

Hippocampal spatio-temporal cognitive maps adaptively guide reward generalization

Mona M. Garvert^{1,3,*}, Tankred Saanum², Eric Schulz², Nicolas W. Schuck³, and Christian F. Doeller^{1,4,5}

¹Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

²Max Planck Institute for Biological Cybernetics, Tübingen, Germany

³Max Planck Institute for Human Development, Berlin, Germany

⁴Kavli Institute for Systems Neuroscience, Centre for Neural Computation, The Egil and Pauline Braathen and Fred Kavli Centre for Cortical Microcircuits, Jebsen Centre for Alzheimer's Disease, NTNU, Trondheim, Norway

⁵Institute of Psychology, Leipzig University, Leipzig, Germany

*garvert@cbs.mpg.de

ABSTRACT

The brain forms cognitive maps of relational knowledge, an organizing principle thought to underlie our ability to generalize and make inferences. However, how can a relevant map be selected in situations where a stimulus is embedded in multiple relational structures? Here, we find that both spatial and temporal cognitive maps influence generalization in a choice task, where spatial location determines reward magnitude. Mirroring behavior, the hippocampus not only builds a map of spatial relationships but also encodes temporal distances. As the task progresses, participants' choices become more influenced by spatial relationships, reflected in a strengthening of the spatial and a weakening of the temporal map. This change is driven by orbitofrontal cortex, which represents the evidence that an observed outcome is generated from the spatial rather than the temporal map and updates hippocampal representations accordingly. Taken together, this demonstrates how hippocampal cognitive maps are used and updated flexibly for inference.

Introduction

As humans we live in complex, ever-changing environments that often require us to select appropriate behaviors in situations never faced before. Luckily, our environment is replete with statistical structure and our experiences are rarely isolated events¹. This allows us to predict outcomes that were never experienced directly by generalizing information acquired about one state of the environment to related ones². Indeed, humans and other animals generalize across spatially or perceptually similar stimuli³⁻⁶ as well as across stimuli forming associative structures such as those acquired in a sensory preconditioning task^{7,8}. Generalization also occurs in reinforcement learning tasks where the same latent state determines the outcome associated with choosing different stimuli^{9,10}.

For generalization to be possible, an appropriate neural representation of stimulus relationships is required. Many studies have shown that spatial relationships, such as distances between landmarks, are represented in a hippocampal cognitive map^{11,12}, which enables flexible goal-directed behavior beyond simple stimulus-response learning¹³. More recently, it has been suggested that the same organizing principle might also underlie the representation of relationships between non-spatial states such as perceptual¹⁴⁻¹⁹ or temporal relationships between stimuli²⁰⁻²², or associative links between objects²³⁻²⁶. Interestingly, cognitive maps even

form incidentally and in the absence of conscious awareness²³. This suggests that the hippocampus automatically extracts the embedding of a stimulus in multiple relational structures²⁷, even for stimulus features that are not directly task-relevant²⁸.

If stimuli are part of multiple relational structures such as space and time, this raises the question how the representation that is most beneficial for reward maximisation and generalization can be selected²⁹. One region implicated in this process is the orbitofrontal cortex (OFC), known to represent task states in situations where these are not directly observable^{24,30}. Little is known, however, about how information in the OFC about the task-relevance of different maps relates to corresponding changes in the representation of cognitive maps in the hippocampus^{31,32}.

Here, we combined virtual reality with computational modeling and functional magnetic resonance imaging (fMRI) to show that participants represent spatial as well as temporal stimulus relationships in hippocampal maps. The degree to which each map was represented neurally determined the degree to which it was used for generalization in a subsequent choice task, even though only the spatial location determined the magnitude of rewards. Notably, the neural representation of each map and its influence on choice changed over the course of the choice task through an OFC signal reflecting the evidence that the spatial rather than the temporal map

66 caused the observed outcome. Together, our results provide a
67 computational and neural mechanism for the representation
68 and adaptive selection of hippocampal cognitive maps during
69 choice.

70 Results

71 Participants used knowledge about stimulus relationships to generalize value

72 To examine how humans use information about stimulus relationships for generalization and inference, forty-eight healthy
73 human participants (mean age 26.8 ± 3.8 years, 20–34 years
74 old, 27 male) took part in a 3-day experiment that involved
75 first learning to locate 12 monster stimuli in a virtual arena,
76 followed by a choice task in which spatial knowledge could
77 be used for predicting rewards (Figure 1a).

78 On day 1, participants performed multiple exploration
79 blocks in which they were instructed to remember the location
80 of the stimuli while freely navigating in the arena (Figure
81 1c, d). Stimuli became visible when they were approached,
82 but were otherwise invisible. Exploration policies differed
83 substantially between individuals (Figure 2a, Supplementary
84 Figure S1). As a result, participants experienced different
85 temporal relations between the monsters, which could also
86 deviate from the spatial distances between objects. For example,
87 some participants visited objects in a stereotyped order,
88 whereas others navigated mostly around the border of the
89 arena or systematically scanned the environment from top to
90 bottom (Figure 2a).

91 After each exploration block, participants performed an
92 object location memory task. Participants were first teleported
93 to a random location in the arena and instructed to then
94 navigate to the hidden location of a presented object.
95 Feedback indicated how far away the current position was
96 from the correct stimulus location. The session terminated
97 when the replacement error averaged across all monsters in a
98 block was below $3vm$ ($vm =$ virtual meter; $3vm$ correspond to
99 $< 10\%$ of the arena’s diameter) and at least five and at most
100 ten blocks had been completed. At the end of the learning
101 phase, participants could position the stimuli in the correct
102 location (Supplementary Figure S2a). Before and after each
103 imaging session on days 2 and 3, participants also performed
104 one block of the object location memory task without feedback.
105 The replacement error did not differ between sessions
106 (Figure 2b), confirming that no new learning took place. In
107 a spatial arena task at the end of the 3-day study, participants
108 also accurately reproduced the stimulus arrangement when
109 instructed to drag-and-drop stimuli imagining a top-down
110 view on the spatial arena (Figure 1g). Participants thus
111 learned the spatial arrangement of the stimuli well.

112 In a choice task performed in the MRI scanner on day
113 3, participants were presented with two stimuli simultaneously
114 and instructed to select the one that was associated
115 with a higher reward (Figure 1f). Participants were told
116 that the reward magnitude was determined by the stimulus
117 location in space (Figure 1a). Participants did initially

120 not know which locations were rewarding, but they could
121 combine their knowledge about the stimulus relationships
122 with previously experienced reward contingencies to infer
123 the rewards of stimuli they had not yet experienced. In
124 order to decorrelate spatial distance and reward relationships,
125 we introduced two contexts with different reward
126 distributions (Figure 1a). Participants performed alternating
127 choice blocks for each context, with the context signaled
128 by the background color. Participants learned to perform
129 the task rapidly (Figure 2c) and their choices were a
130 function of the difference in value between the stimuli presented
131 on the left and the right on the screen in both contexts
132 (Figure 2d, context 1: $t(47) = 10.0, p < 0.001$, context 2:
133 $t(47) = 12.1, p < 0.001$).

134 To test whether participants could use their knowledge
135 about the stimulus relationships to generalize, two stimuli
136 per context were never presented during the choice task (“inference
137 stimuli”, Figure 1a, b). A value rating at the end of the study
138 (Figure 1g) showed that participants were able to infer which
139 of the two inference stimuli had a higher value in each context
140 (Figure 2e; repeated measures ANOVA, $F(1, 46) = 21.4, p < 0.001$),
141 reflecting that they combined their knowledge about the stimulus
142 location with knowledge about associated rewards of nearby stimuli.
143 The error between the true inference values and the value ratings
144 was larger in participants who struggled to reproduce the spatial
145 map as indicated by a larger error between the true z-scored
146 spatial distances and the z-scored distances in the arena task
147 (“Map reproduction error”, $r = 0.37, p = 0.01$,
148 Figure 2f). This demonstrates that participants exploited
149 knowledge about stimulus relationships to infer unseen values.
150

151 Cognitive maps of spatial and temporal stimulus relationships explain generalization

152 The fact that participants could successfully infer the values of the
153 inference stimuli suggests that they formed a representation of the
154 stimulus relationships. But stimulus relationships were learned during
155 free exploration, which was typically non-random and differed
156 substantially between participants (Figure 2a, Supplementary Figure
157 S1). This means that the experienced temporal distances between
158 the objects differed meaningfully from their spatial distances in
159 most participants (Supplementary Figures S6 and S4b). Intelligent
160 agents should keep track of both the spatial distance as well as
161 the temporal relationships between objects, since either feature
162 may become relevant for generalization. We therefore reasoned
163 that the brain may extract two relational maps: one reflecting
164 spatial distances between stimuli and the other one reflecting
165 temporal relationships.
166

167 To test explicitly to what extent generalization was guided by
168 the spatial or temporal maps – or a combination of both – we
169 fitted Gaussian process (GP) models to participants’ choices
170 (see Online Methods). The GP predicts rewards for a novel
171 stimulus based on the rewards associated




Figure 1. Experimental design. **a** Spatial position of monsters during the navigation tasks and value distribution associated with the same monsters in context 1 and 2 in the choice task. Darker colors indicate higher values. Numbered circles indicate the location of inference objects that were never presented during the choice task. **b** True values of the four inference objects. **c** Tasks performed on the three subsequent days, see text. **d** Exploration and object location memory tasks. In the exploration task, participants navigated around a virtual arena with button presses corresponding to forward, backward, right and left movements. Monsters appeared when they were approached, but were never all visible at the same time. In the object location memory task, participants were instructed to navigate to the position of a cued monster (each monster cued once in each block). Feedback indicated how far away the positioned object was from the correct object location. On day 1, participants performed between five and ten blocks (depending on performance) of the exploration and the object location memory task in alternation. On subsequent days, only one block of the object location memory task was performed before and after scanning without feedback. **e** Picture viewing task performed in the scanner. Participants were presented with monsters one after another. When two monsters appeared, participants were instructed to choose the monster that was closer in space to the preceding monster (map symbol) or the monster that was more similar in value to the preceding monster (coins symbol, day 3 only). On day 2, the background color was irrelevant for the task, on day 3 it indicated the context determining the stimulus values. **f** Choice task performed in the scanner. Participants were instructed to maximize accumulated points by choosing the monster associated with a higher reward. Participants were told that the monsters had different values in two different contexts, and that the relevant context was signalled by the background color. The values associated with each monster in the two contexts were learned in alternation, with ten blocks of context 1 followed by ten blocks of context 2, and so forth. **g** At the end of day 3, four post-tests were performed: Participants indicated for each monster how many points they would receive in each of the two contexts and how much they liked each monster. They were then asked to arrange the monsters in terms of their similarity in a circle in such a way that monsters that were considered similar were positioned near each other (Arena task 1). Lastly, participants were instructed to imagine a top-down view of the arena they had navigated around and to place the monsters in the corresponding location (Arena task 2).




Figure 2. Behavioral results. **a** Trajectories of three example participants during the exploration phase on day 1. Purple dots indicate the stimulus locations and black lines the participant trajectories. See all participants' trajectories in Supplementary Figure S1. **b** replacement error for days 2 and 3, before (pre) and after (post) the scanning session. The replacement error was defined as the Euclidean distance between the true location and the drop location. The replacement error did not differ significantly between those four sessions (all $p > 0.05$), see object positioning at the end of the learning phase on day 1 in Supplementary Figure S2. **c** Percent correct of choices over the course of the choice task. Trials are divided into ten sub-blocks of ten trials each with a constant context. **d** Probability of choosing the right option as a function of the difference in value between the right and the left option, separately for each context. **e** Value rating for the inference stimuli at the end of the study. **f** Correlation between the the map reproduction error (root-mean-square error between the true z-scored spatial distances and the z-scored distances in the arena task) and the root-mean-square error for the inference ratings. Data in **b**, **c** and **e** are plotted as group-level whisker-boxplots (center line, median; box, 25th to 75th percentiles; whiskers, $1.5 \times$ interquartile range; error bars in **d** denote standard error of the mean. Circles and transparent lines in **b-f** represent individual participant data. ** $p < 0.01$, ns = not significant




Figure 3. Modeling results suggest that participants generalized over spatial and temporal stimulus relationships. **a** Graph representation corresponding to the three example exploration paths in Figure 2a. **b** Model comparison: Model frequency represents how often a model prevailed in the population. The error bars represent the standard deviation of the estimated Dirichlet distribution. The winning model generalizes values according to a combination of spatial and temporal relationships between stimuli. **c** Inference performance as predicted by the model. Depicted are the inferred values for the inference objects in analogy to the participant ratings in 2e. **d** Relationship between inference error predicted by the model and actual inference error in participants' value ratings. **e** Model comparison for the value ratings for the inference objects at the end of the study. The winning model generalizes values according to a combination of spatial and temporal relationships between stimuli. **f** Correlation between the spatial and temporal effects on choice behavior. **g** Correlation between the relative spatial weight as estimated by the model and inference error. Data in **c** are plotted as group-level whisker-boxplots (center line, median; box, 25th to 75th percentiles; whiskers, $1.5 \times$ interquartile range; crosses, outliers). Circles and transparent lines represent individual participant data. ** $p < 0.01$, P_{exc} = exceedance probability.

ated with all other stimuli, weighted by their similarity to the novel stimulus. Since the similarity function determines how the GP generalizes, we can express hypotheses about what cognitive map participants use by pairing GPs with similarities implied by spatial or temporal maps.

