www.sciencemag.org/content/352/6292/1464/suppl/DC1

Supplementary Materials for

Organizing conceptual knowledge in humans with a gridlike code

Alexandra O. Constantinescu,* Jill X. O'Reilly, Timothy E. J. Behrens*

*Corresponding author. Email: alexandra.constantinescu@magd.ox.ac.uk (A.O.C.); behrens@fmrib.ox.ac.uk (T.E.J.B.)

Published 17 June 2016, *Science* **352**, 1464 (2016)
DOI: 10.1126/science.aaf0941

This PDF file includes:

Other Supporting Online Material for this manuscript includes the following:
(available at www.sciencemag.org/content/352/6292/1464/suppl/DC1)

Movie S1

Materials and Methods

Participants and experimental design

28 participants (16 men; mean age 26) reported to be in good health with no history of neurological or psychiatric disease and with normal or corrected-to-normal eye vision participated in the fMRI experiment at the Oxford Centre for Functional MRI of the Brain (FMRIB). Each subject participated in 1-4 separate fMRI sessions, across two separate days, at least one week apart. 22 of them participated in both days, whereas the remaining 6 subjects participated only in the first day and did not return for the second testing day. We split in half the data of the subjects who were scanned only once for subsequent analyses that required cross-validation between separate session. In total, we acquired 94 separate datasets. One participant was excluded from all further analyses because he did not learn how to use the controller for navigation in abstract space. Four participants were trained less than the rest of the group. To avoid a confounding effect of the training regimes, we excluded them from the correlations with behavior. All remaining participants were included in further analyses with the exception of one participant who was excluded from the consistency analysis in the prefrontal cortex as a result of parameter estimates being more than 4 s.d. away from the group mean. Subjects were firstly trained outside the MRI scanner to navigate in a continuous abstract "bird space". They had to remember which locations were paired with an outcome and then, on a separate day, they did the task in the scanner. We created two versions of the task in the scanner and half of the participants performed one version and the other half performed the other version. All participants gave written informed consent and were paid for participating, as approved by the ethics committee of the University of Oxford (ref. number 2013-171).

Behavioral training

We developed a set of novel tasks for navigation in an abstract space that were designed to be analogous to those used for navigation in a physical space. We trained the participants extensively to be able to navigate precisely in this abstract space. Using keyboard button presses, subjects learned to morph the birds, to explore the abstract space and learn the locations of the outcomes.

Each location in this abstract space was represented by a bird stimulus. A trajectory through the bird space was equivalent to morphing one bird stimulus into another, that is, watching the neck and legs stretching vertically. The direction $\boldsymbol{\theta}$ of the trajectory was the ratio between how much the neck and legs were changing relative to each other. Subjects were instructed to learn which bird shape was associated with each outcome and we purposefully did not instruct them that these associations could actually be organized in a continuous 2D "bird space". Movie S1 illustrates a recording of the task and the rationale of our experimental design.

Match task:

Firstly, subjects played the *match task* where they were trained to use a controller to morph the bird silhouette shown on the left side of the screen to match the bird on the right side of the screen (Fig. S1). The two birds varied only in two dimensions: the lengths of the neck and legs. Thus, the participants could morph the bird by changing the lengths of the neck and legs. To ensure that participants integrated between these two dimensions simultaneously, we precisely instructed them to morph the bird by choosing the ratio between how much the neck and legs change relatively to each other only once and as correctly as possible (as opposed to changing the lengths of the neck and legs separately). To ensure that subjects did not do this task spatially, we

created a non-spatial controller for choosing this ratio. The participants pressed buttons on the keyboard to choose the ratio and then to morph the bird with that ratio. The morphing was continuous and they saw a bird shrinking or stretching its neck and legs. Each of the two black bars of the controller signified how much that dimension would change. If the black bar was positioned on the midline, the corresponding dimension did not change. The closer to the top line the bar was, the more that dimension would increase. Conversely, the closer to the bottom line the bar was, the more that dimension would decrease.

Explore task:
After the participants learned how to perform the *match task* correctly, they continued with the *explore task* (Fig. S2). In this game, subjects needed to find six outcomes (images of Santa Claus, a Christmas Tree, a gingerbread man, a sledge, a bell and a teddy-bear) by freely morphing the bird. Each outcome was associated with a different bird (Fig. 1A). Participants knew they found an outcome because it would pop up next to its associated bird. Critically, this ensured that subjects became familiar with the entire bird space and not just with the stimuli paired with outcomes and that they had a clear representation of what one unit distance on each axis (neck, legs) looked like. Thus, we did not simply train the participants on the mere associations between the six birds and the six outcomes (i.e. only the locations paired with outcomes in the abstract space). This approach is analogous to exploration tasks in open environments in rodents(29). Indeed, participants did not know a priori which birds were paired with outcomes, hence, they had to morph through many different unpaired birds to find the paired ones (i.e. they explored many locations in the abstract space that were not paired with outcomes).
The locations of the outcomes were identical in all participants. This allowed us to identify learning patterns of the environment across participants, and rule out the possibility that inter-individual differences between participants were caused by the outcome locations. We carefully selected these locations to rule out any potential hexagonal arrangements.

Collect task:
When subjects knew the stimulus-outcome associations well, they continued to train with the "collect task". In this task, we periodically tested their knowledge of the bird space by asking them to find specific outcomes from random start positions, by correctly choosing the neck:legs ratio only once. Thus, they had to recall the correct neck/legs ratio associated with that outcome, and generate the abstract trajectory to that bird correctly in one go.

Recall task during scanning (Fig. 1):
Subjects were trained with the version of the task that would also be used subsequently during scanning (Movie S1). During fMRI scanning, we wanted to present participants with a controlled set of trajectories, thus, we devised a task in which participants watched videos of birds morphing into different birds according to pre-defined neck:legs ratios. Each trial corresponded to one trajectory and consisted of several stages: morphing, imagination and, in a subset of trials, choice. During the morphing stage, the participants saw only a single trajectory in the abstract space.

Next, they were instructed to imagine the bird continued to morph with the same neck:legs ratio, for the same amount of time and the same speed, and recall if any of the imagined bird stimuli was associated with an outcome. This means they did not only need to estimate equivalents of angles of moving in space, but also equivalents of distances. In other words, we split the trajectories in half, and subjects saw the first half (visual morphing of the bird) and had to imagine the second half (imagined morphing of the bird). Finding grid-like representations during imagined navigation in physical space is indeed possible, as shown in a recent study(11).

We deliberately designed the trajectories to sample the 360° space uniformly, across all trials, as well as separately for trials paired our unpaired with outcomes.

Finally, subjects could choose one out of three possible outcomes (outcome 1, outcome 2 or no outcome) by pressing one out of three keys on a button box. To control for the effect of the outcomes, we designed two versions of the task. In one version, these outcomes were displayed during the entire duration of the trial and subjects made a choice immediately after the imagination period. In the other task version, participants only saw the outcomes at the end of the imagination period, with a 50% probability, and then were asked to make a choice. There were no significant differences between the two versions of the task, at either the neural or behavioral levels, therefore we pooled data from both versions. Overall, the accuracy on the task was 72.8 ± 1.02% (mean ± SEM across participants – chance is 33%), suggesting that they learned how to navigate the bird space.

