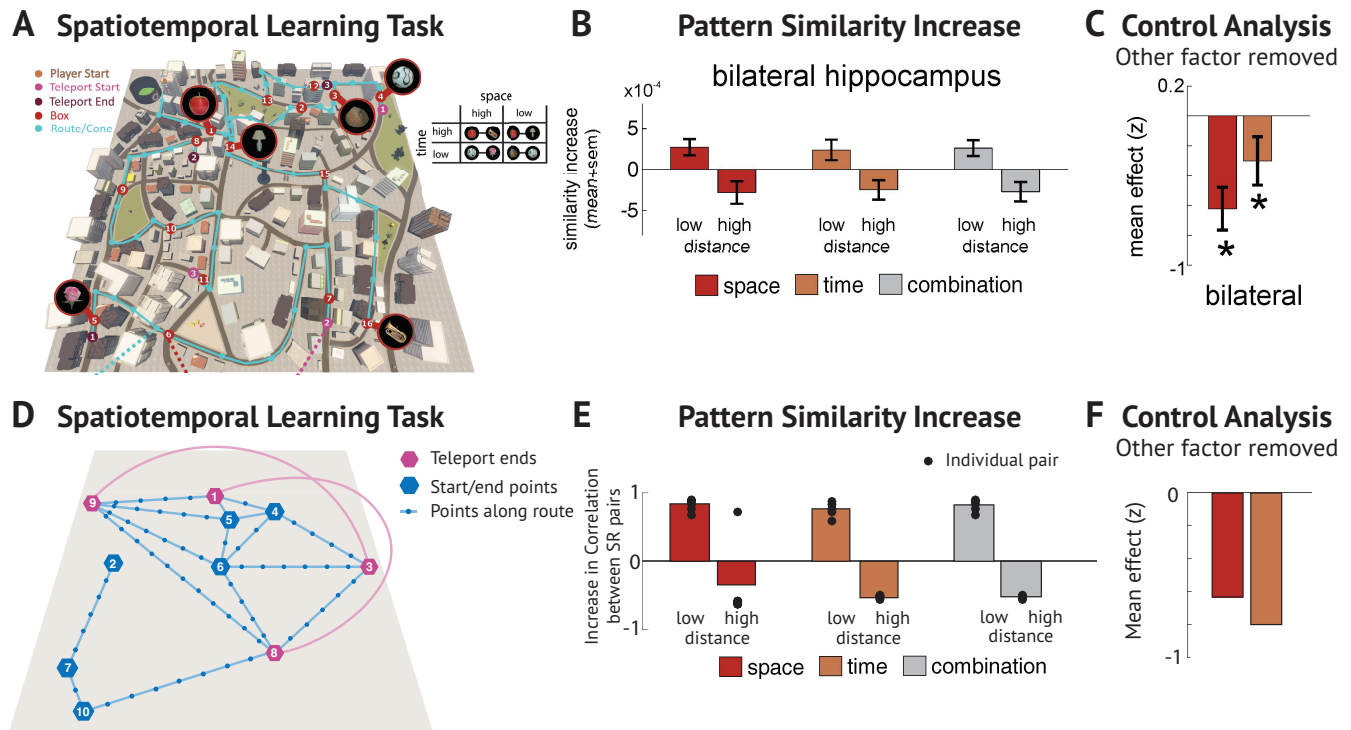


Figure 8. *Hippocampal representations in spatio-temporal task.* (A) Deuker *et al.*³⁷ trained subjects on a spatio-temporal navigation task. Subjects were told to objects scattered about the map. It is possible to take a “teleportation” shortcut between certain pairs of states (pink and purple), and other pairs of states are sometimes joined only by a long, winding path. Nearness in time is therefore partially decoupled from nearness in space. (B) The authors find significant increase in hippocampal representational similarity between nearby states and a decrease for distant states. This effect holds when states are nearby in space, time, or both. (C) Since spatial and temporal proximity are correlated, the authors controlled for the each factor and measured the effect of the remaining factor on the residual. (D-F) Simulation of experimental results in panels A-C.