Specifically, generalizing using a spatial cognitive map corresponds to pairing the GP with a similarity function that decays with Euclidean distance. Generalizing using a temporal cognitive map corresponds to pairing the GP with a similarity function that decays with temporal distance. We constructed these temporal similarities based on individual participants' navigation runs from day 1: Using their stimulus visitation history from the exploration phase, we computed each participants' successor representation³³, reflecting the expected number of visits of any stimulus s' given a starting stimulus s . This can be transformed into a probability that two stimuli are visited in direct succession (see Online Methods). We then computed temporal similarities based on the diffusion distance^{6,34} implied by these transition probabilities (Figure 3a).

Finally, kernel functions can be added or multiplied together³⁵ to model function learning where generalization may be guided by a combination of multiple similarity functions^{36,37}. As such, the hypothesis that both the spatial and temporal maps guide generalization together is captured in the spatio-temporal GP, which uses the additive composition of the *spatial* and the *temporal* similarities to generalize.

To test which map best explained how participants generalized rewards, we created three GP models that generalized based on either spatial, temporal or spatio-temporal relationships between monsters. Then, for each trial, we made each GP model predict the reward of both monsters, conditioning the GPs on all monster-reward pairs observed in the relevant context up to that point. We also compared these models to a "mean tracker" model that assumes participants only learn about directly experienced stimulus-reward associations, without generalization (see Online Methods).

To fit our models to participants' choices, we entered the predicted difference in reward between the two presented monsters in a mixed-effect logistic regression model with random slopes per participant³⁸, and determined the maximum likelihood hyper-parameters using grid search. We then computed model frequency based on the leave-one-out cross-validated log-likelihood (leaving one trial out) for each model³⁹.

The model generalizing based on the compositional, spatio-temporal similarities explained participants' choices best (Figure 3b; model frequency = 0.681, $XP > 0.999$, see Supplementary Figure S3 for full modeling results). This model performed substantially better than the temporal model (model frequency = 0.08), the spatial model (model frequency = 0.23) and the mean tracker (model frequency = 0.005). The model also reproduced the difference in value rating for the high- and the low-inference stimuli (Figure 3c; repeated measures ANOVA, $F(1, 47) = 2602.3, p < 0.001$).

Across participants, the root-mean-square error between true values and values predicted by the winning model was highly correlated with the root-mean-square error between the true values and the value ratings provided by participants (Figure 3d, $r = 0.85, p < 0.001$).

Furthermore, participants' value ratings for the inference objects at the end of the study were also predicted best by a spatio-temporal model (Figure 2e). This demonstrates that behavior in two independent parts of the study, the choice task and the inference test, was influenced by both spatial and temporal knowledge about stimulus relationships. Notably, the value ratings for the stimuli whose values could be directly sampled were best predicted by the mean tracker model, rather than the spatio-temporal GP (Supplementary Figure S3a). This suggests that participants evoked specific memories of stimulus-reward associations where possible, but relied on the spatio-temporal map when they needed to construct values of stimuli which were not directly experienced (Supplementary Figure S3c).

We estimated effect sizes for the spatial and the temporal component as the participant-specific random effects in a model where the spatial and temporal predictors competed to explain variance in participants' choices. Spatial weights were defined as the relative contribution of the spatial compared to the temporal predictor. Both the spatial and the temporal relationships had non-zero influence on choice behavior and the effect sizes were negatively correlated (Figure 3f, $r = -0.45, p = 0.001$), suggesting that participants tended to rely predominantly on one of the two maps for guiding choice. Consistent with the fact that the spatial, but not the temporal relationships, were relevant for generalization, participants whose choices were driven more by the spatial relationships compared to the temporal ones performed better in the inference test (Figure 3g, $r = -0.43, p = 0.003$).

Spatial and temporal stimulus relationships represented in the hippocampal system influence choice

Our modeling results suggest that participants generalized values based on both the spatial and temporal relationships experienced between stimuli during the exploration phase. To investigate the neural representation of these relationships, we scanned participants before the choice task on day 2 and after the choice task on day 3 using fMRI. During these imaging sessions, stimuli on the two background colors were presented in random order (Figure 2e). Once after each stimulus on each background color (i.e. in 24 of 144 trials), participants were presented with two stimuli and instructed to either report which one was closer in space or more similar in value in the given context (on day 3 only) to the preceding stimulus. Participants performed this task well above chance (correct performance on day 2: $81 \pm 10\%$ (distance judgement); day 3: $78 \pm 12\%$ (distance judgement) and $68 \pm 14\%$ (value judgement), mean \pm standard deviation,




Figure 4. Spatial and temporal cognitive maps in the hippocampal formation are related to generalization and inference. **a** Whole-brain analysis showing a cross-stimulus enhancement effect in the scanning session after the choice task (session 3) that scales with spatial distance. For illustration purposes, voxels thresholded at $p > .01$ (uncorrected) are shown; only the right hippocampal cluster survives for multiple comparisons. **b** Correlation between the spatial cross-stimulus enhancement effect extracted from the right hippocampal ROI depicted in **a** (thresholded at $p < 0.001$) and the spatial effects governing decisions in the choice task. **c** Correlation between the spatial cross-stimulus enhancement effect extracted from the right hippocampal ROI depicted in **a** and the root-mean-square error between ratings for the inference stimuli and their true value. **d** Correlation between temporal cross-stimulus enhancement effect extracted from the right hippocampal ROI depicted in **a** and the temporal effects governing decisions in the choice task. **e** Correlation between the temporal cross-stimulus enhancement effect extracted from the right hippocampal ROI depicted in **a** and the root-mean-square error between ratings for the inference stimuli and their true value. **f** Whole-brain analysis where spatial effects (top) and temporal effects (bottom) describing generalization during choice are entered as second-level covariates for the spatial and temporal cross-stimulus enhancement effects. Both analyses reveal significant clusters in the hippocampal formation. **g** Whole-brain analysis where the inference error is entered as second-level covariate for the spatial and temporal cross-stimulus enhancement effects. This analysis reveals a negative effect for the spatial map and a positive effect for the temporal map in the hippocampal formation. **h** Mediation path diagram for inference error as predicted by the hippocampal map and spatial effects. **a**, **f** and **g** are thresholded at $p < 0.01$, uncorrected for visualization. ** $p < 0.01$; *** $p < 0.001$

all $p < 0.001$).

We used fMRI adaptation^{40,41} to investigate the representational similarity of the 12 stimuli. This technique uses the amount of suppression or enhancement observed when two stimuli are presented in direct succession as a proxy for the similarity of the underlying neural representations. In line with previous work demonstrating similar effects for graph-like structures²³, we hypothesized that in regions encoding a cognitive map of the stimulus relationships, the size of the cross-stimulus adaptation effect should scale with spatial or temporal distance between stimuli. Based on previous work, we expected the hippocampal formation to be a candidate region for representing such cognitive maps²³ and therefore focus on a bilateral region comprising the hippocampus, the entorhinal cortex and the subiculum (see mask used for small-volume correction in Supplementary Figure S5). We tested for adaptation effects by including spatial and temporal distances as parametric modulators in the same general linear model (GLM).

We found a significant cross-stimulus enhancement effect that scaled with spatial distance in session 3 (after the choice task) in the right hippocampal formation (Figure 4a, $t(47) = 3.86, p = 0.045, [24, -28, -16]$). A cluster in the left hippocampal formation trended in the same direction ($t(47) = 3.63, p = 0.08, [-12, -36 - 6]$). No voxels survived the conservative correction procedure for the temporal relations. One reason for this could be that different participants represented the spatial and temporal aspects to different degrees, with a stronger representation of the spatial map across the group as a whole. Indeed, in most participants (44 out of 48), the spatial component contributed more to generalization during choice than the temporal component ($t(47) = 9.9, p < 0.001$). We therefore investigated whether the strength of the neural representation predicted the degree to which an individual was influenced by either spatial or temporal distances in the choice task.

To test this, we first extracted parameter estimates for the spatial and temporal maps from the above-identified region of interest (ROI) in the right hippocampal formation showing a cross-stimulus enhancement effect that scaled with spatial distance (Figure 4a, masking threshold $p < 0.001$). A significant correlation with the spatial and temporal effects on choice behavior confirmed a relationship the neural representation of the respective maps in this region and generalization behavior (Figure 4b, d), spatial: $r = 0.37, p = 0.01$, temporal: $r = 0.40, p = 0.005$). We also found that the representation of the spatial, but not the temporal map in this ROI can be linked to performance in the later, independent inference test that depended on spatial knowledge (spatial: $r = -0.44, p = 0.002$, temporal: $r = 0.06, p = 0.7$, Figure 4c, e).

To investigate whether the relationship between spatial and temporal influences on behavior and neural map representation is specific to the hippocampus, we included spatial and temporal effects on choice behavior as covariates

on the second level in the GLM that was used to identify spatial and temporal cross-stimulus enhancement effects above. For both spatial and temporal maps we found precisely localized clusters in the hippocampal formation, where the spatial and temporal fMRI effects were larger the stronger the respective map's influence on behavior (Figure 4f, spatial: $t(47) = 4.45, p = 0.009, [22, -28, -18]$, temporal: $t(47) = 4.19, p = 0.02, [26, -20, -28]$, $t(47) = 4.14, p = 0.02, [28, -14, -16]$ and $t(47) = 3.91, p = 0.04, [-28, -16, -13]$). Furthermore, the representation of the spatial map in the hippocampus was stronger and the representation of the temporal map was weaker in individuals who made smaller inference errors (Figure 4g, spatial: $t(47) = 5.08, p = 0.002, [32, -14, -25]$ and $t(47) = 4.95, p = 0.002, [-32, -14, -22]$, temporal: $t(47) = 4.53, p = 0.007, [-32, -12, -25]$). This suggests that participants who represented the spatial map more strongly in the hippocampal formation also generalized more according to spatial distances in the choice task and performed better in the inference task, with the reverse pattern for the temporal relationships.

To test whether the hippocampal spatial map formally mediated the impact of the neural representation on inference performance, we related the parameter estimates for the spatial map extracted from the right hippocampal ROI to both the spatial effects as estimated from behavior in the choice task as well as the inference performance using single-level mediation^{42,43}. The path model jointly tests the relationship between the neural representation of the spatial map and the degree to which spatial relationships influenced generalization in the choice task (path a), the relationship between spatial weights in the choice task and inference performance (path b), and a formal mediation effect (path ab) that indicates that each explains a part of the inference performance effect while controlling for effects attributable to the other mediator. All three effects were significant (path $a = 0.26, SE = 0.10, p = 0.01$; path $b = -3.40, SE = 0.92, p = 0.003$; path $ab = -0.86, SE = 0.42$; path $c = -1.07, SE = 0.45, p = 0.02$; path $c' = -1.93, SE = 0.54, p < 0.001$, Figure 4h). This confirms that it is the representation of a hippocampal cognitive map that is critical for guiding spatial generalization and inference during the choice task and the inference test. Furthermore, despite the fact that the spatial and the temporal kernel were correlated in most participants (average Pearson's $r = 0.58 \pm 0.12$), the neural effect as well as the degree to which behavior was influenced by either component could not be explained by a correlation between spatial and temporal kernels (Supplementary Figure S6).