Post-scanning debriefings

Even though the stimuli were not 2D spatial, it would have been possible for the subjects to solve the problem using conscious and explicit analogy to 2D space. Such a spatial strategy would go beyond the scope of previous reports of the function of grid cells but would retain space as a 2 dimensional anchor for the grid. Alternatively, subjects may have had no conscious knowledge of the spatial arrangement of the birds. To gain some handle on this interesting question, we asked subjects after scanning to describe the strategy they used to do the task and to remember the associations between the birds and the outcomes. For example, some strategies for remembering what bird gave the Christmas Tree were: visualizing the precise lengths of the neck and legs as a mental picture, using internal metrics such as the bird has 4 units of neck and 1 unit of legs, verbally repeating in their minds “long neck, short legs”, finding similarities between the characteristics of the bird and the outcome such as the bird had a long neck (trunk) and short legs (base) like a true Christmas Tree and, also, finding similarities between different birds associated with outcomes (e.g. the Christmas Tree bird had longer neck but similar legs with the Teddy Bear bird). Thus, no participant described being aware that birds could be associated spatially, on a two-dimensional map. For certainty, we subsequently showed them the map of the abstract space (Fig. 1B) and asked them whether this spatial strategy was “easier”, “harder” or “the same” as their own strategy. Two participants reported that the spatial strategy had the same level of difficulty as their own strategy, however they clarified that their approach was not spatial and were unaware that the birds could be arranged into a 2D map. Everyone else reported that the spatial strategy was easier or harder than their own. Overall, no subject reported doing the task spatially.

MRI data acquisition

We acquired T2-weighted functional images on a 3 Tesla scanner. We used a gradient-echo echo-planar imaging (EPI) pulse sequence that sets the slice angle of 30° relative to the anterior-posterior commissure line, minimizing the signal loss in the orbitofrontal cortex region(30). We acquired 45 slices, 3mm thick, in an interleaved order. We used the following parameters: repetition time (TR) = 3000 ms, echo time (TE) = 30 ms, flip angle = 87°, field of view (FoV) = 192mm, voxel size = 3 x 3 x 3 mm^3. To correct for deformations in the inferior prefrontal cortex, we also acquired a field map with dual echo-time images covering the whole brain, with the following parameters: TR = 500 ms, TE1 = 5.19 ms, TE2 = 7.65 ms, flip angle = 60°, FoV = 224mm, voxel size = 3.5 x 3.5 x 3 mm^3. To facilitate an accurate registration of the EPIs to the standard space, we also acquired a T1-weighted structural image using a magnetization-prepared rapid gradient echo sequence (MPRAGE) with the following parameters: TR = 2040 ms, TE = 4.7 ms, flip angle = 8°, FoV = 192mm, voxel size = 1 x 1 x 1 mm^3. Stimulus presentation and subject button presses were registered and time-locked to the fMRI data.

Data analysis

We analyzed whole-brain MRI data with FMRIB's Software library (FSL)(31) and we performed all the other analyses in Matlab.

Behavioral analyses

We tested for measures of learning during the "explore" task, where subjects navigated freely to look for the outcomes. We concatenated the data from all training sessions in each subject, and then split it into five equal parts (quantiles). We then computed the amount of time spent in each part of the environment during each quantile and plotted it as color-coded trajectory maps. We then estimated the amount of time spent navigating at the edges ("time at edges") and at the locations paired with outcomes ("time at outcomes"), relative to the total time spent navigating. At the group level, we averaged these measures across participants.

In the "collect" task, subjects were precisely instructed to find specific outcomes. We computed the mean number of transitions needed to find an outcome, in each trial. A perfect performance is equivalent to making only one transition to find the target outcome, thus, choosing the correct neck:legs ratio only once and generating the abstract trajectory to the target bird in one go. In contrast, a poor performance means making more than one transition. Therefore, we also computed the percentage of trials where participants made only one transition. Finally, a precise trajectory would have a small angle error, defined as the angle between the ideal trajectory angle and the angle of the first transition. Thus, we next computed the percentage of trials where participants had an angle error < 15°.

In the recall task in the scanner, we computed the performance as percentage of correct responses in each scanning session.

Pre-processing of functional images

We segmented brain matter from nonbrain(32). We corrected functional data for motion artefacts and we removed low frequency signals using a high-pass filter at

1/100 Hz. We smoothed the data using a gaussian filter of 7 mm full width at half maximum and we corrected for slice time acquisition differences. We also corrected for geometric distortions in the EPI images using the acquired fieldmaps(33). We registered the EPI images to the high-resolution anatomical image using boundary-based reconstruction and then we normalized them into standard space (Montreal Neurological Institute – MNI152) using non-linear registration. Because of the notable breathing- and susceptibility-related artefacts in the enthorinal cortex, we ran a separate analysis where we pre-cleaned the data with FMRIB's ICA tool, FIX (34,35).

Measures of motion estimation

We computed four parameters to assess motion in each participant and identify movements bigger than the voxel size (3mm). First, we computed the mean displacement, defined as the amount of head motion (sum of the absolute translational movement in the x, y and z directions) relative to the middle time point across the entire scanning session, averaged across all time points.

Second, we examined the total displacement (TD) from the beginning to the end of the scanning session.

Third, we examined the data for rapid movements because these are the ones that could have disrupted signal intensities. We computed the mean framewise displacement, as the amount of head motion relative to the previous time point, averaged across all time points.

Finally, we computed the maximum framewise displacement (MaxFD) to address the concern of sudden movement more conservatively. We looked for subjects who moved more than 3mm at any time point.

fMRI whole brain analyses

After pre-processing, we modeled the fMRI time series using two general linear models (GLMs). Both models included regressors for the main effects of the morphing stage and the response stage of the trial and six nuisance regressors to account for motion-related artifacts. Each individual model had various parametric modulations of the regressors for the morphing stage (see below). All these regressors were convolved with the FSL default hemodynamic response function and filtered by the same high pass filter as the fMRI data before entering the GLM.

GLM 1: hexagonal modulation analysis (Fig. 2):

We created two parametrically modulated regressors(9): the sine and cosine of the direction of each trajectory in each trial, $\theta(t)$, with a periodicity of 60°, that is, $\sin(6\theta(t))$ and $\cos(6\theta(t))$. The factor 6 means that these regressors will produce coefficients with high amplitude $\left(\sqrt{\beta_{sin}^2 + \beta_{cos}^2}\right)$ for brain regions that are sensitive to hexagonal symmetry. To calculate the magnitude of the hexagonal modulation, we used an F test to investigate which brain regions were significantly modulated by a linear combination of these two regressors: $\beta_{sin}*\sin(6\theta) + \beta_{cos}*\cos(6\theta)$. We then transformed this statistic to a Z-statistic in each subject, using an asymptotic

approximation (http://www.fmrib.ox.ac.uk/analysis/techrep/tr00mj1/tr00mj1/) and, together with a behavioral covariate, we performed a 1-sample t-test across the group. For visualization purposes, the statistic images were thresholded using clusters determined by Z-statistic = 3.1 and the supra-threshold clusters were corrected for family-wise error using a cluster significance threshold of p = 0.05. For the entorhinal cortex, we used a more lenient cluster forming threshold at Z = 2.3 and p = 0.05 because this brain region is prone to signal loss and, thus, it is difficult to image with fMRI (see below for discussion of these tests).
This test has the major benefit of identifying periodic signals across the whole brain without knowing the phase of these signals. However, although the reported clusters survive multiple comparison correction across the whole brain, this test is unusual and should not be used for statistical inference because it can overestimate the Z-scores. Thus, we used it instead as a method to identify informative brain regions that we could then use to create ROIs for subsequent unbiased tests.
In detail, the first level statistic that is being brought to the group level relies on the mean and variance of the first-level estimates in such a way that mis-estimating the variance can cause a bias in the mean (for example if the variance is under-estimated, the F-statistic will be overestimated, and the expectation of the transformed Z-statistic will be greater than zero even in null data). Estimating unbiased variances from autocorrelated fMRI data is notoriously difficult. Therefore, the reader should be cautious in interpreting the results of this test as statistically significant. In the current manuscript, we use this test as a method for generating candidate regions of interest for tests that are certainly unbiased (consistency tests). This allows us to extract the grid angle from these ROIs in one dataset and test it in another. These subsequent tests do not use first level variances and are therefore unbiased tests performed with the standard fMRI GLM machinery.