Krupic *et al.* (2015) Effects of environmental geometries

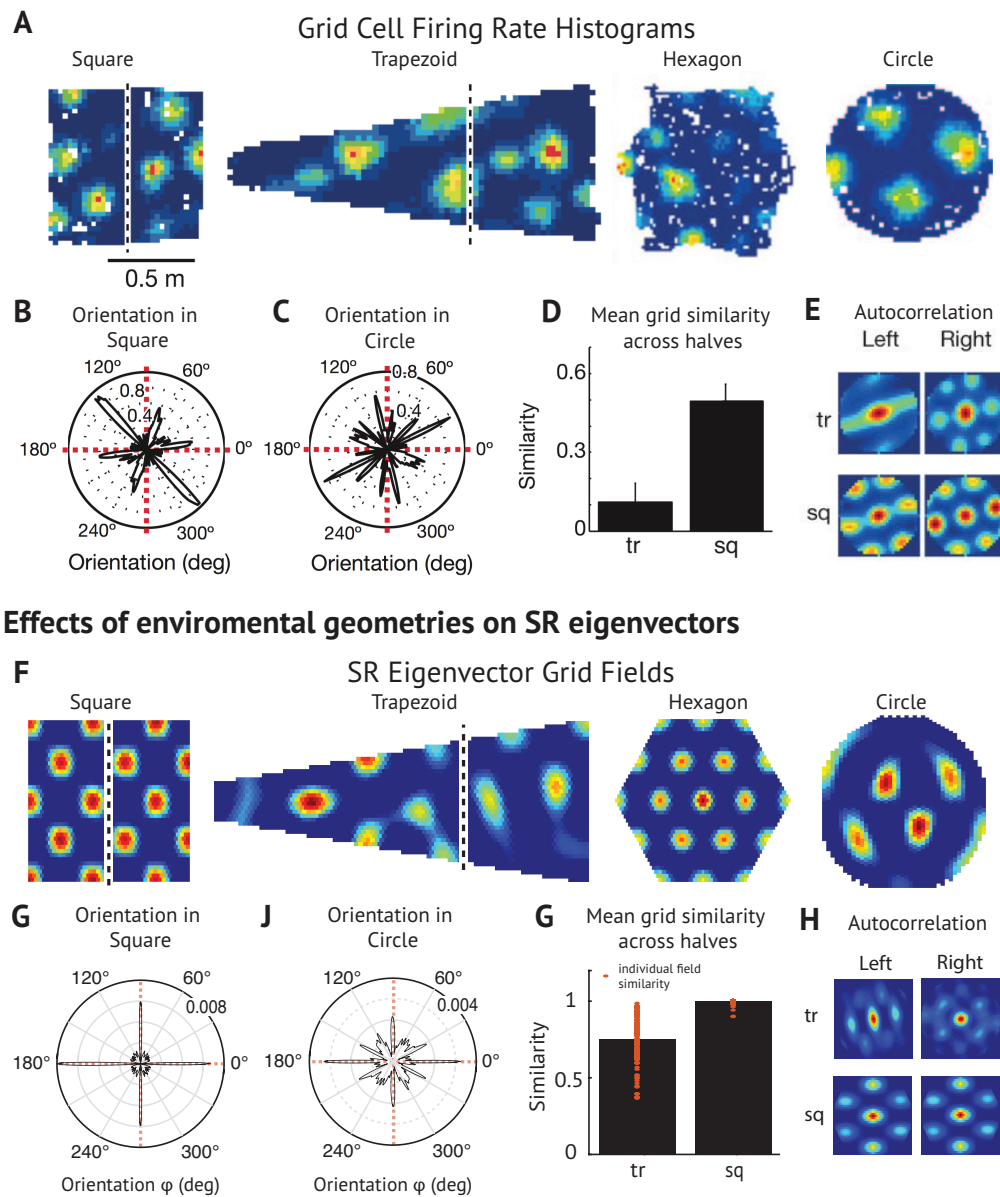


Figure 9. *Grid fields in geometric environments.* (A) Grid fields recorded in a variety of geometric environments³⁹. Grid fields in trapezoid and square environments are split at the dividing line shown for split-halves analysis. (B,C) Grid fields in the square environment had more consistent orientations with respect to boundaries and distal cues than in the square environment. (D) While grid fields tend to be similar on both halves of a square (sq) environment, they tend to be less similar across halves of the irregular trapezoidal (tr) environment. (E) Autocorrelograms for different halves of trapezoidal and square environments in circular windows used for split-halves anal. (F-H) Simulations of experimental results in panels A-E.