The representation of cognitive maps adapts to the task demands

In the choice task, rewards associated with the monsters were determined by their location in space and participants who had a better neural representation of the spatial map




Figure 5. Hippocampal cognitive maps adapted to the task demands. **a** Logistic functions for each participant fitted to how individual spatial weights changed over trials. Curves are colored according to a participant's relative performance in the choice task. **b** Correlation between the slopes of the estimated logistic function depicted in **a** and the inference error. **c** Correlation between the slopes of the logistic function and the change in the hippocampal spatial enhancement effect extracted from the ROI depicted in Figure 4a). **d** Change in spatial and temporal enhancement cross-stimulus enhancement effects in the ROI depicted in Figure 4a). Because the ROI was defined based on the existence of a spatial enhancement effect in session 3, the spatial effect is biased and displayed for visualization only. **e** Correlation between the change in the hippocampal spatial and temporal enhancement effects. Both were extracted from the ROI depicted in Figure 4a). **f** Whole-brain analysis depicting the update in spatial weights at the time of feedback. **g** Whole-brain analysis depicting voxels where the increase in the spatial cross-stimulus enhancement effect across participants correlates with the size of the hippocampal spatial weight update during the choice task as shown in **f**. **h** Whole-brain analysis depicting voxels where the difference in unsigned prediction errors as computed based on the temporal versus the spatial map correlates with with the size of the hippocampal spatial weight update during the choice task as shown in **f**. **i** Mediation path diagram for the change in the hippocampal spatial cross-stimulus enhancement effect extracted from the ROI depicted in Figure 4a as predicted by the OFC evidence integration signal and the hippocampal spatial weight update. **f-h** are thresholded at $p < 0.01$, uncorrected for visualization. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

392 performed better in the inference tasks (Figure 4e). Yet, 447
393 we also found evidence for a lingering effect of experienced 448
394 temporal distances on choice. 449

395 We hypothesized that individuals adjust the degree to 450
396 which they rely on one over the other map for guiding choice 451
397 depending on the observed outcome contingencies. Indeed, 452
398 a logistic function fitted to how individual weights changed 453
399 over trials showed that in most participants the temporal 454
400 component explained generalization behavior in the choice 455
401 task better initially, but as the choice task progressed, spatial 456
402 knowledge became more influential (Figure 5a). The slope 457
403 of this logistic function was particularly steep for participants 458
404 who performed better in the choice task (Figure 5a) as well 459
405 as in the inference test (Figure 5b, $r = -0.44$, $p = 0.002$). 460

406 We reasoned that this might reflect changes in the rep- 461
407 resentation of the neural map over the course of the choice 462
408 task. If this is the case, then participants who showed a larger 463
409 increase in the contribution of spatial knowledge on choices, 464
410 i.e. a steeper slope in the logistic regression, should also 465
411 show a larger increase in the neural representation of the 466
412 spatial map from day 2 (before the choice task) to day 3 467
413 (after the choice task). To test this, we extracted parameter 468
414 estimates from the same region of interest we used for the 469
415 analyses in Figure 4 for session 2 (before the choice task) 470
416 and session 3 (after the choice task) and correlated the dif- 471
417 ference with the slope of the logistic function. The positive 472
418 relation we found suggests that participants whose behavior 473
419 in the choice task was characterized by marked increases in 474
420 the reliance on the spatial map during choice also showed 475
421 a larger increase in the neural representation of the spatial 476
422 map (Figure 5c, $r = -0.44$, $p = 0.002$). In the same region, 477
423 the temporal map decreased significantly across participants 478
424 (Figure 5d, $t(47) = -2.1$, $p = 0.04$) and the change in the 479
425 spatial map representation was negatively correlated with 480
426 the change in the temporal map representation (Figure 5e, 481
427 $r = -0.62$, $p < 0.001$), suggesting that in participants where 482
428 the spatial map representation became stronger, the temporal 483
429 map representation became weaker. 484

430 We reasoned that this change in representation might 485
431 be driven by a neural signal reflecting the degree to which 486
432 either map was task-relevant during the choice task. To test 487
433 this hypothesis, we first set up a GLM which included a 488
434 parametric regressor that reflected the difference in the de- 489
435 gree to which the spatial map influenced choice from one 490
436 trial to the next. This identified a region in the left hip- 491
437 pocampus tracking the trial-by-trial change in the degree 492
438 to which the spatial dimension guided choice (Figure 5f, 493
439 $t(47) = 4.14$, $p = 0.02$, $[-18, -32, -18]$). 494

440 If this neural weight update signal led to an increase in 495
441 the neural representation of the relevant map, then partic- 496
442 ipants with stronger hippocampal weight updating signals 497
443 should display a larger change in hippocampal representa- 498
444 tion of the spatial map from day 2 to day 3. To test where 499
445 the spatial weight updating signal correlated with a change 500
446 in the spatial map representation, we looked for changes 501

in the spatial map representation from session 2 to session 3 across the whole brain, and included the parameter estimates extracted from the hippocampal ROI reflecting the spatial weight update as a covariate. This analysis revealed a significant positive effect in the left hippocampal formation (Figure 5g, $p = 0.018$, $t(47) = 4.21$, $[18, -14, -25]$), suggesting that participants whose hippocampus tracked the spatial weight updates during the choice task also updated the representation of the spatial map in the hippocampus.

The changes in the composition of the hippocampal map likely reflect a representation learning process that was driven by the experienced reward contingencies in the choice task. To test whether any brain region tracks the evidence that the observed outcomes were generated by either of the two maps, we calculated the trial-wise unsigned prediction errors for each outcome separately for the spatial and the temporal map. The difference between these two prediction errors indicates how much more expected an outcome was according to the spatial as compared to the temporal map. We then set up a GLM that modeled this difference between spatial and temporal prediction errors at feedback time. Areas reflecting the evidence for the spatial over the temporal map should respond positively on trials where the spatial map made more accurate predictions than the temporal maps. We reasoned that, if there is a relationship between this signal and the spatial updating signal, then participants whose hippocampal weight updating signal was stronger should also show more of such an evidence tracking signal, and therefore included the parameter estimate extracted from the hippocampal ROI as a covariate. The only region where an evidence integration signal covaried with the hippocampal updating signal was the medial orbitofrontal cortex (Figure 5h, $p = 0.03$, $[14, 46, -13]$, family-wise error corrected on the cluster level).

In line with the observation that the OFC adapts behavior by changing associative representations in other brain regions⁴⁴, the orbitofrontal evidence signal may thus align task representation with observed outcomes. By signalling the degree to which either map is task-relevant, spatial weights may be updated during the choice task, which in turn leads to an update of the spatial map representation itself. To test this assumption, we investigated whether the spatial weight update in the hippocampus formally mediated the relationship between the evidence integration signal in the OFC and the hippocampal changes in the spatial map representation. The fact that the OFC signal and the hippocampal spatial weight update was significant (path $a = 0.7$, $SE = 0.3$, $p = 0.02$) is not surprising, since the ROI was identified based on voxels where the corresponding covariate explains some variance. However, the effect of the spatial weight updating signal on the change in representation remains significant if we control for the OFC signal (path $b = 14.0$, $SE = 4.6$, $p = 0.0003$). Furthermore, there is a relationship between the OFC signal and the change in hippocampal map representation (path $c = 13.1$, $SE = 6.6$, $p = 0.03$), which can be fully accounted

502 for by the hippocampal weight update (path $c' = 2.9, SE =$ 556
503 $6.5, p = 0.6$, path $ab = 10.23, SE = 5.71, p = 0.007$, Figure 557
504 4h). Hence, participants with the largest OFC evidence 558
505 integration signal at feedback time exhibited the largest updates 559
506 in spatial weights in the hippocampus, which in turn related 560
507 to a larger change in the neural representation of the spatial 561
508 map. This suggests a role for OFC signal in adjusting the 562
509 use of an appropriate map to the current task demands, and 563
510 an associated behavioral change. 564

511 Discussion 565

512 The hippocampal formation is known to organize relation- 566
513 ships between events in cognitive maps, thought to be critical 567
514 for generalization and inference. However, the neural and 568
515 computational mechanisms underlying the ability to use cog- 569
516 nitive maps for generalization remains unknown in situations 570
517 where stimuli are embedded in multiple relational structures. 571
518 Here, we combined virtual reality, computational modeling 572
519 and fMRI to demonstrate that the hippocampus extracts both 573
520 spatial and temporal stimulus relationships from experience 574
521 during navigation in a virtual arena. The strength of each 575
522 neural representation was related to the degree to which it 576
523 influenced behavior in an independent choice task. Notably, 577
524 the OFC tracked the evidence that outcomes observed in 578
525 the choice task were consistent with the predictions made 579
526 by the spatial and the temporal cognitive map, which led to 580
527 corresponding adjustments of the hippocampal map repre- 581
528 sentation. 582

529 Participants learned to locate stimuli in a virtual arena. 583
530 Because most individuals chose non-random behavioral poli- 584
531 cies for exploring the arena, stimulus relationships could 585
532 be characterized both in terms of spatial distance as well as 586
533 temporal co-occurrence. We found that the hippocampal for- 587
534 mation extracted both types of relationships and represented 588
535 those in clusters well-known to represent distances to goals⁴⁵, 589
536 goal direction signals⁴⁶ as well as associative distances be- 590
537 tween objects forming a non-spatial graph²³. Notably, the 591
538 degree to which either map was represented in this region 592
539 determined the degree to which participant's generalization 593
540 behavior in a later choice task was influenced by the cor- 594
541 responding map. This demonstrates a clear link between 595
542 hippocampal map representations and their use for guiding 596
543 generalization in decision making. It also shows that this 597
544 system efficiently deals with higher-dimensional relational 598
545 structures and can combine information from multiple maps 599
546 for guiding choice. 600

547 We found substantial inter-individual differences in terms 601
548 of the degree to which participants represented the spatial 602
549 and temporal relationships a stimulus was embedded in neu- 603
550 rally, and were influenced by those dimensions during choice. 604
551 Indeed, in participants whose choices were influenced by 605
552 the spatial or the temporal map, we found a cross-stimulus 606
553 enhancement effect for spatial or temporal stimulus relation- 607
554 ships, respectively. In participants whose choices were not 608
555 influenced by those dimensions, on the other hand, the oppo- 609
610

site was true: responses to a stimulus were suppressed if the 611
preceding stimulus was close in space or time. Often, repe- 612
tition suppression effects are more common than repetition 613
enhancement effects in fMRI adaptation paradigms⁴⁰. How- 614
ever, behavioral relevance can influence the directionality 615
of an fMRI adaptation effect. For example, while repetition 616
suppression effects are typically observed in the hippocam- 617
pus when a stimulus that is irrelevant for the task at hand 618
is repeated, repetition enhancement effects can be observed 619
in the same region when a stimulus is task-relevant⁴⁷. It 620
is therefore conceivable that what a participant considered 621
the relevant stimulus dimension was enhanced, while the 622
irrelevant dimension was suppressed. In the context of our 623
experiment, it was more adaptive to generalize along spatial 624
rather than temporal distances, since spatial distances were 625
used for creating reward contingencies in the first place. The 626
more a participant therefore succeeded in enhancing the spa- 627
tial dimension and suppressing the temporal dimension, the 628
better they performed in the task. 629

630 Furthermore, participant choices became increasingly 631
more influenced by spatial relational knowledge as the choice 632
task progressed, suggesting that which map is used for guid- 633
ing choice can be adaptively adjusted to the current task 634
demands. This effect was driven by an OFC evidence in- 635
tegration signal, indexing the difference in accuracy of the 636
predictions made by a spatial compared to a temporal model 637
at feedback time. Participants whose OFC responded more 638
strongly also showed a larger spatial weight updating signal 639
in the hippocampus at feedback, which was in turn related to 640
a stronger increase in the representation of the spatial map 641
from before to after the choice task. This suggests that the 642
OFC tracks the evidence that the currently observable state of 643
the world was driven by either of the two maps, and updates 644
the degree to which either influences behavior accordingly. 645

646 Our findings are consistent with the proposed function of 647
the OFC to represent state spaces, in particular in situations 648
where the current state of the world is not readily observable 649
and must be inferred⁴⁸. The OFC is also typically involved in 650
situations where participants need to adjust their behaviour 651
when outcome contingencies change³⁰ or when memory 652
responses require an arbitration between hippocampal and 653
striatal inputs⁴⁹. For example, reversal learning or outcome 654
devaluation, where previously acquired cue–outcome and 655
response–outcome associations need to be adapted, rely on 656
an intact OFC⁵⁰. 657

658 Importantly, our results also shed light on the interaction 659
between OFC and the hippocampus. In line with previous ob- 660
servations indicating a relation between state representations 661
in OFC and the hippocampus^{31,51,52}, our results indicate that 662
OFC might play an active role in learning state presentations 663
in the hippocampus through experience⁵³. Future experi- 664
ments should assess whether similar adjustments can also 665
be observed when temporal rather than spatial stimulus re- 666
lationships govern the reward distribution, or when rewards 667
are governed by a compositional, spatio-temporal map. 668

In conclusion, our results suggest that the hippocampus represents different dimensions of experienced relationships between stimuli such as space and time in parallel cognitive maps. The degree to which each one is used for guiding choice is governed by an OFC evidence integration signal. The OFC drives a spatial updating signal in the hippocampus, which is in turn related to a change in the representation of the spatial map. This provides a mechanistic insight into the way in which appropriate stimulus dimensions are selected for guiding decision making in multi-dimensional environments.