GLM 2: hexagonal consistency analysis (Fig. 3A):

We created a parametric regressor $\cos(6[\boldsymbol{\theta}(t) - \boldsymbol{\varphi}])$, where $\boldsymbol{\theta}(t)$ is the trajectory angle for each trial t within one session and $\boldsymbol{\varphi}$ is the grid angle from another session acquired on the same day(9,10) (Fig. S3). This regressor had higher values for grid-like activity with the same grid angle as in the other session. Thus, if the trajectories from one session aligned with the predicted grid angle taken from a different session, then this regressor would have high values. This resulted in subject-specific parameter estimates for each experimental session.

In the next stage, participant-specific linear contrasts of these parameter estimates were averaged across the experimental sessions and, together with a behavioral covariate represented by the behavioral accuracy, they were entered into a series of one-sample t-tests, each constituting a group-level statistical parametric map. The statistic images were thresholded using clusters determined by Z-statistic = 2.3 and the supra-threshold clusters were corrected for family-wise error using a cluster significance threshold of p = 0.05.

To estimate the grid angle $\boldsymbol{\varphi}$, we used a similar approach to the one for navigating in physical space(9). We first averaged the coefficients for the sine (β_{sin}) and cosine (β_{cos}) regressors across all voxels in the ROI. Then, we calculated the grid angle $\boldsymbol{\varphi}$ (varying between 0 and 59°) as $\arctan[(\beta_{sin}/\beta_{cos})]/6$ within regions of interest (ROIs),

where arctan was mapped into the 360° space, varying between -30° to 30°, according to the signs of β_{sin} and β_{cos}.

ROI analyses (Fig. 3-4)
We defined all ROIs from the whole-brain quadrature filter analysis for hexagonal modulation (GLM 1 and Fig. 2), which was orthogonal to the contrasts of interest. This allowed statistical tests to be performed in an unbiased fashion. To make it clear why this is the case: in the null case, two scans that both had high hexagonal magnitude would have randomly oriented hexagonal phase (grid angle) and so all analyses that relied on aligned grid angle would result in mean zero signal. The ROIs were 5-mm spheres centered on the voxels with the highest hexagonal magnitude activity (GLM 1) in the ventromedial prefrontal (vmPFC) and the entorhinal cortex (ERH) (Fig. 2A). We used these ROIs as seeds for estimating the values of the grid angle and to test for grid angle consistency between separate experimental sessions.

To examine the pattern underlying the hexagonal magnitude effect, we ran separate regression analyses for trajectories that were aligned and misaligned with the grid, using fMRI time series from the vmPFC and ERH ROIs. We aligned the trajectory angles to the grid angle φ and we split these aligned trajectories into 12 equal bins of 30°. We thus created 12 separate regressors for trajectories that belonged to one of these 12 bins, resampled them at a resolution of 100ms and convolved them with the hemodynamic response function. To match the resolution of the fMRI timeseries, we then resampled the regressors to the duration of the fMRI repetition time. Next, we ran 12 separate regression analyses in each individual experimental session. The 12 resulting regression coefficients were then averaged separately across various sessions, and expressed as mean ± sem (Fig. 3B-C and Fig. 4). Given our strong *a priori* hypotheses of hexagonal symmetry, we performed a t-test on the differences between the resulting betas aligned and misaligned to the grid angle φ, across subjects, to investigate if these differences were significantly different from zero.

We tested for consistency within-day (Fig. 3), across-day (Fig. 4A), and all data combined within- and across-day (Fig. 4B-C). Again, to allow statistical unbiased testing, we aligned data from one session to the grid angle from the other session.

Control analyses

fMRI data:
To test for the specificity of the six-fold symmetric sinusoidal modulation, we used the same approaches as above, but using control models with directional periodicities of 90°, 72°, 51.4° and 45° (that is, four-fold, five-fold, seven-fold and eight-fold rotational symmetry).

Supplementary Text

Ruling out alternative explanations for the fMRI hexagonally symmetric signal

The trial-by-trial fMRI activity resembles a periodic wave with six peaks only if the angles of the trajectories the participants are running at are sorted in an ascending order (modulo 60) from 0° to 360° (Fig. 1E, Fig. S4A). However, participants did not move along trajectories in a sorted order, but in a random fashion and the signal looks random when the order of the trajectory angles is shuffled (Fig. S4C). Moreover, the trial-to-trial fMRI activity is different from day-to-day in each participant because they navigated along various random angles in each day. These responses are also different from person-to-person because each subject has their own grid angle. For example, if two participants have grid angles at 30° from each other, then trajectories that are aligned in one person, thus leading to fMRI responses with big amplitudes, will be misaligned in the other person, leading to fMRI responses with low amplitudes.

Therefore, potential confounds of this signal need not resemble a periodic wave with six peaks regardless of the order of the trials. Instead, the artifacts need to be consistently and significantly happening with a hexagonal periodicity that depends precisely on the order of the trajectories of moving in space, be aligned to one grid angle in one participant and to another grid angle in another participant. It is therefore unlikely that motion or physiological artifacts (such as breathing and heart rate) would cause the hexagonal pattern. For example, subjects would need to move more during aligned versus misaligned trajectories, according to the random order these trajectories are presented, which is different from day-to-day, and according to the orientation of the grid, which is different from person-to-person. The hexagonal pattern of fMRI activity is therefore difficult to explain with unspecific fMRI artefacts.

Subject estimates of motion

The average mean displacement across participants was 0.4 ± 0.043 mm, thus less than 1 voxel-diameter (3 mm).

We found 7 subjects who had a total displacement > 3mm. This is a very conservative measure of subject movement as it penalizes many incremental movements that are well corrected by linear methods. Nevertheless, we tested using region of interest analyses whether our results survived the exclusion of these 7 subjects. Despite the inevitable loss of statistical power entailed by removal of data from more than 25% of the participants, we were able to confirm the majority our findings. All results presented except for two (the correlation between behavior and hexagonal consistency in vmPFC and the hexagonal consistency vmPFC within-day) remained significant in this re-analysis (Table S1). We next focused on measures of rapid motion of the kind that cause a bigger problem for fMRI.

The average mean framewise displacement across participants was 0.031 ± 0.0045

mm, thus less than 3mm. Indeed, participants did not tend to make sudden head movements throughout scanning.

We found that 3 subjects made 1-2 rapid head movements in one scanning session, as assessed by the maximum framewise displacement (MaxFD). It is very unlikely that these events could have potentially affected the observed hexagonally symmetric signal because they did not occur frequently, with hexagonal periodicity, and they would have needed to be precisely aligned with the grid angle. Despite this improbable situation, we tested again using region of interest analyses whether our results survived the removal of these 3 subjects. All the presented results remained significant in this re-analysis, except for one, which approached significance (the correlation between behavior and hexagonal signal modulation in vmPFC).