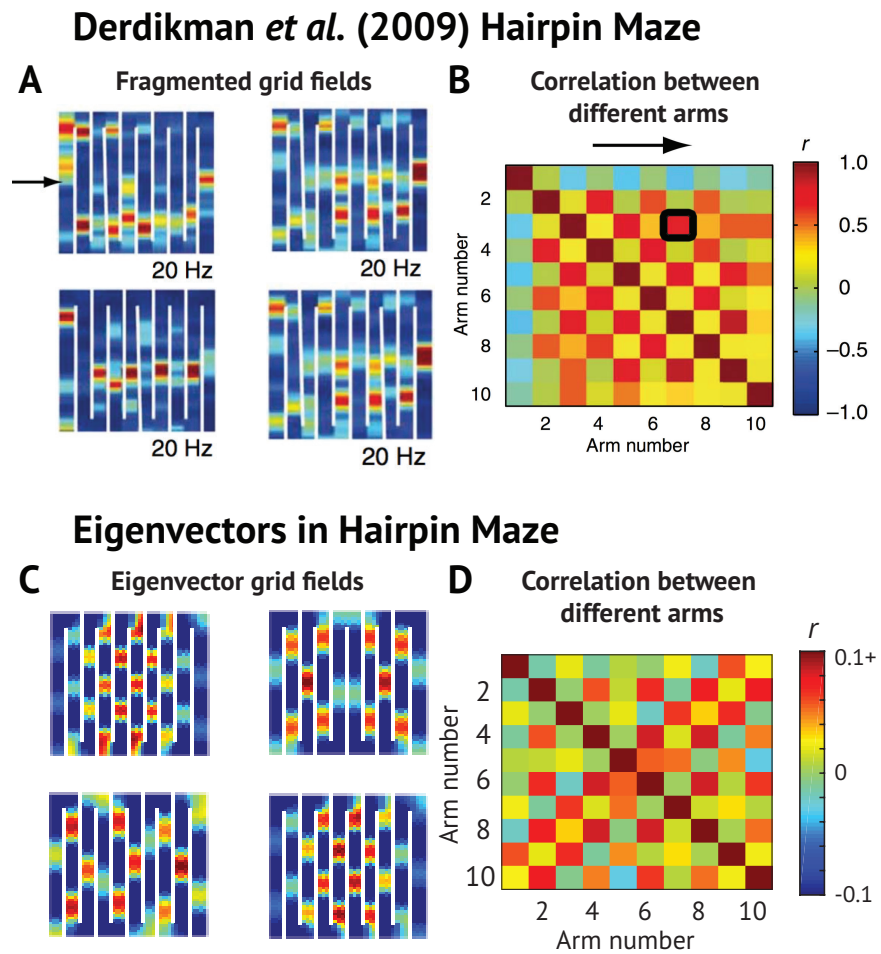
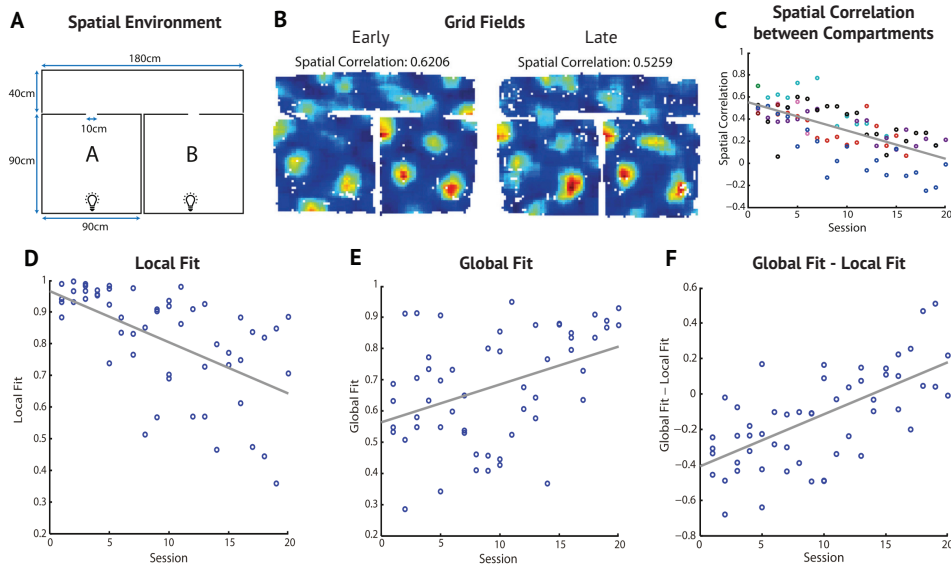