Online Methods

Participants

52 neurologically and psychiatrically healthy participants took part in this study (mean age 26.8 ± 3.8 years, 20-34 years old, 27 male). Participants were recruited using the participant database of the Max Planck Institute for Human Cognitive and Brain Sciences. Due to a scanner defect, three participants could not complete the last day. One participant was excluded due to problems during the preprocessing. 48 participants therefore entered the analyses. Two of those participants did not do the arena task at the end of the experiment, but their data was included in all other analyses. The study was approved by the ethics committee at the Medical Faculty at the University of Leipzig (221/18-ek) and all participants gave written informed consent prior to participation.

Experimental procedure

The experiment consisted of three parts performed on three subsequent days. On day 1, participants learned the stimulus distribution in a virtual arena. On day 2, we assessed the stimulus representation in the fMRI scanner. On day 3, participants performed a choice task to learn the rewards associated with each stimulus in the scanner. Afterwards, we again assessed the stimulus representations in the scanner. The sessions are described in more detail below. The exploration and object location memory task were coded using the virtual reality software package Vizard (Version 4, Santa Barbara, CA: WorldViz LLC). All other tasks were written in custom-written Matlab scripts using Psychtoolbox. Imaging data was preprocessed using fmriprep. Imaging and behavioural analyses were carried out with Matlab.

Day 1

Participants were first familiarized with the stimuli by being presented with the monsters one-by-one on the screen. They could click through the stimuli to proceed to the next one. Participants were then instructed that they would be asked to learn where each monster belongs in space, and that this knowledge would be important for collecting points in later sessions. Monsters were distributed in a circular arena with a virtual radius of 15m (Figure 1A). Which monster was presented in which location was randomized across participants.

5 distinct trees were located behind the wall surrounding the arena, which functioned as landmarks. The location of the trees was randomized in such a way that one tree occurred at a random position in every 72deg block in each participant. Tree locations were fixed across all experimental sessions.

Participants then learned the location of stimuli in space by navigating around a virtual arena (Figure 1E) in multiple blocks. Each block consisted of an exploration phase and an object location memory task. In the exploration phase, participants navigated around the arena in any way they liked and for as long as they wanted. Whenever a participant approached a monster (i.e. they entered a 3 m radius around the monster location), it became visible and slowly turned around its own axis. This means that participants never saw all monsters at the same time. After each exploration phase, participants performed an object location memory task. In this task, participants were cued with a monster and had to navigate to the corresponding location (Figure 1F). Feedback indicated how close to the correct location a monster was positioned (<3m, <5m, <7m, <9m, >9m). In each block, every monster had to be positioned once. The order was randomized. If performance reached a pre-specified performance criterion of <3m drop error averaged across all monsters (corresponding to <10% error) in a block, the session terminated if a participant had completed at least five blocks. Participants performed a minimum number of 5 and a maximum number of ten blocks of this task to ensure that they had a good knowledge of the stimulus distribution.

Day 2

Before the scanning session, participants had another opportunity to explore the monster locations freely, followed by one more round of the object location memory task with feedback.

Subsequently, we assessed the monster representations in the scanner using a picture viewing task. Here, participants were presented in the fMRI scanner with the monsters in a random order on a red or a blue background. Participants were instructed to view the images attentively. Occasionally (once after each monster on each background color), two monsters were presented simultaneously and participants had to indicate which of the two monsters was located closer in space to the monster they had seen immediately before the two monsters. Participants received no feedback. The purpose of this task was to ensure that participants would always evoke the location a monster was embedded in during the stimulus presentations. Correct answers were rewarded with 0.10 EUR. Participants were instructed that the background color was irrelevant for performing the task. Each monster was presented 6 times on each background color (red, blue) per block, resulting in 144 stimulus presentations in each block. Participants completed three blocks of this task. Stimulus sequences were generated pseudo-randomly using a genetic algorithm with the following constraints: Each stimulus in each context occurred the same number of times per block and no monster-monster transition was

718 presented more than once.

719 After the scanning session, another round of the object
720 location memory task was performed without feedback to
721 assess participants' memory for the monster locations.

722 **Day 3**

723 Before the scanning session, another round of the object
724 location memory task was performed without feedback to
725 assess participants' memory for the monster locations.

726 In the scanner, participants then performed a choice task.
727 Here, they were presented with pairs of monsters and in-
728 structed to select the monster that would lead to the highest
729 reward. The reward distribution was related to the position
730 of the monsters in space and the context as indicated by the
731 background color (Figure 1H). Participants were instructed
732 that they would receive similar amounts of points for mon-
733 sters located near each other in space. They learned the two
734 value distributions in a blocked fashion, with ten trials of
735 choices in context 1 alternating with ten trials of choices
736 in context 2. Background colors and contexts were coun-
737 terbalanced across participants. Value distributions were
738 selected such that pairwise spatial distances and pairwise
739 value differences across both contexts were not significantly
740 correlated and that the overall value across all objects was
741 similar across the two contexts.

742 Two objects in each context ('inference objects') could
743 never be chosen during the choice task (Figure 1B). These
744 were later used to assess whether participants were able to
745 combine information about rewards with information about
746 the relationship between monsters to infer stimulus values
747 that were never directly experienced. Critically, the value of
748 one inference object per context was high (71 and 72) and
749 the value of the other inference objects was low (3 and 13).

750 After the choice task, three blocks of the picture viewing
751 task were performed in the scanner. This time, the back-
752 ground colour indicated the relevant context and participants
753 were instructed to think about each monster's location in
754 space and its associated value. Occasionally (once after each
755 monster on each background color), two monsters were pre-
756 sented simultaneously and participants had to indicate which
757 of the two monsters was located closer in space to the mon-
758 ster they had seen immediately before the two monsters or
759 which monster had a more similar value. Which task was to
760 be performed was indicated with a symbol presented above
761 the two options. Correct answers were rewarded with 0.10
762 EUR. Stimulus sequences were the same as on day 2.

763 After the scanning session, another round of the object
764 location memory task was performed without feedback to
765 assess participants memory for the monster locations. This
766 was followed by four brief tasks. (1) Participants had to
767 indicate on a sliding scale from 0 to 100 how many points
768 they would receive for each monster in each context, (2)
769 Participants rated on a scale from "not at all" to "very much"
770 how much they liked each monster, (3) Participants arranged
771 monsters in an arena according to their similarity (Arena task
772 1), and (4) according to their spatial location (Arena task 2).

773 In each task, the order in which monsters were presented
774 was randomized across participants.

775 **Reimbursement**

776 Participants were paid a baseline fee of 9€/hour for the be-
777 havioral parts of the experiment and 10€/hour for the fMRI
778 sessions. In addition, participants could earn a monetary
779 bonus depending on performance. Points accumulated dur-
780 ing the choice blocks were converted into money (100 points
781 = 0.1€). Furthermore, each correct choice during the mon-
782 ster presentation block was rewarded with 0.10€.

783 **Behavioral analysis**

784 **Object positioning task.** The replacement error in the ob-
785 ject location memory task was defined as the Euclidean dis-
786 tance between the drop location and the true object location.
787 It was reported relative to the arena diameter.

788 **Choice task.** A correct choice was the choice corre-
789 sponding to the object with the higher value.

790 **Inference task.** The inference error was defined as the
791 root-mean-square error between the true inference values
792 and the error ratings provided by a participant at the end of
793 the study.

794 **Arena task.** The map reproduction error was defined as
795 the root-mean-square error between the true z-scored spa-
796 tial distances between the monsters in the virtual arena and
797 the z-scored distances between the monster positions in the
798 arena task. We z-scored the distances to ensure that they had
799 a comparable range.

800 **Modeling**

We used Gaussian process regression to model reward learn-
ing and generalization in the choice task. Gaussian pro-
cesses (GPs) define probability distributions over functions
 $f \sim \mathcal{N}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where $m(\mathbf{x})$ is the mean function,
giving the expected function values $\hat{\mathbf{y}}$ at input points \mathbf{x} , and
 $k(\mathbf{x}, \mathbf{x}')$ the covariance function, or kernel, defining how simi-
lar any pair of input points, \mathbf{x} and \mathbf{x}' , are. GPs can be updated
to posterior distributions over functions by conditioning on a
set of observed function outputs \mathbf{y} . Here the posterior mean
function is given by

$$801 \quad m_{\text{post}}(\mathbf{x}) = \mathbf{k}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}^T \quad (1)$$

802 where \mathbf{k} is the kernel matrix containing the covariance be-
803 tween training points and the evaluation points, \mathbf{K} is the
804 kernel matrix containing the covariance between all training
points, and σ^2 is a diagonal variance matrix.

The hypothesis that generalization is guided by a spatial
cognitive map corresponds to equipping a GP model with
a Gaussian (or Radial Basis Function) kernel, representing
similarity as an exponentially decaying function of squared
Euclidean distance. The Gaussian kernel defines similarity
as follows:

$$805 \quad k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\lambda^2}\right) \quad (2)$$

where σ_f^2 is a parameter controlling the degree to which the predictions differ from the mean, and λ is the lengthscale parameter, controlling how strongly input point similarity decays with distance. We obtained estimates of stimuli locations for every participant by performing path integration on their navigation runs.

To construct a kernel that corresponds to the hypothesis that temporal relations guided generalization, we started by computing a successor matrix \mathbf{M} for every participant³³. Each entry in the successor matrix $\mathbf{M}(s, s')$ (Equation 3) contains the expected discounted number of future visits of stimulus s' , starting from a visit to stimulus s

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

$$\hat{\mathbf{M}}(s, :) \leftarrow \hat{\mathbf{M}}(s, :) = \eta [\mathbf{1}_s + \gamma \hat{\mathbf{M}}(s', :) - \hat{\mathbf{M}}(s, :)] \quad (4)$$

where γ is the discount factor and \mathbb{I} is the indicator function. The successor matrix can be approximated from a participant's stimulus visitation history using a simple temporal-difference updating rule⁵⁴ (Equation 4), where $\hat{\mathbf{M}}(s, :)$ is the row corresponding to stimulus s , $\mathbf{1}_s$ is a vector of zeros except for the s th component which is a 1, and η is the learning rate. From \mathbf{M} we computed the transition matrix \mathbf{T} using the following equation (see Supplementary Note section for derivation):

$$\mathbf{T} = \frac{\mathbf{M}^{-1} - \mathbf{I}}{-\gamma} \quad (5)$$

where \mathbf{I} is the identity matrix. We enforced that \mathbf{T} was symmetric by taking the pairwise maximum of the entries of its upper and lower triangles. From \mathbf{T} , which describes the relevant participant's probabilities of walking directly from one stimulus to another, we computed the diffusion kernel³⁴ \mathbf{K} , embodying the hypothesis that temporal relations guide generalizations (Equation 6).

$$\mathbf{K} = \exp(-\lambda \mathbf{L}) \quad (6)$$

Here \exp is matrix exponentiation, \mathbf{L} is the normalized graph Laplacian which equals $\mathbf{I} - \mathbf{T}$, and λ is a lengthscale parameter analogous to that of the Gaussian kernel (Equation 2). To obtain the compositional kernel we took the average of the Gaussian and the diffusion kernel⁵⁵, and to implement the mean tracker we used a GP model whose kernel was the identity matrix \mathbf{I} .