To sum up, we re-analyzed data after excluding subjects who moved more than 3mm in the scanner. Although the occasional results changed when we remove these subjects, in each case the vast majority of assessments survived. The statistical tests from these two re-analyses, one that is common within FSL and SPM analysis standards (MaxFD) and another more conservative measure (TD) are presented in table S1.

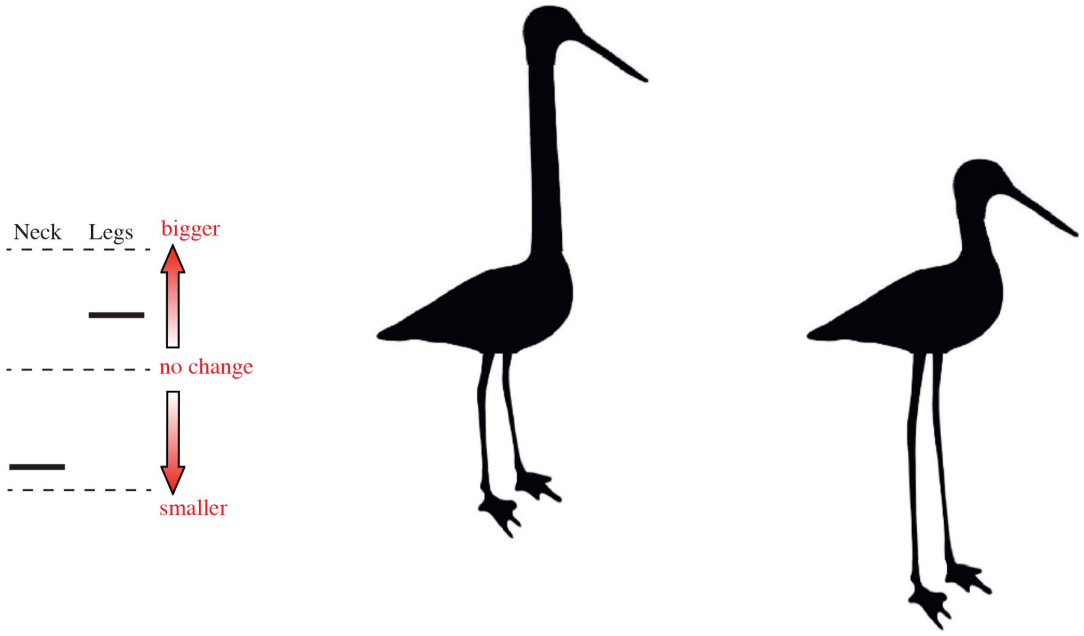

Fig. S1.

Example screenshot of the match task. Here, to morph the bird on the left into the one on the right, the neck needed to decrease more than the legs needed to increase. The correct neck:legs ratio to achieve this is shown on the controller on the left side of the screen (notice the position of the two black bars).

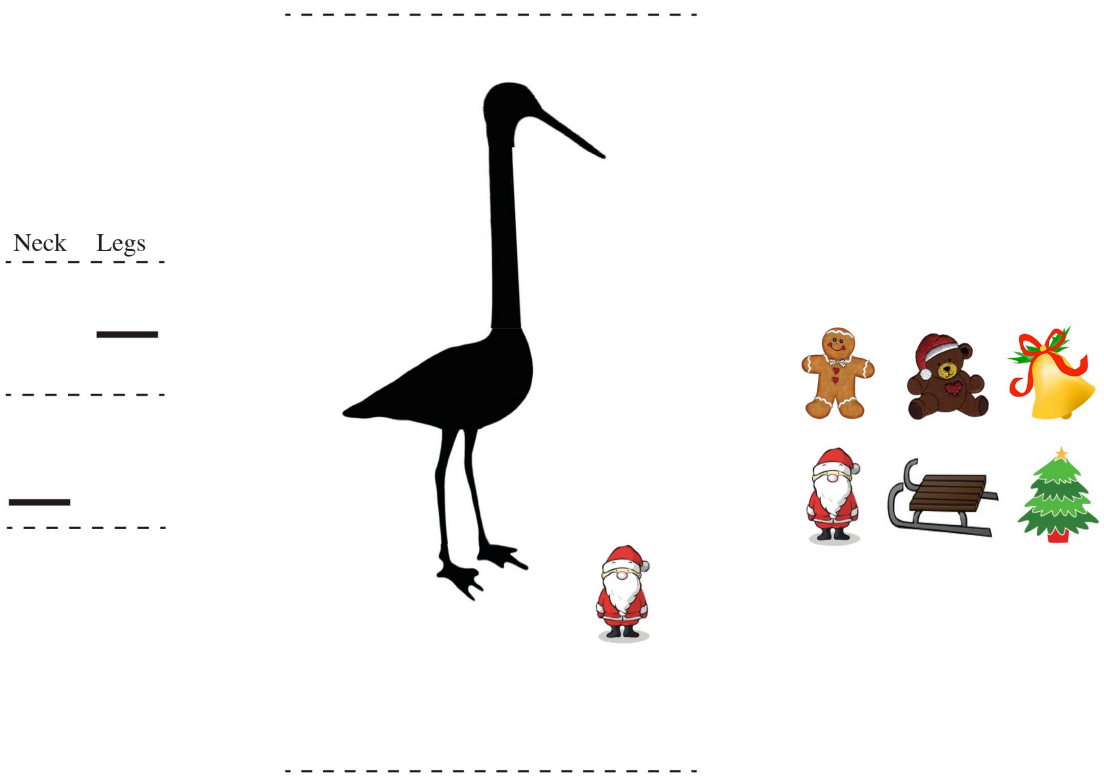

Fig. S2
Example screenshot in the explore task where the subject found Santa Claus.

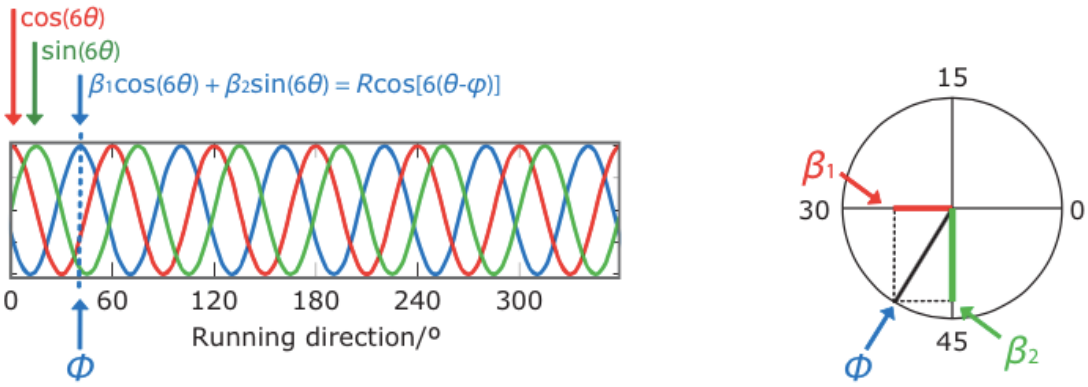

Fig. S3

Estimation of the grid angle (reproduced with permission from Nature, Doeller et al, 2010).