Figure 10. *Grid fragmentation in compartmentalized maze.* (A) Barriers in the hairpin maze cause grid fields to fragment repetitively across arms⁴⁰. (B) Spatial correlation between activity in different arms. The checkerboard pattern emerges because grid fields frequently repeat themselves in alternating arms. (C-D) Simulations of the experimental results in panels A-B.

Carpenter *et al.* (2015) Grid fields in multi-compartment environment



Eigenvector Grid fields learned in multi-compartment environment

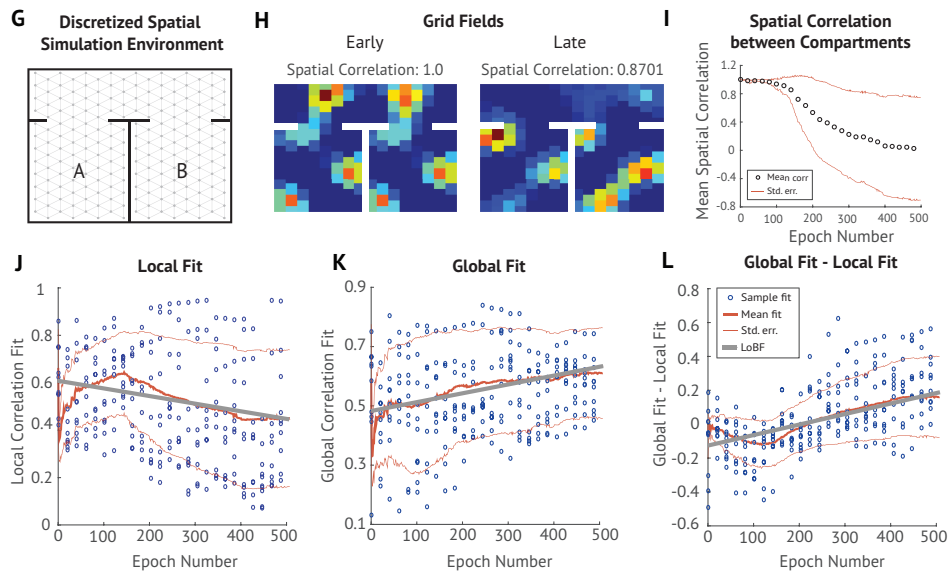


Figure 11. *Grid fields in multi-compartment environment.* (A) Multi-compartment environment employed by Carpenter and colleagues⁴¹. (B) Example grid fields early and late in training. (C) Spatial correlation between grid fields in compartments A and B across sessions. (D-F) To explain this decline in inter-compartment similarity, Carpenter and colleagues fit a local model (grid constrained to replicate between the two compartments) and a global model (single continuous grid spanning both compartments). They found that the local fit decreased across sessions, while the global fit increased, and correspondingly the difference between the two models increased. (G-L) Simulation of experimental results in panels A-F. In I-J, the blue circles indicate individual samples, the thick red line denotes the mean, the thin red lines denote one standard deviation from the mean, and the thick gray lines are lines of best fit.