To obtain the various GP models' estimates of stimuli's rewards at any given trial in the choice task, we conditioned them on all previously observed stimuli's rewards for the relevant context up to that point, and computed the posterior mean using Equation 1. The differences in estimated rewards were used as single predictors of participant choices in a logistic mixed-effects model with a participant-specific random slope³⁸, implemented in R using the `lme4`⁵⁶ package. We optimized hyper-parameters to minimize the log-likelihood of producing the choice data using a grid-search.

For the Gaussian kernel, we optimized the lengthscale λ , for the diffusion kernel we optimized the learning rate η , and set the discount rate parameter γ to 0.9 and the lengthscale λ to 1. For the compositional, spatio-temporal kernel, we optimized both the Gaussian kernel's lengthscale and the learning rate. The variance in Equation 1 was set 0.01 to improve numerical stability for matrix inversion. Using the best-fitting hyperparameter configurations, we performed a leave-one-out cross-validation (LOO-CV) procedure and obtained each model's cross-validated log-likelihood of producing every choice in the data set. We then computed the posterior model frequencies and exceedance probabilities⁵⁷, reported in Figure 3B.

We used the same procedure for modelling participants' value judgements. Here, we made the GP models predict the values of all stimuli, based on all reward-observations the participants had made, respectively. The GPs were equipped with the best-fitting hyper-parameters (see Supplementary Note section) from the choice task. We then sought to predict participants' value judgements for the different stimuli using the various value estimates as single predictors (plus an intercept) in separate linear mixed-effects models with a participant-specific random slope. We split the value judgements into two sets: One containing the value judgements of the inference objects, and another containing the value judgements of all monsters except the inference objects. Again, we performed LOO-CV to obtain model-specific log-likelihoods for all value judgements in the two data sets. Since the mean tracker could not generate predictions for the inference object any different from its prior mean function (which was 0), we used the average of the mean tracker's value predictions for the non-inference objects as a baseline model. From the cross-validated log-likelihoods we computed the corresponding sets of model frequencies and exceedance probabilities.

To compute the effects of the spatial and temporal components on each participant's choice behaviour, we fitted mixed-effects logistic regression models like the ones described above, using the estimated value differences generated by the spatial and temporal maps as individual predictors (using their respective best-fitting hyper-parameters) in the same model. Since the two predictors were correlated, we created two such models, one where the spatial value difference was the main predictor, and the second predictor was the temporal *minus* the spatial predictor, and a second model where this relation was inverted⁵⁸. We aggregated the unsigned mixed effects (random effects plus the fixed effects) across these two models for all participants, which left us with the effects for the two maps. To compute the spatial weights, we calculated how big the spatial effects were in proportion to the total effects (spatial + temporal effects). The temporal weights were consequently 1 minus the spatial weights. To compute the slopes, we first obtained a weight for the spatial map for all trials, and for all participants. We computed these weights by estimating two models similar to the ones used to estimate participant-specific effects, this time including

883 an interaction term with trial number as well. To obtain trial-
884 specific spatial weights for all participants, we estimated how
885 likely the spatial \times trial interaction predictor was at predict-
886 ing each individual choice compared to the temporal \times trial
887 interaction predictor, aggregating over our two models. We
888 then fitted logistic slopes to each participant's spatial weight
889 time series, predicting single participants' spatial weights
890 from trial number, using logistic regression.

891 **fMRI data acquisition and pre-processing**

892 Visual stimuli were projected onto a screen via a computer
893 monitor. Participants indicated their choice using an MRI-
894 compatible button box.

895 MRI data were acquired using a 32-channel head coil
896 on a 3 Tesla Siemens Magnetom SkyraFit system (Siemens,
897 Erlangen, Germany). fMRI scans were acquired in axial
898 orientation using T2*-weighted gradient-echo echo planar
899 imaging (GE-EPI) with multiband acceleration, sensitive to
900 blood oxygen level-dependent (BOLD) contrast^{59,60}. Echo-
901 planar imaging (EPI) with sampling after multiband excita-
902 tion achieves temporal resolution in the sub-second regime
903 whilst maintaining a good slice coverage and spatial resolu-
904 tion^{59,60}. We collected 60 transverse slices of 2-mm thick-
905 ness with an in-plane resolution of 2×2 mm, a multiband
906 acceleration factor of 3, a repetition time of 2 s, and an echo
907 time of 23.6 ms. Slices were tilted by 90 deg relative to the
908 rostro-caudal axis. The first five volumes of each block were
909 discarded to allow for scanner equilibration. Furthermore,
910 a T1-weighted anatomical scan with $1 \times 1 \times 1$ mm reso-
911 lution was acquired. In addition, a whole-brain field map
912 with dual echo-time images (TE1 = 5.92 ms, TE2 = 8.38
913 ms, resolution $2 \times 2 \times 2.26$ mm) was obtained in order to
914 measure and later correct for geometric distortions due to
915 susceptibility-induced field inhomogeneities.

916 **Anatomical data preprocessing**

917 Results included in this manuscript come from preprocessing
918 performed using *fMRIPrep* 1.4.0^{61,62} (RRID:SCR_016216),
919 which is based on *Nipype* 1.2.0^{63,64} (RRID:SCR_002502).

920 A total of 2 T1-weighted (T1w) images were
921 found within the input BIDS dataset. All of
922 them were corrected for intensity non-uniformity
923 (INU) with *N4BiasFieldCorrection*⁶⁵, distributed
924 with ANTs 2.2.0⁶⁶. The T1w-reference was then
925 skull-stripped with a *Nipype* implementation of the
926 *antsBrainExtraction.sh* workflow (from ANTs),
927 using OASIS30ANTs as target template. Brain tissue seg-
928 mentation of cerebrospinal fluid (CSF), white-matter (WM)
929 and gray-matter (GM) was performed on the brain-extracted
930 T1w using *fast*⁶⁷. A T1w-reference map was computed
931 after registration of 2 T1w images (after INU-correction)
932 using *mri_robust_template*⁶⁸.

933 Brain surfaces were reconstructed using *recon-all*⁶⁹,
934 and the brain mask estimated previously was refined
935 with a custom variation of the method to reconcile
936 ANTs-derived and FreeSurfer-derived segmentations of

937 the cortical gray-matter of Mindboggle⁷⁰. Volume-
938 based spatial normalization to one standard space
939 (MNI152NLin6Asym) was performed through nonlinear
940 registration with *antsRegistration* (ANTs 2.2.0), using
941 brain-extracted versions of both T1w reference and the
942 T1w template. The following template was selected for
943 spatial normalization: *FSL's MNI ICBM 152 non-linear*
944 *6th Generation Asymmetric Average Brain Stereotaxic Reg-*
945 *istration Model*⁷¹ [RRID:SCR_002823; TemplateFlow ID:
946 MNI152NLin6Asym].

947 **Functional data preprocessing**

948 For each of the 7 BOLD runs per subject (across all tasks
949 and sessions), the following preprocessing was performed.
950 First, a reference volume and its skull-stripped version were
951 generated using a custom methodology of *fMRIPrep*. A
952 deformation field to correct for susceptibility distortions
953 was estimated based on a field map that was co-registered
954 to the BOLD reference, using a custom workflow of *fM-*
955 *RIPrep* derived from D. Greve's *epidewarp.fsl* script
956 and further improvements of HCP Pipelines⁷². Based on
957 the estimated susceptibility distortion, an unwarped BOLD
958 reference was calculated for a more accurate co-registration
959 with the anatomical reference. The BOLD reference was
960 then co-registered to the T1w reference using *bbregister*
961 (FreeSurfer) which implements boundary-based registra-
962 tion⁷³. Co-registration was configured with nine degrees
963 of freedom to account for distortions remaining in the BOLD
964 reference. Head-motion parameters with respect to the
965 BOLD reference (transformation matrices, and six corre-
966 sponding rotation and translation parameters) are estimated
967 before any spatiotemporal filtering using *mcflirt*⁷⁴.

968 BOLD runs were slice-time corrected using *3dTshift*
969 from AFNI 20190105⁷⁵. The BOLD time-series (including
970 slice-timing correction when applied) were resampled onto
971 their original, native space by applying a single, compos-
972 ite transform to correct for head-motion and susceptibility
973 distortions. These resampled BOLD time-series will be re-
974 ferred to as *preprocessed BOLD in original space*, or just
975 *preprocessed BOLD*. The BOLD time-series were resampled
976 into standard space, generating a *preprocessed BOLD run*
977 *in ['MNI152NLin6Asym'] space*. First, a reference volume
978 and its skull-stripped version were generated using a custom
979 methodology of *fMRIPrep*.

980 Additionally, several confounding time-series were calcu-
981 lated based on the *preprocessed BOLD*: framewise displace-
982 ment (FD), DVARS and three region-wise global signals.
983 FD and DVARS are calculated for each functional run, both
984 using their implementations in *Nipype*⁷⁶. The three global
985 signals are extracted within the CSF, the WM, and the whole-
986 brain masks. Additionally, a set of physiological regressors
987 were extracted to allow for component-based noise correc-
988 tion *CompCor*⁷⁷. Principal components are estimated after
989 high-pass filtering the *preprocessed BOLD* time-series (us-
990 ing a discrete cosine filter with 128s cut-off) for the two
991 *CompCor* variants: temporal (tCompCor) and anatomical

992 (aCompCor). tCompCor components are then calculated 1046
993 from the top 5% variable voxels within a mask covering the 1047
994 subcortical regions. This subcortical mask is obtained by 1048
995 heavily eroding the brain mask, which ensures it does not 1049
996 include cortical GM regions. For aCompCor, components 1050
997 are calculated within the intersection of the aforementioned 1051
998 mask and the union of CSF and WM masks calculated in 1052
999 T1w space, after their projection to the native space of each 1053
1000 functional run (using the inverse BOLD-to-T1w transforma- 1054
1001 tion). Components are also calculated separately within the 1055
1002 WM and CSF masks. For each CompCor decomposition, the 1056
1003 k components with the largest singular values are retained, 1057
1004 such that the retained components' time series are sufficient 1058
1005 to explain 50 percent of variance across the nuisance mask 1059
1006 (CSF, WM, combined, or temporal). The remaining com- 1060
1007 ponents are dropped from consideration. The head-motion 1061
1008 estimates calculated in the correction step were also placed 1062
1009 within the corresponding confounds file. The confound time 1063
1010 series derived from head motion estimates and global signals 1064
1011 were expanded with the inclusion of temporal derivatives 1065
1012 and quadratic terms for each⁷⁸. 1066

1013 Frames that exceeded a threshold of 0.5 mm FD or 1.5 1067
1014 standardised DVARS were annotated as motion outliers. All 1068
1015 resamplings can be performed with a *single interpolation* 1069
1016 *step* by composing all the pertinent transformations (i.e. head- 1070
1017 motion transform matrices, susceptibility distortion correc- 1071
1018 tion when available, and co-registrations to anatomical and 1072
1019 output spaces). Gridded (volumetric) resamplings were per- 1073
1020 formed using `antsApplyTransforms` (ANTs), configu- 1074
1021 red with Lanczos interpolation to minimize the smoothing 1075
1022 effects of other kernels⁷⁹. Non-gridded (surface) resamplings 1076
1023 were performed using `mri_vol2surf` (FreeSurfer). 1077

1024 **fMRI data analysis**

1025 We implemented three types of event-related general linear 1080
1026 models (GLMs) in SPM 12 to analyze the fMRI data. All 1081
1027 GLMs included a button press regressor as a regressor of 1082
1028 no interest. All regressors were convolved with a canonical 1083
1029 haemodynamic response function. Because of the sensitiv- 1084
1030 ity of the blood oxygen level-dependent signal to motion 1085
1031 and physiological noise, all GLMs included frame-wise dis- 1086
1032 placement, six rigid-body motion parameters (three transla- 1087
1033 tions and three rotation), six anatomical component-based 1088
1034 noise correction components (aCompCorr) and four cosine 1089
1035 regressors estimated by `fmripred` as confound regressors for 1090
1036 denoising. Each block was modeled separately within the 1091
1037 GLMs. 1092

1038 The first GLM contained separate onset regressors for 1093
1039 each of the twelve objects. By modeling each object sepa- 1094
1040 rately, we could account for any object-specific differences 1095
1041 in activity driving the main effects and focus on distance- 1096
1042 dependent modulations that ride on top of those object- 1097
1043 specific differences in activation. Each onset regressor was 1098
1044 accompanied by two parametric regressors. These corre- 1099
1045 sponded to the distance to the object presented immediately 1100

before the current object according to the spatial kernel and
distance to the immediately preceding object according to the
temporal kernel. Both parametric regressors were zscored,
but not orthogonalized, so that any shared variance would
be discarded. Trials where the same object was repeated
were modeled separately and objects immediately following
a choice were excluded. Furthermore, the GLM contained an
onset regressor for the choice trials. This was accompanied
by two parametric regressors, reflecting chosen and an un-
chosen distance between the two objects and the preceding
object. Each of the three blocks were modeled separately.