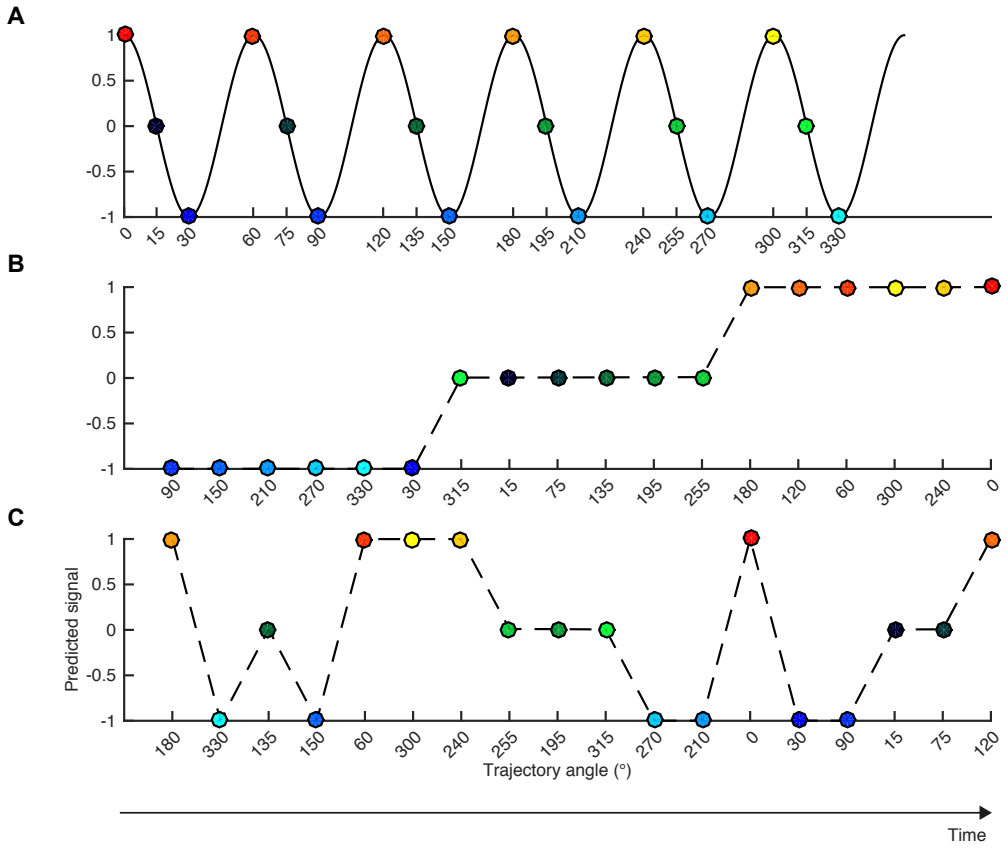

Fig. S4

(A) The hexagonal symmetry pattern in the fMRI response becomes evident when the trajectory angles are sorted in an ascending order. **(B)** When the angles are ordered first as completely misaligned with the grid angle φ, then right in the middle between aligned and misaligned, and then completely aligned with the grid angle, the fMRI response will show a step-wise pattern, from a small response, to no change, and to a big response. **(C)** When shuffled randomly, the fMRI response also becomes random.

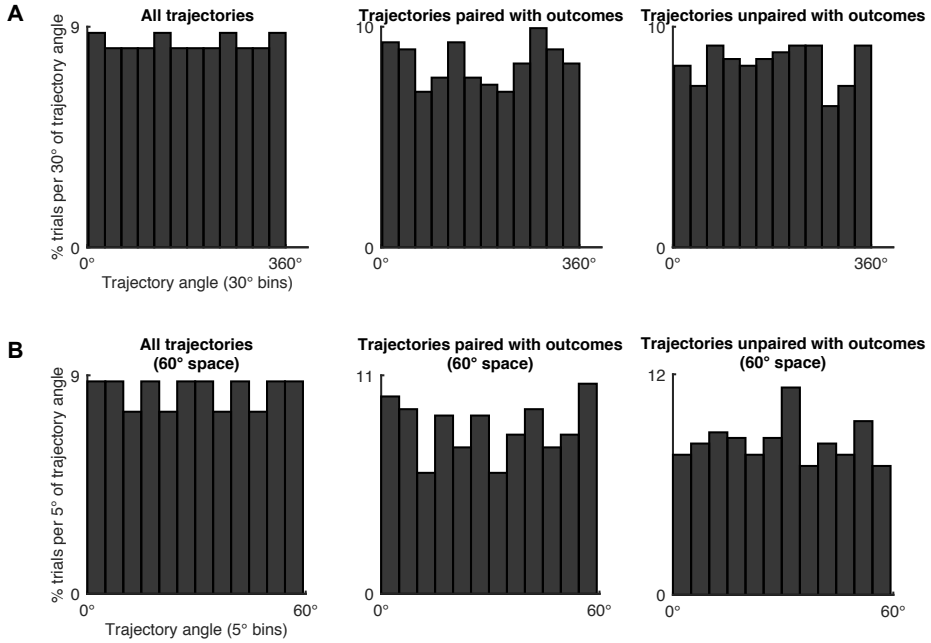

Fig. S5

Distribution of trajectory angles. **(A)** Trajectory angles were evenly sampled across the 360° space (varying between 0°-360°). Histograms are shown for trajectory angles passing through all locations in conceptual space, and categorized as a function of passing through an outcome or not (Rayleigh's test, p=1, p=0.626 and p=0.64, respectively). **(B)** To investigate the distribution of angles relative to hexagonal symmetry, we transformed the angles into a 60° space (varying between 0°-60°), using the modulo operation, which computes the remainder after division of the angles by 60. For example, a 70° trajectory in 360° space corresponds to a 10° trajectory in 60° space. We found that the trajectory angles were evenly sampled across the 60° space for all locations in sum, and separately for the locations paired and unpaired with outcomes (Rayleigh's test, p=1, p=0.623 and p=0.637, respectively).

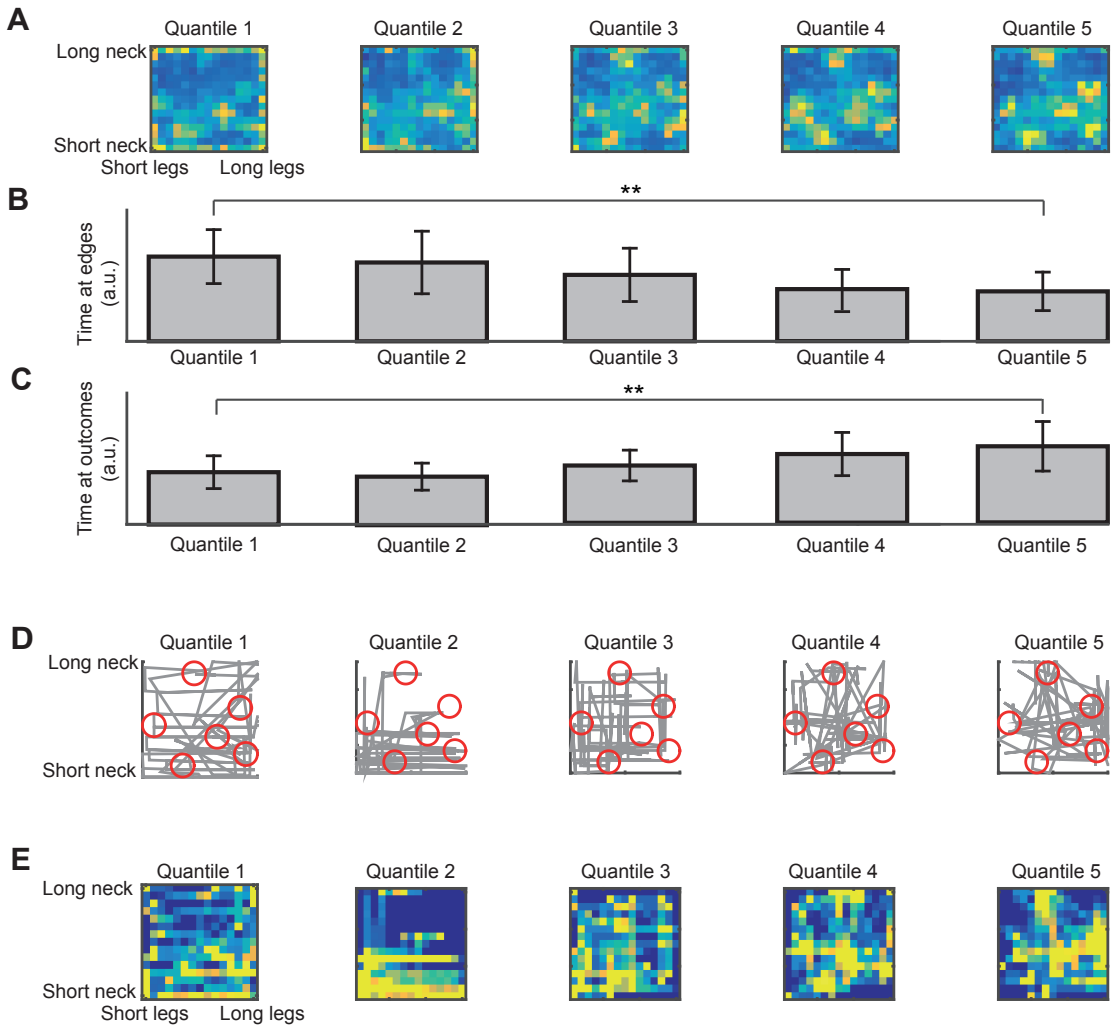

Fig. S6

Analysis of the explore task at the group level (A-C) and in one representative subject (D-E). (A) Color-coded trajectory maps illustrating the amount of time spent in each part of the environment (yellow is maximum, dark blue is 0). Note how subjects first navigate in most parts of the abstract space, choosing random trajectories to find the outcomes, and spend most of their time near the boundaries of bird-space. As training progresses, (B) they spend less time at the edges (times at edges, quantile1 vs quantile5, t_{22}=3.1776, ** p<0.01) and (C) more time navigating directly between locations paired with outcomes (time at outcomes, quantile1 vs quantile5, t_{22}=-3.1773, ** p<0.01). These results demonstrate that the identities of the birds paired with outcomes were well learnt in the subject cohort, conceptually equivalent to learning object locations in a 2D spatial environment. This means that subjects did not only need to indicate if an outcome occurred somewhere on the imagined trajectory, but also at which distance. (D) Trajectories (grey lines) with locations paired with outcomes (red circles) in one representative subject. (E) Corresponding color-coded trajectory maps. Note how this subject spends more time at the edges in quantile1 and explore more the lower area of the environment in quantile 2.

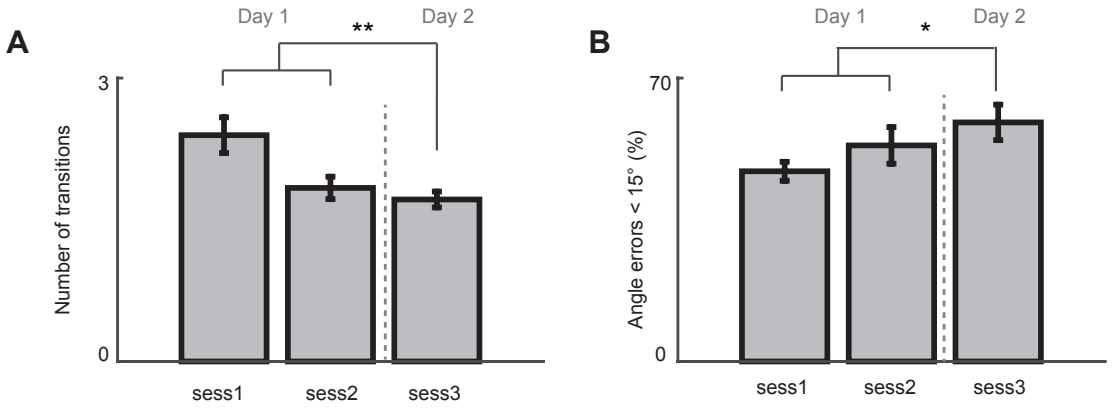

Fig. S7

Measures of learning during the collect task. Participants made significant improvements in training day 2 compared to training day 1. (A) Subjects made fewer transitions (t_{33}=-2.89, ** p<0.01) in day 2 than day1, and (B) the percentage of angle errors < 15° increased (t_{33}=2.37, * p<0.5).

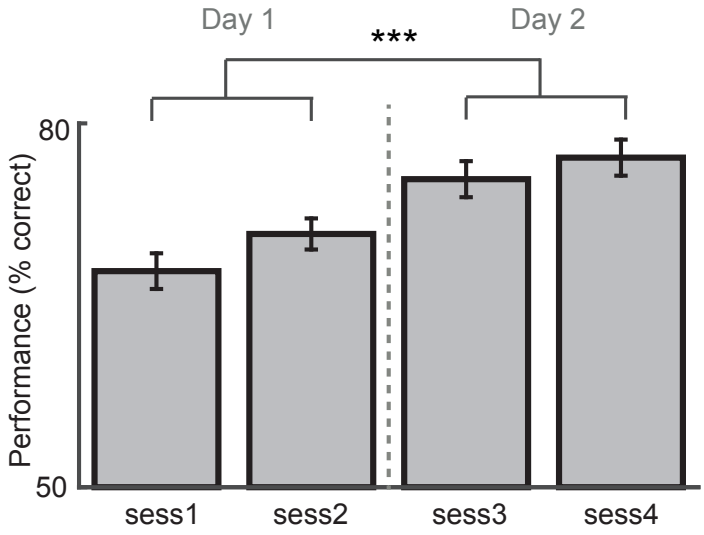

Fig. S8

Learning in the recall task. Effects of fresh training on subsequent performance in the scanner. Participants made significantly more correct responses in day 2 compared to day 1 (t_{41}=3.89, *** p<0.001).