The second and third GLM modeled events during the
choice task. Here, three onset regressors were included,
one indicating the choice period, the second one indicating
feedback times and the third one corresponding to button
presses. The duration of each event corresponded to the
actual duration during the experiment. The choice period
regressor was accompanied by two parametric modulators
reflecting chosen and unchosen values of the objects as esti-
mated by the winning model. Both were demeaned, but not
orthogonalized.

In the second GLM instead, the feedback regressor was
accompanied by a spatial weight updating signal. A trial-
by-trial estimate of the influence of the spatial map on the
choices was estimated, and the demeaned trial-by-trial dif-
ference was included as a parametric modulator.

In the third GLM, the feedback regressor was accompa-
nied by a parametric regressor reflecting a prediction error
difference signal. The reward prediction error was estimated
separately for the spatial and the temporal map, and the de-
meaned difference between the absolute prediction errors
was included as a parametric regressor.

The contrast images of all participants from the first level
were analysed as a second-level random effects analysis.
We report all our results in the hippocampal formation, as
this was our a priori ROI, at an uncorrected cluster-defining
threshold of $p < 0.001$, combined with peak-level family-
wise error (FWE) small-volume correction at $p < 0.05$. For
the SVC procedure, we used a mask comprising hippocampus,
entorhinal cortex and subiculum (Supplementary Figure
S5). Activations in other brain regions were only consid-
ered significant at a level of $p < 0.001$ uncorrected if they
survived whole-brain FWE correction at the cluster level
($p < 0.05$). Results in the orbitofrontal cortex in 5h are
reported at a cluster-defining threshold of $p < 0.01$ uncor-
rected, combined with a whole-brain FWE-corrected signifi-
cance at the cluster level of $p < 0.05$. While we used masks
to correct for multiple comparisons in our ROI, all statistical
parametric maps presented in the manuscript are unmasked
and thresholded at $p < 0.01$ for visualization.

To relate neural effects to behavioral parameters and to
each other, we defined the following ROIs: spatial hippocam-
pal map in session 3 from GLM 1, Figure 4a; hippocampal
spatial weight update from GLM 2, Figure 5f; change in
hippocampal map representation from session 2 to session

1101 3 with hippocampal spatial weight update as covariate from
1102 GLM 1, Figure 5g; and OFC evidence integration signal
1103 with hippocampal spatial weight update as covariate from
1104 GLM 3 5h. All voxels exceeding a threshold of $p < 0.001$
1105 were included in an ROI if the cluster survived correction for
1106 multiple comparisons.

1107 To estimate how much an effect co-varied with behav-
1108 ioral effects, we included spatial and temporal weights, re-
1109 spectively (Figure 4f), as well as the inference error (Figure
1110 4g) as a covariate on the second level and tested for signifi-
1111 cant effects. In Figure 5g and h, we included the parameter
1112 estimate reflecting the size of the hippocampal spatial weight
1113 update signal (Figure 5f) as a covariate.

1114 Mediation analysis

1115 We used the Mediation and Moderation Toolbox^{42,43} to per-
1116 form two single-level mediation analyses (Figures 4h and
1117 5i). The total effect of the independent variable X on the
1118 dependent variable Y is referred to as path c. That effect is
1119 then partitioned into a combination of a direct effect of X on
1120 Y (path c'), and an indirect effect of X on Y that is transmit-
1121 ted through a mediator M (path ab). We also estimated the
1122 relationship between X and M (path a) as well as between M
1123 and Y (path b). This last path "b" is controlled for X, such
1124 that paths "a" and "b" correspond to two separable processes
1125 contributing to Y. We determined two-tailed uncorrected p
1126 values from the bootstrap confidence intervals for the path
1127 coefficients⁴³.

1128 To test whether the spatial weights mediate the effect of
1129 hippocampal spatial map on the inference error, we defined X
1130 as each individual's parameter estimate from the hippocam-
1131 pal ROI encoding the spatial map (ROI based on Figure 4a).
1132 The mediator M corresponded to each participant's spatial
1133 weight as estimated by the model fit to the choice data. The
1134 outcome variable Y was defined as a participant's inference
1135 error.

1136 To test for a significant mediation linking the OFC evi-
1137 dence integration signals (X) to the change in hippocampal
1138 spatial map (Y), we extracted parameter estimates from an
1139 orbitofrontal ROI tracking the evidence that an outcome is
1140 predicted by either of the two maps (X, ROI based on Figure
1141 5h) and related this to the change in spatial representation
1142 in the left hippocampus (Y, ROI based on Figure 5g) via the
1143 spatial updating signal in the right hippocampus (M, ROI
1144 based on Figure 5f).

1145 Data availability

1146 Source data to reproduce the figures and unthresholded
1147 group-level statistical brain maps from neuroimaging analy-
1148 ses will be made openly available upon publication.

1149 Code availability

1150 Task, analysis and computational modeling code will be
1151 made publicly available on github upon publication.

References

1. Shepard, R. N. Toward a universal law of generaliza-
tion for psychological science. *Science* **237**, 1317–1323
(1987).
2. Gershman, S. J. & Daw, N. D. Reinforcement learning
and episodic memory in humans and animals: An inte-
grative framework. *Annu. Rev. Psychol.* **68**, 101–128,
DOI: 10.1146/annurev-psych-122414-033625 (2017).
3. Guttman, N. & Kalish, H. I. Discriminability and stimu-
lus generalization. *J. experimental psychology* **51**, 79
(1956).
4. Hanson, H. M. Effects of discrimination training on
stimulus generalization. *J. experimental psychology* **58**,
321 (1959).
5. Kahnt, T. & Tobler, P. N. Dopamine regulates stimu-
lus generalization in the human hippocampus. *Elife* **5**,
e12678 (2016).
6. Wu, C. M., Schulz, E., Garvert, M. M., Meder, B. &
Schuck, N. W. Similarities and differences in spatial
and non-spatial cognitive maps. *PLoS computational
biology* **16**, e1008149 (2020).
7. Barron, H. C. *et al.* Neuronal computation underlying
inferential reasoning in humans and mice. *Cell* **183**,
228–243 (2020).
8. Brogden, W. J. Sensory pre-conditioning. *J. Exp. Psy-
chol.* **25**, 323 (1939).
9. Baram, A. B., Muller, T. H., Nili, H., Garvert, M. M.
& Behrens, T. E. J. Entorhinal and ventromedial pre-
frontal cortices abstract and generalize the structure of
reinforcement learning problems. *Neuron* **109**, 713–723
(2021).
10. Wimmer, G. E., Daw, N. D. & Shohamy, D. Generaliza-
tion of value in reinforcement learning by humans. *Eur.
J. Neurosci.* **35**, 1092–1104 (2012).
11. Morgan, L. K., MacEvoy, S. P., Aguirre, G. K. & Ep-
stein, R. A. Distances between real-world locations are
represented in the human hippocampus. *J. Neurosci.* **31**,
1238–1245, DOI: 10.1523/JNEUROSCI.4667-10.2011
(2011). [https://www.jneurosci.org/content/31/4/1238.
full.pdf](https://www.jneurosci.org/content/31/4/1238.full.pdf).
12. O'keefe, J. & Nadel, L. *The hippocampus as a cognitive
map* (Oxford: Clarendon Press, 1978).
13. Tolman, E. C. Cognitive maps in rats and men. *Psychol.
review* **55**, 189 (1948).
14. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E.
Organizing conceptual knowledge in humans with a
gridlike code. *Science* **352**, 1464–1468 (2016).
15. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-
spatial dimension by the hippocampal–entorhinal cir-
cuit. *Nature* **543**, 719–722, DOI: 10.1038/nature21692
(2017).