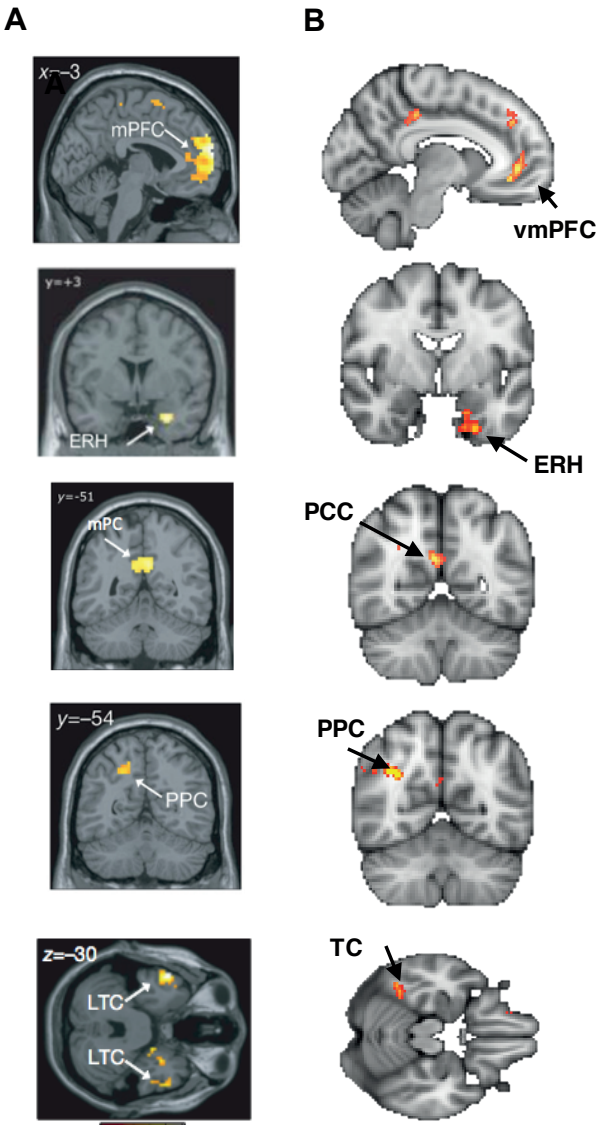

Fig. S9

The brain network for navigation in physical space (**A**) and conceptual space (**B**). Abbreviations: ventromedial prefrontal cortex (vmPFC), entorhinal cortex (ERH), posterior cingulate cortex (PCC), posterior parietal cortex (PCC), temporal cortex (TC). A is adapted by permission from Nature, Doeller et al, 2010.

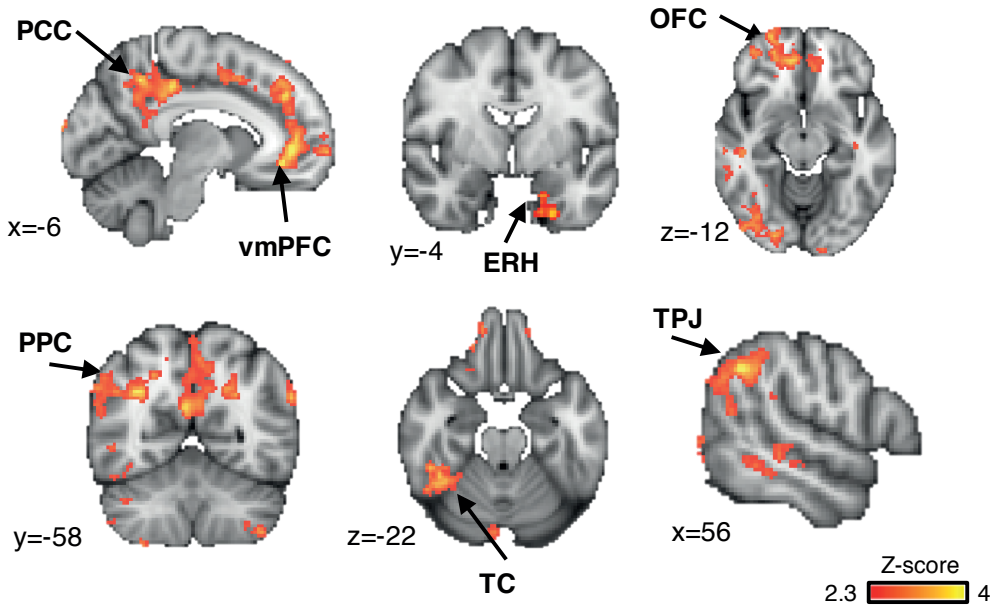

Fig. S10

Same network of brain regions as in Fig. 2A from the main text, but whole-brain cluster corrected at a more lenient cluster threshold Z=2.3 and p<0.05, for illustrative purposes. Note that, as in the main text, this analysis is subject to potential bias caused by the mis-estimation of the autocorrelation of fMRI data(26).

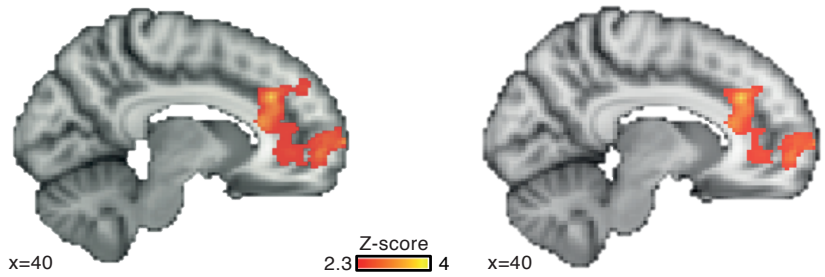

Fig. S11:
The vmPFC consistency map (left and also shown originally in Fig. 3A) did not change after including a confound regressor that modeled out the effect of outcomes (right).

Test	**Presented in**	**Value *including* subjects after motion correction**	**Value *excluding* subjects with MaxFD > 3mm**	**Value *excluding* subjects with TD > 3mm**
Correlation between behavior and hexagonal signal modulation in vmPFC	Fig. 2B	r=0.432, p=0.039	r=0.41, p=0.07*	r=0.4, p=0.09*
Correlation between behavior and hexagonal consistency in vmPFC	Fig. 3A, right	r=0.431, p=0.039	r=0.49, p=0.02	r=0.32, p=0.19**
Hexagonal consistency vmPFC, within-day	Fig. 3B, left	t(26)=2.61; p<0.05	t(23)=2.17; p<0.05	t(20)=1.54, p=0.13**
Control periodicities, vmPFC, within-day	Fig. 3B ,right	All p>0.26	All p>0.3	All p>0.29
Hexagonal consistency ERH, within-day	Fig. 3C, left	t(27)=2.36; p<0.05	t(24)=2.3, p<0.05	t(21)=2.65; p<0.01
Control periodicities, ERH, within-day	Fig. 3C, right	All p>0.15	All p>0.18	All p>0.31
Hexagonal consistency vmPFC, cross-day	Fig. 4A, left	t(20)=3.65; p<0.01	t(17)=2.97; p<0.01	t(15)=2.77, p<0.05
Control periodicities, vmPFC, cross-day	Fig. 4A, right	All p>0.18	All p>0.3	All p>0.12
Hexagonal consistency vmPFC, within- and cross-day	Fig. 4B, left	t(20)=3.41, p<0.01	t(17)=2.66; p<0.05	t(15)=2.35,p<0.05
Control periodicities, vmPFC, within- and cross-day	Fig. 4B, right	All p>0.29	All p>0.6	All p>0.54
Hexagonal consistency, cross region ERH --> vmPFC	Fig. 4C, left	t(21)=2.04, p=0.053	t(18)=2.21; p<0.05	t(16)=2.32, p<0.05
Control periodicities, cross region ERH -->	Fig. 4C, right	All p>0.8	All p>0.55	All p>0.25

vmPFC				

Table S1.

Statistical tests from the data re-analyses after excluding participants who moved more than 3mm. *approaching significance, **no longer significant.

Movie S1

Experimental design. Example trajectories for movement in bird-space at various orientations correspond to morphing of visual stimuli with specific neck:legs ratios. Note that the trajectories were dissociated from the properties of the visual scene, such that trajectories with the same orientation were formed by different stimuli, and trajectories with different orientations could pass through the same stimulus. We designed multiple trajectories in bird-space, and each trajectory corresponded to one trial. Participants had to choose what outcome (Christmas symbol) they received in a trial. They also had the option of choosing "nothing" (black square), meaning that the trajectory did not go through an outcome. We then looked for the hexagonal coding pattern, depending on whether the trajectories were aligned or misaligned with the grid.

References and Notes

1. D. Kumaran, J. J. Summerfield, D. Hassabis, E. A. Maguire, Tracking the emergence of conceptual knowledge during human decision making. *Neuron* **63**, 889–901 (2009). Medline doi:10.1016/j.neuron.2009.07.030

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

3. E. C. Tolman, Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948). Medline doi:10.1037/h0061626

4. J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map* (Oxford Univ. Press, 1978).

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

6. H. Eichenbaum, N. J. Cohen, Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* **83**, 764–770 (2014). Medline doi:10.1016/j.neuron.2014.07.032

7. T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, E. I. Moser, Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005). Medline doi:10.1038/nature03721

8. J. Jacobs, C. T. Weidemann, J. F. Miller, A. Solway, J. F. Burke, X. X. Wei, N. Suthana, M. R. Sperling, A. D. Sharan, I. Fried, M. J. Kahana, Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat. Neurosci.* **16**, 1188–1190 (2013). Medline doi:10.1038/nn.3466

9. C. F. Doeller, C. Barry, N. Burgess, Evidence for grid cells in a human memory network. *Nature* **463**, 657–661 (2010). Medline doi:10.1038/nature08704

10. L. Kunz, T. N. Schröder, H. Lee, C. Montag, B. Lachmann, R. Sariyska, M. Reuter, R. Stirnberg, T. Stöcker, P. C. Messing-Floeter, J. Fell, C. F. Doeller, N. Axmacher, Reduced grid-cell-like representations in adults at genetic risk for Alzheimer's disease. *Science* **350**, 430–433 (2015). Medline doi:10.1126/science.aac8128

11. A. J. Horner, J. A. Bisby, E. Zotow, D. Bush, N. Burgess, Grid-like Processing of Imagined Navigation. *Curr. Biol.* **26**, 842–847 (2016). Medline doi:10.1016/j.cub.2016.01.042

12. C. Barry, R. Hayman, N. Burgess, K. J. Jeffery, Experience-dependent rescaling of entorhinal grids. *Nat. Neurosci.* **10**, 682–684 (2007). Medline doi:10.1038/nn1905

13. H. Stensola, T. Stensola, T. Solstad, K. Frøland, M. B. Moser, E. I. Moser, The entorhinal grid map is discretized. *Nature* **492**, 72–78 (2012). Medline doi:10.1038/nature11649

14. T. Stensola, H. Stensola, M.-B. Moser, E. I. Moser, Shearing-induced asymmetry in entorhinal grid cells. *Nature* **518**, 207–212 (2015). Medline doi:10.1038/nature14151

15. M. D. Fox, M. E. Raichle, Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* **8**, 700–711 (2007). Medline doi:10.1038/nrn2201

16. J. R. Binder, R. H. Desai, W. W. Graves, L. L. Conant, Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb. Cortex* **19**, 2767–2796 (2009). Medline doi:10.1093/cercor/bhp055

17. D. L. Schacter, D. R. Addis, D. Hassabis, V. C. Martin, R. N. Spreng, K. K. Szpunar, The future of memory: Remembering, imagining, and the brain. *Neuron* **76**, 677–694 (2012). Medline doi:10.1016/j.neuron.2012.11.001

18. D. Hassabis, E. A. Maguire, Deconstructing episodic memory with construction. *Trends Cogn. Sci. (Regul. Ed.)*. **11**, 299–306 (2007).

19. J. A. Clithero, A. Rangel, Informatic parcellation of the network involved in the computation of subjective value. *Soc. Cogn. Affect. Neurosci.* **9**, 1289–1302 (2014). Medline doi:10.1093/scan/nst106

20. R. Saxe, S. Carey, N. Kanwisher, Understanding other minds: Linking developmental psychology and functional neuroimaging. *Annu. Rev. Psychol.* **55**, 87–124 (2004). Medline doi:10.1146/annurev.psych.55.090902.142044

21. R. G. Benoit, K. K. Szpunar, D. L. Schacter, Ventromedial prefrontal cortex supports affective future simulation by integrating distributed knowledge. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 16550–16555 (2014). Medline doi:10.1073/pnas.1419274111

22. R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, I. Fried, Invariant visual representation by single neurons in the human brain. *Nature* **435**, 1102–1107 (2005). Medline doi:10.1038/nature03687

23. D. Aronov, R. Nevers, D. W. Tank, CA1 firing fields represent an abstract coordinate during non-spatial navigation. *Proceedings of the Meeting of Computational and Systems Neuroscience*, **T-29**, 41 (2016).

24. B. J. Kraus, M. P. Brandon, R. J. Robinson 2nd, M. A. Connerney, M. E. Hasselmo, H. Eichenbaum, During Running in Place, Grid Cells Integrate Elapsed Time and Distance Run. *Neuron* **88**, 578–589 (2015). Medline doi:10.1016/j.neuron.2015.09.031

25. T. Davis, R. A. Poldrack, Quantifying the internal structure of categories using a neural typicality measure. *Cereb. Cortex* **24**, 1720–1737 (2014). Medline doi:10.1093/cercor/bht014

26. Materials and methods are available as supplementary materials on *Science* Online.

27. Y. Dordek, D. Soudry, R. Meir, D. Derdikman, Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife* **5**, 682 (2016). Medline doi:10.7554/eLife.10094

28. K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, "Design Principles of the Hippocampal Cognitive Map," *NIPS Proceedings, Advances in Neural Information Processing Systems*, 2528–2536 (2014).

29. D. Dupret, J. O'Neill, B. Pleydell-Bouverie, J. Csicsvari, The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* **13**, 995–1002 (2010). Medline doi:10.1038/nn.2599

30. N. Weiskopf, C. Hutton, O. Josephs, R. Deichmann, Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. *Neuroimage* **33**, 493–504 (2006). Medline doi:10.1016/j.neuroimage.2006.07.029

31. S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, P. M. Matthews, Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* **23** (suppl. 1), S208–S219 (2004). Medline doi:10.1016/j.neuroimage.2004.07.051

32. S. M. Smith, Fast robust automated brain extraction. *Hum. Brain Mapp.* **17**, 143–155 (2002). Medline doi:10.1002/hbm.10062

33. M. Jenkinson, P. Bannister, M. Brady, S. Smith, Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825–841 (2002). Medline doi:10.1006/nimg.2002.1132

34. G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser, L. Griffanti, S. M. Smith, Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* **90**, 449–468 (2014). Medline doi:10.1016/j.neuroimage.2013.11.046

35. L. Griffanti, G. Salimi-Khorshidi, C. F. Beckmann, E. J. Auerbach, G. Douaud, C. E. Sexton, E. Zsoldos, K. P. Ebmeier, N. Filippini, C. E. Mackay, S. Moeller, J. Xu, E. Yacoub, G. Baselli, K. Ugurbil, K. L. Miller, S. M. Smith, ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* **95**, 232–247 (2014). Medline doi:10.1016/j.neuroimage.2014.03.034