- 1203 **16.** Nau, M., Navarro Schröder, T., Bellmund, J. L. S. & 1255
 1204 Doeller, C. F. Hexadirectional coding of visual space in 1256
 1205 human entorhinal cortex. *Nat. Neurosci.* **21**, 188–190, 1257
 1206 DOI: 10.1038/s41593-017-0050-8 (2018). 1258
- 1207 **17.** Theves, S., Fernández, G. & Doeller, C. F. The hip- 1259
 1208 pocus maps concept space, not feature space. *J.* 1260
 1209 *Neurosci.* **40**, 7318–7325, DOI: 10.1523/JNEUROSCI. 1261
 1210 0494-20.2020 (2020). [https://www.jneurosci.org/](https://www.jneurosci.org/content/40/38/7318.full.pdf)
 1211 [content/40/38/7318.full.pdf](https://www.jneurosci.org/content/40/38/7318.full.pdf). 1262
- 1212 **18.** Theves, S., Fernandez, G. & Doeller, C. F. The hip- 1264
 1213 pocus encodes distances in multidimensional fea- 1265
 1214 ture space. *Curr. Biol.* **29**, 1226–1231.e3, DOI: [https://](https://doi.org/10.1016/j.cub.2019.02.035)
 1215 doi.org/10.1016/j.cub.2019.02.035 (2019). 1266
- 1216 **19.** Vigano, S. & Piazza, M. Distance and direction 1268
 1217 codes underlie navigation of a novel semantic space 1269
 1218 in the human brain. *J. Neurosci.* **40**, 2727–2736, 1270
 1219 DOI: 10.1523/JNEUROSCI.1849-19.2020 (2020). [https://](https://www.jneurosci.org/content/40/13/2727.full.pdf)
 1220 www.jneurosci.org/content/40/13/2727.full.pdf. 1271
- 1221 **20.** Deuker, L., Bellmund, J., Schröder, T. N. & Doeller, 1272
 1222 C. An event map of memory space in the hippocampus. 1273
 1223 *eLife* **5**, DOI: 10.7554/eLife.16534 (2016). 1274
- 1224 **21.** Bellmund, J. L. S., Polti, I. & Doeller, C. F. Sequence 1275
 1225 memory in the hippocampal–entorhinal region. *J. Cogn.* 1276
 1226 *Neurosci.* **32**, 2056–2070, DOI: 10.1162/jocn_a_01592 1277
 1227 (2020). 1278
- 1228 **22.** Eichenbaum, H. Time cells in the hippocampus: a new 1279
 1229 dimension for mapping memories. *Nat. Rev. Neurosci.* 1280
 1230 **15**, 732–744, DOI: 10.1038/nrn3827 (2014). 1281
- 1231 **23.** Garvert, M. M., Dolan, R. J. & Behrens, T. E. A 1282
 1232 map of abstract relational knowledge in the human 1283
 1233 hippocampal–entorhinal cortex. *Elife* **6**, e17086 (2017). 1284
- 1234 **24.** Schuck, N., Cai, M., Wilson, R. & Niv, Y. Human Or- 1286
 1235 bitofrontal Cortex Represents a Cognitive Map of State 1287
 1236 Space. *Neuron* **91**, 1402–1412, DOI: 10.1016/j.neuron. 1288
 1237 2016.08.019 (2016). 1289
- 1238 **25.** Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk- 1290
 1239 Browne, N. B. & Botvinick, M. M. Neural represen- 1291
 1240 tations of events arise from temporal community struc- 1292
 1241 ture. *Nat. Neurosci.* **16**, 486–492, DOI: 10.1038/nn.3331 1293
 1242 (2013). 1294
- 1243 **26.** Schapiro, A., Kustner, L. & Turk-Browne, N. Shaping 1295
 1244 of object representations in the human medial tempo- 1296
 1245 ral lobe based on temporal regularities. *Curr. Biol.* **22**, 1297
 1246 1622–1627, DOI: 10.1016/j.cub.2012.06.056 (2011). 1298
- 1247 **27.** Nieh, E. H. *et al.* Geometry of abstract learned knowl- 1298
 1248 edge in the hippocampus. *Nature* **595**, 80–84, DOI: 1299
 1249 10.1038/s41586-021-03652-7 (2021). 1300
- 1250 **28.** Shahar, N. *et al.* Credit assignment to state-independent 1301
 1251 task representations and its relationship with model- 1302
 1252 based decision making. *Proc. Natl. Acad. Sci.* **116**, 1303
 1253 15871–15876, DOI: 10.1073/pnas.1821647116 (2019). 1304
 1254 <https://www.pnas.org/content/116/32/15871.full.pdf>. 1305
- 29.** Niv, Y. Learning task-state representa-
 tions. *Nat. Neurosci.* **22**, 1544–1553, DOI:
 10.1038/s41593-019-0470-8 (2019).
- 30.** Wikenheiser, A. M. & Schoenbaum, G. Over the river,
 through the woods: cognitive maps in the hippocam-
 pus and orbitofrontal cortex. *Nat. Rev. Neurosci.* **17**,
 513–523, DOI: 10.1038/nrn.2016.56 (2016).
- 31.** Schuck, N. W. & Niv, Y. Sequential replay of nonspatial
 task states in the human hippocampus. *Science* **364**,
 eaaw5181, DOI: 10.1126/science.aaw5181 (2019).
- 32.** Wittkuhn, L., Chien, S., Hall-McMaster, S. & Schuck,
 N. W. Replay in minds and machines. *Neurosci. &*
Biobehav. Rev. **129**, 367–388, DOI: [https://doi.org/10.](https://doi.org/10.1016/j.neubiorev.2021.08.002)
 1016/j.neubiorev.2021.08.002 (2021).
- 33.** Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J.
 The hippocampus as a predictive map. *Nat. neuroscience*
20, 1643 (2017).
- 34.** Kondor, R. I. & Lafferty, J. Diffusion kernels on graphs
 and other discrete structures (2002).
- 35.** Duvenaud, D., Lloyd, J., Grosse, R., Tenenbaum, J. &
 Zoubin, G. Structure discovery in nonparametric res-
 gression through compositional kernel search. In *Inter-*
national Conference on Machine Learning, 1166–1174
 (PMLR, 2013).
- 36.** Saanum, T., Schulz, E. & Speekenbrink, M. Composi-
 tional generalization in multi-armed bandits. *PsyArXiv*
 (2021).
- 37.** Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speeken-
 brink, M. & Gershman, S. J. Compositional inductive
 biases in function learning. *Cogn. psychology* **99**, 44–79
 (2017).
- 38.** Gershman, S. J. Uncertainty and exploration. *Decision*
6, 277 (2019).
- 39.** Rigoux, L., Stephan, K. E., Friston, K. J. & Daunizeau,
 J. Bayesian model selection for group studies—revisited.
Neuroimage **84**, 971–985 (2014).
- 40.** Barron, H. C., Garvert, M. M. & Behrens, T. E. Rep-
 etition suppression: a means to index neural represen-
 tations using bold? *Philos. Transactions Royal Soc. B:*
Biol. Sci. **371**, 20150355 (2016).
- 41.** Grill-Spector, K. Selectivity of adaptation in single units:
 implications for fmri experiments. *Neuron* **49**, 170–171
 (2006).
- 42.** Wager, T. D., Davidson, M. L., Hughes, B. L., Lindquist,
 M. A. & Ochsner, K. N. Prefrontal-subcortical pathways
 mediating successful emotion regulation. *Neuron* **59**,
 1037–1050, DOI: [https://doi.org/10.1016/j.neuron.2008.](https://doi.org/10.1016/j.neuron.2008.09.006)
 09.006 (2008).
- 43.** Atlas, L. Y., Bolger, N., Lindquist, M. A. & Wager,
 T. D. Brain mediators of predictive cue effects on
 perceived pain. *J. Neurosci.* **30**, 12964–12977, DOI:

- 10.1523/JNEUROSCI.0057-10.2010 (2010). <https://www.jneurosci.org/content/30/39/12964.full.pdf>.
44. Schoenbaum, G., Roesch, M. R., Stalnaker, T. A. & Takahashi, Y. K. A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nat. Rev. Neurosci.* **10**, 885–892, DOI: 10.1038/nrn2753 (2009).
45. Howard, L. R. *et al.* The hippocampus and entorhinal cortex encode the path and euclidean distances to goals during navigation. *Curr. Biol.* **24**, 1331–1340, DOI: 10.1016/j.cub.2014.05.001 (2014).
46. Chadwick, M., Jolly, A., Amos, D., Hassabis, D. & Spiers, H. A goal direction signal in the human entorhinal/subicular region. *Curr. Biol.* **25**, 87–92, DOI: <https://doi.org/10.1016/j.cub.2014.11.001> (2015).
47. Segaert, K., Weber, K., de Lange, F. P., Petersson, K. M. & Hagoort, P. The suppression of repetition enhancement: A review of fmri studies. *Neuropsychologia* **51**, 59–66, DOI: <https://doi.org/10.1016/j.neuropsychologia.2012.11.006> (2013).
48. Schuck, N. W., Wilson, R. & Niv, Y. Chapter 12 - a state representation for reinforcement learning and decision-making in the orbitofrontal cortex. In Morris, R., Bornstein, A. & Shenhav, A. (eds.) *Goal-Directed Decision Making*, 259–278, DOI: <https://doi.org/10.1016/B978-0-12-812098-9.00012-7> (Academic Press, 2018).
49. Doeller, C. F., King, J. A. & Burgess, N. Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proc. Natl. Acad. Sci.* **105**, 5915–5920, DOI: 10.1073/pnas.0801489105 (2008). <https://www.pnas.org/content/105/15/5915.full.pdf>.
50. Gallagher, M., McMahan, R. W. & Schoenbaum, G. Orbitofrontal cortex and representation of incentive value in associative learning. *J. Neurosci.* **19**, 6610–6614, DOI: 10.1523/JNEUROSCI.19-15-06610.1999 (1999). <https://www.jneurosci.org/content/19/15/6610.full.pdf>.
51. Wikenheiser, A. M., Marrero-Garcia, Y. & Schoenbaum, G. Suppression of ventral hippocampal output impairs integrated orbitofrontal encoding of task structure. *Neuron* **95**, 1197–1207.e3, DOI: <https://doi.org/10.1016/j.neuron.2017.08.003> (2017).
52. Boorman, E., Rajendran, V., O’Reilly, J. & Behrens, T. Two anatomically and computationally distinct learning signals predict changes to stimulus-outcome associations in hippocampus. *Neuron* **89**, 1343–1354, DOI: <https://doi.org/10.1016/j.neuron.2016.02.014> (2016).
53. Zhou, J. *et al.* Evolving schema representations in orbitofrontal ensembles during learning. *Nature* **590**, 606–611, DOI: 10.1038/s41586-020-03061-2 (2021).
54. Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J. & Daw, N. D. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS computational biology* **13**, e1005768 (2017).
55. Schulz, E., Franklin, N. T. & Gershman, S. J. Finding structure in multi-armed bandits. *Cogn. psychology* **119**, 101261 (2020).
56. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *J. Stat. Software, Articles* **67**, 1–48, DOI: 10.18637/jss.v067.i01 (2015).
57. Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J. & Friston, K. J. Bayesian model selection for group studies. *Neuroimage* **46**, 1004–1017 (2009).
58. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans’ choices and striatal prediction errors. *Neuron* **69**, 1204–1215 (2011).
59. Feinberg, D. *et al.* Multiplexed echo planar imaging for sub-second whole brain fmri and fast diffusion imaging. *PLoS one* **5**, e15710 (2010).
60. Moeller, S. *et al.* Multiband multislice ge-epi at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fmri. *Magn Reson. Med* **63**, 1144–1153 (2010).
61. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* DOI: 10.1038/s41592-018-0235-4 (2018).
62. Esteban, O. *et al.* fmriprep. *Software* DOI: 10.5281/zenodo.852659 (2018).
63. Gorgolewski, K. *et al.* Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front. Neuroinformatics* **5**, 13, DOI: 10.3389/fninf.2011.00013 (2011).
64. Gorgolewski, K. J. *et al.* Nipype. *Software* DOI: 10.5281/zenodo.596855 (2018).
65. Tustison, N. J. *et al.* N4itk: Improved n3 bias correction. *IEEE Transactions on Med. Imaging* **29**, 1310–1320, DOI: 10.1109/TMI.2010.2046908 (2010).
66. Avants, B., Epstein, C., Grossman, M. & Gee, J. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Analysis* **12**, 26–41, DOI: 10.1016/j.media.2007.06.004 (2008).
67. Zhang, Y., Brady, M. & Smith, S. Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Med. Imaging* **20**, 45–57, DOI: 10.1109/42.906424 (2001).
68. Reuter, M., Rosas, H. D. & Fischl, B. Highly accurate inverse consistent registration: A robust approach. *NeuroImage* **53**, 1181–1196, DOI: 10.1016/j.neuroimage.2010.07.020 (2010).

- 1408 **69.** Dale, A. M., Fischl, B. & Sereno, M. I. Cortical surface- 1446
 1409 based analysis: I. segmentation and surface reconstruc- 1447
 1410 tion. *NeuroImage* **9**, 179–194, DOI: 10.1006/nimg.1998. 1448
 1411 0395 (1999). 1449
- 1412 **70.** Klein, A. *et al.* Mindboggling morphometry of hu- 1450
 1413 man brains. *PLOS Comput. Biol.* **13**, e1005350, DOI: 1451
 1414 10.1371/journal.pcbi.1005350 (2017). 1452
- 1415 **71.** Evans, A., Janke, A., Collins, D. & Baillet, S. Brain 1453
 1416 templates and atlases. *NeuroImage* **62**, 911–922, DOI: 1454
 1417 10.1016/j.neuroimage.2012.01.024 (2012). 1455
- 1418 **72.** Glasser, M. F. *et al.* The minimal preprocessing 1456
 1419 pipelines for the human connectome project. *NeuroIm-* 1457
 1420 *age* **80**, 105–124, DOI: 10.1016/j.neuroimage.2013.04. 1458
 1421 127 (2013). 1459
- 1422 **73.** Greve, D. N. & Fischl, B. Accurate and robust brain im- 1460
 1423 age alignment using boundary-based registration. *Neu-* 1461
 1424 *roImage* **48**, 63–72, DOI: 10.1016/j.neuroimage.2009. 1462
 1425 06.060 (2009). 1463
- 1426 **74.** Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Im- 1464
 1427 proved optimization for the robust and accurate linear 1465
 1428 registration and motion correction of brain images. *Neu-* 1466
 1429 *roImage* **17**, 825–841, DOI: 10.1006/nimg.2002.1132 1467
 1430 (2002). 1468
- 1431 **75.** Cox, R. W. & Hyde, J. S. Software tools for analysis 1469
 1432 and visualization of fmri data. *NMR Biomed.* **10**, 171– 1470
 1433 178, DOI: 10.1002/(SICI)1099-1492(199706/08)10:4/ 1471
 1434 5<171::AID-NBM453>3.0.CO;2-L (1997). 1472
- 1435 **76.** Power, J. D. *et al.* Methods to detect, characterize, and 1473
 1436 remove motion artifact in resting state fmri. *NeuroImage* 1474
 1437 **84**, 320–341, DOI: 10.1016/j.neuroimage.2013.08.048 1475
 1438 (2014). 1476
- 1439 **77.** Behzadi, Y., Restom, K., Liau, J. & Liu, T. T. A com- 1477
 1440 ponent based noise correction method (CompCor) for 1478
 1441 BOLD and perfusion based fmri. *NeuroImage* **37**, 90– 1479
 1442 101, DOI: 10.1016/j.neuroimage.2007.04.042 (2007). 1480
- 1443 **78.** Satterthwaite, T. D. *et al.* An improved framework 1481
 1444 for confound regression and filtering for control of mo- 1482
 1445 tion artifact in the preprocessing of resting-state func- 1483
 tional connectivity data. *NeuroImage* **64**, 240–256, DOI: 1484
 10.1016/j.neuroimage.2012.08.052 (2013). 1485
- 79.** Lanczos, C. Evaluation of noisy data. *J. Soc. for Ind. 1486
 Appl. Math. Ser. B Numer. Analysis* **1**, 76–85, DOI: 1487
 10.1137/0701007 (1964). 1488

Acknowledgements


We would like to thank Jacob Bellmund for helpful comments on the manuscript, Joshua Julian for providing example code for the VR experiment and Kerstin Träger and Nicolas Filler for help with data collection. We thank the University of Minnesota Center for Magnetic Resonance Research for the provision of the multiband EPI sequence software.

This work is supported by the Max Planck Society. TS and ES are supported via the Mini Graduate School on “Compositionality in Minds and Machines” from the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy–EXC2064/1–390727645. ES is further supported by an Independent Max Planck Research Group grant awarded by the Max Planck Society. NWS is supported by an Independent Max Planck Research Group grant awarded by the Max Planck Society and a Starting Grant from the European Union (ERC-2019-StG REPLAY-852669). CFD is supported by the Max Planck Society, the European Research Council (ERC-CoG GEOCOG 724836), the Kavli Foundation, the Jebsen Foundation, the Centre of Excellence scheme of the Research Council of Norway – Centre for Neural Computation (223262/F50), The Egil and Pauline Braathen and Fred Kavli Centre for Cortical Microcircuits, and the National Infrastructure scheme of the Research Council of Norway – NORBRAIN (197467/F50).


Author contributions statement

M.M.G., N.W.S. and C.F.D. conceived the experiment, M.M.G. developed the tasks and acquired the data, all authors planned the analyses, M.M.G. and T.S. analyzed the data, T.S. and E.S. performed the computational modeling, all authors discussed the results, M.M.G. and T.S. wrote the manuscript with input from all authors.


1 Supplementary Information




Supplementary Figure S1. Exploration paths on day 1 in each individual. Each panel represents the exploration trajectories concatenated across exploration blocks on day 1 in one participant. Purple indicates the stimulus locations and black the participant's trajectory.




Supplementary Figure S2. Object positioning after learning. **a** Each panel displays the data for one object. Yellow indicates the true object position. Black indicates the drop location for each participant. The replacement error is defined as the Euclidean distance between the true location and the drop location. Visualized is the data from the last object location memory task block on day 1, i.e. at the end of learning. **b** Linear regression of values on replacement error. On day 2 as well as on day 3 before the choice task, there was no relationship between values participants learned to associate with each object and replacement error (all p values > 0.05). This is not surprising, since participants only learned the value associations on day 3. On day 3 after the choice task, the replacement error was smaller the higher the reported value of an object ($t(47) = -2.9, p = 0.005$). The difference between value-dependent performance pre and post choice was also significant on day 3 ($t(47) = 2.26, p = 0.03$), but not on day 2 ($t(47) = 0.27, p = 0.79$). This suggests that participants' memory expression was more accurate around valuable objects compared to less valuable ones after participants learned to associated objects with values. We used the average values that participants reported at the end of the study on day 3 as predictors. For inference objects, only the value experienced in the other context was considered.




Supplementary Figure S3. Full modeling results. **a** Model frequencies for predicting participants' value ratings for the experienced objects at the end of the study. The winning model does not generalize about value. **b** Model AIC differences for the choice task. **c** Models' McFadden's R^2 for the choice task. This statistic quantifies how likely a model is to produce the data relative to a random model, where a score of 1 means that the model is infinitely more likely to produce the data, and a score of 0 means the model is as likely as the random model. The dashed line represents the score of a model that uses the true value difference between options as a predictor. This model was only tested on trials where participants had observed the value of both options, and whose score therefore approximates an upper bound on how accurately one can predict participant choices assuming one has access to their beliefs about value and perfect memory, relative to a chance levels. **d** Model AIC differences for predicting participants' value ratings for the inference objects.



Supplementary Figure S4. Model recovery. **a** The pairwise correlation between all participants' temporal kernels estimated with a learning rate of 0.4125, which gave the best fit for the temporal model. We flattened each participant's 12×12 temporal kernel matrix into a vector and computed Pearson's correlation coefficient r between all pairs of vectors. **b** Predicting reward generalization using participants' own temporal kernel yields substantially better fits to their choice behaviour (blue line) than predicting generalization using another randomly picked participant's temporal kernel (red line). Error bars are standard deviations of negative log-likelihood of 10 sampled random assignments. See Supplementary Note section for procedure. **c** To verify that the predictive performance of the spatio-temporal model was not an artifact of the kernel composition procedure per se, we compared the spatio-temporal model against a model using a composition of a spatial kernel and the identity matrix. The spatio-temporal kernel produced a substantially better fit to participant choices, indicating that both the spatio-temporal model's components captures something important about how participants generalized. **d** We performed a model recovery analysis for our computational models, using their own best-fitting hyper-parameters. We first simulated choice behaviour from our models based on the choices that the participants encountered in the experiment. As such, we obtained 4800 simulated decisions from our models. The temporal and spatio-temporal models used the temporal kernel of the participant at the corresponding trial. Choices were made deterministically to maximize expected reward, where the expected reward was estimated from previous observations. After each choice, the models received a reward which they used to condition predictions about rewards for subsequent trials. We then computed how likely each model was to produce the simulated choice behaviour from all other models, including its own choice behaviour. We were able to recover each model's behaviour successfully. The entries in **d** show each model's posterior probability of generating all simulated choice data sets, assuming a uniform prior. All models were by far the most likely to produce their own choice data.



Supplementary Figure S5. Anatomically defined region of interest used for small-volume correction. The mask comprises the bilateral hippocampus, entorhinal cortex and subiculum



Supplementary Figure S6. The correlation between spatial and the temporal kernels is not related to behavioral performance measures or hippocampal map representations. The correlation between the spatial and the temporal kernel is plotted against and percent correct in the choice task (a), inference error (b), spatial effect on choice behavior (c), temporal effect on choice behavior (d) and fMRI cross-stimulus enhancement effect in the hippocampus for spatial (e) and temporal distances (f). Parameter estimates in e and f are extracted from the region of interest depicted in Figure 4a. None of the correlations reach significance (all $p > p0.2$).

2 Supplementary Methods

3 Deriving the temporal kernel

Given a participant's exploration run from day 1, we want a method for obtaining a transition matrix $\mathbf{T}(s, s')$, whose entries reflect the participant's propensity for venturing directly from stimuli s to stimuli s' . The successor representation (SR) is captured in the matrix \mathbf{M} , where entries $\mathbf{M}(s, s')$ equal the expected discounted number of future visits to stimulus s' , starting from s . If we know the transition matrix \mathbf{T} governing the one-step transition probabilities between every pair of stimuli, we can define the SR matrix \mathbf{M} as the following infinite sum of \mathbf{T} raised to the power of t

$$\mathbf{M} = \sum_{t=0}^{\infty} \gamma^t \mathbf{T}^t \quad (7)$$

4 where γ is the discount factor. This infinite sum can be computed analytically with matrix inversion

$$\mathbf{M} = (\mathbf{I} - \gamma \mathbf{T})^{-1} \quad (8)$$

5 where \mathbf{I} is the identity matrix. Since we can compute the SR matrix \mathbf{M} analytically from the transition matrix \mathbf{T} , we can
6 attempt to recover the transition matrix from the SR matrix. This is fairly simple using matrix algebra. Since taking the matrix
7 inverse of an inverted matrix gives us the uninverted matrix, $\mathbf{A}^{-1-1} = \mathbf{A}$, we obtain

$$\mathbf{M}^{-1} = (\mathbf{I} - \gamma \mathbf{T}) \quad (9)$$

8 From Equation 9 we obtain \mathbf{T} by subtracting the identity matrix \mathbf{I} , and dividing by $-\gamma$.

$$\mathbf{M}^{-1} - \mathbf{I} = -\gamma \mathbf{T} \quad (10)$$

$$\frac{\mathbf{M}^{-1} - \mathbf{I}}{-\gamma} = \mathbf{T} \quad (11)$$

9 leaving us with the transition matrix \mathbf{T} , which is such that performing an infinite random walk on it produces the SR matrix
10 asymptotically.

11 Temporal relations explain reward generalization in the choice task

12 We sought to verify that the particular exploration trajectory a participant took on day 1 actually influenced how that participant
13 generalized about value, and that the predictive performance of the temporal and the spatio-temporal model could not be
14 attributed to other, more general properties of the temporal kernels, for instance, that they are generally similar to the spatial
15 kernel. To test this, we shuffled the assignments of the temporal kernels, so that each participant would have their choices
16 predicted based on a kernel computed from an exploration trajectory they *themselves* had not taken. If participant choices
17 and generalization were really driven by their specific temporal interaction with the stimuli, then the predictive performance
18 of a model based on the shuffled kernels should be substantially worse than the performance of a model using the correct
19 exploration trajectories. We made the assignments symmetric (for a select pair of participant, their temporal kernels were
20 swapped), and unique (no two participants could be assigned the same temporal kernel). As can be observed in Figure S4a,
21 there were several temporal kernels that were substantially correlated with each other. We reasoned that swapping correlated
22 kernels would yield smaller differences in predictive performance. We therefore sought to generate our shuffled assignments so
23 that the overall correlation would be as small as possible. To do this, we sampled new kernels for each participant based on their
24 inverse correlation r^{-1} to the participant's true kernel. We sampled from a distribution obtained through a softmax transform

$$p_i(K_j) = \frac{\exp(r_{ij}^{-1}/\tau)}{\sum_j^M \exp(r_{ij}^{-1}/\tau)} \quad (12)$$

25 where K_j is the kernel of participant j , M is the number of participant minus participant i and those already assigned, and
26 τ is the temperature parameter. τ plays a key role here, as it allows us to control the degree to which we sample exclusively
27 from the least correlated kernels, as opposed to more uniformly from all other kernels. As τ increases, the distribution gets
28 more uniform. We collected negative log-likelihoods from the temporal model predicting participant choices (Figure S4b).

29 We sampled 10 kernel assignments for 10 evenly spaced values for τ between 0.01 and 3, leaving us with 100 samples of
30 shuffled assignments, where the assignments were d with various degrees of uniformity. Substantiating the hypothesis that
31 participant-specific temporal relations guide generalization, we observe that for all values of τ , the shuffled assignments (the
32 red line) produce substantially worse fits (negative log-likelihood) to the choice data on average than the model using each
33 participant’s true temporal kernel (dashed blue line). Moreover, we observe that this loss is at its highest on average when we
34 sample kernels more concentrated based on inverse correlations (lower τ), as opposed to more uniformly (higher τ) from the
35 set of all kernels.

36 **Hyper-parameters**

37 The successor representation was learnt with temporal-difference learning, using a discount rate γ of 0.9. The signal variance
38 parameter σ_f^2 of the Gaussian kernel (Equation 2) was set to 1, and the observation noise parameter σ^2 (Equation 1) was set
39 to 0.01. The lengthscale parameter λ of the diffusion kernel (Equation 6) was set to 1. For the spatial model, the best-fitting
40 lengthscale λ was 1.242. For the temporal model, the best-fitting learning rate η was 0.4125. For the spatio-temporal model,
41 the best-fitting lengthscale λ was 2.05, and the best-fitting learning rate η was 0.01. To create the kernel matrices used as
42 predictors in the fMRI analyses, we used the spatial kernel with a lengthscale of 2.05 and the temporal kernel with a learning
43 rate of 0.01, which gave the best fit for the spatio-temporal model. These best-fitting hyper-parameter configurations were used
44 in modelling value ratings, and for the model recovery.