Flexible and abstract neural representations of structural knowledge

Mark S[2]., Baram A.[1], Schwartenbeck P.[2], Hahamy A[1], Samborska V[2], Behrens T[1,2,3].

1. Wellcome Centre for Integrative Neuroimaging, University of Oxford, Oxford, UK.

2. Wellcome Centre for Human Neuroimaging, University College London, London, UK.

3. Sainsbury Wellcome Centre for Neural Circuits and Behaviour, University College London, London, UK

Abstract

Humans' ability for generalisation is outstanding. Our previous research has demonstrated that humans can generalise over tasks that share the same statistical rules (structural form) that govern the relational structure between states within a task. Studies suggest that the brain represents non-spatial tasks' structure similarly to the representation of spatial tasks and have highlighted the entorhinal cortex (EC) as a possible neural substrate for this structural generalisation. EC spatial representation generalises across environment size and shape. However, in non-spatial domains generalisation was only shown across tasks with the exact same underlying graphs. Here we show with fMRI that EC representations generalise across complex non-spatial tasks that share a hexagonal grid structural form but differ in their size and sensory stimuli, i.e. their only shared feature is the rule governing their statistical structure.

Introduction

If you grew up in a small town, arriving in a big city might come as a shock. However, you'll still be able to make use of your previous experiences, despite the difference in the size of the environment: When trying to navigate the busy city streets, your knowledge of navigation in your home town is crucial. For example, it's useful to know the constraints that a 2D topological structure exerted on distances between locations. When trying to make new friends, it's useful to remember how people in your home town tended to cluster in groups, with popular individuals perhaps belonging to several groups. Indeed, the statistical rules (termed "structural form", (Kemp and Tenenbaum 2008)) that govern the relationships between elements (states) in the environment are particularly useful for generalisation to novel situations, as they do not depend on the size, shape or sensory details of the environment (Mark *et al.* 2020). Such generalisable features of environments are proposed to be part of the "cognitive map" encoding the relationships between their elements (Tolman 1948; Behrens *et al.* 2018; Mark *et al.* 2020).

The most studied examples of such environments are spatial 2D tasks. In all spatial environments, regardless of their size or shape, the relations between states (in this case locations) are subject to the same Euclidean statistical constraints. The spatial example is particularly useful because neural spatial representations are well-characterised. Indeed, one of the most celebrated of these - grid cells in the entorhinal cortex (EC) - has been

suggested as (part of) a neural substrate for spatial generalisation (Behrens *et al.* 2018; Whittington *et al.* 2022). This is because (within a grid module) grid cells maintain their coactivation structure across different spatial environments (Fyhn *et al.* 2007; Yoon *et al.* 2013). In other words, the information embedded in grid cells generalises across 2D spatial environments (including environments of different shapes and sizes). Following a surge of studies showing that EC spatial coding principles are also used in non-spatial domains (Constantinescu, O'Reilly and Behrens 2016; Garvert, Dolan and Behrens 2017; Bao *et al.* 2019; Park *et al.* 2020), we have recently shown that EC also generalises over non-spatial environments that share the same statistical structure (Baram *et al.* 2021). Importantly, in that work the graphs that described the same-structure environments were isomorphic - i.e. there was a one-to-one mapping between states across same-structure environments.

What do we mean when we say the EC has "generalisable representations" in spatial tasks? and how can we probe these representations in complex non-spatial tasks? Between different spatial environments, each grid cell realigns: its firing fields might rotate and shift (Fyhn *et al.* 2007). Crucially, this realignment is synchronized within a grid module population (Yoon *et al.* 2013; Gardner *et al.* 2022), such that the change in the grid angle and phase of all cells is the same. This means that cells that have neighboring firing fields in one environment will also have neighboring firing fields in another environment- the coactivation structure is maintained (Yoon *et al.* 2013; Gardner *et al.* 2022). A mathematical corollary is that grid cells' activity lies in the same low-dimensional subspace (manifold, (Yoon *et al.* 2013; Gardner *et al.* 2022)) in all spatial environments. This subspace remains even during sleep, meaning the representation is stably encoded (Burak and Fiete 2009; Gardner *et al.* 2019).

We have recently developed an analysis method, referred to as "subspace generalisation", which allows for the quantification of the similarities between linear neural subspaces, and used it to probe generalisation in cell data (Samborska *et al.* 2022). Unlike other representational methods for quantifying the similarity between activity patterns (like RSA, used in Baram et al. (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte 2017)), this method has the ability to isolate the shared features underlying tasks that do not necessarily have a straightforward cross-task mapping between states, such as when the sizes of tasks underlying graphs are different. Here, we use it to quantify generalisation in such a case, but on fMRI data of humans solving complex abstract tasks rather than on cell data. We designed an abstract associative-learning task in which visual images were assigned to nodes on a graph and were presented sequentially, according to their relative ordering on the graph. The graphs belonged to two different families of graphs, each governed by a different set of statistical regularity rules (structural forms (Kemp and Tenenbaum 2008)) – hexagonal (triangular) lattice graphs, and community structure graphs. There were two graphs of each structural form. Crucially, the graph size and embedded images differed within a pair of graphs with the same structural form (fig 2b), allowing us to test generalisation due to structural form across both environment size and sensory information.

We first validate our approach by showing that subspace generalisation detects the known generalisation property of entorhinal grid cells when rodents free-forage in two different spatial environments – property that has inspired our study's hypothesis. We then validate the use of this method for fMRI signals in our task by showing it detects known properties of visual encoding in the visual cortex in our task. Finally, and most importantly, we show that EC generalises its voxelwise correlation patterns over hexagonal graphs of different size and stimuli, exactly as grid cells do in space. This result, however, did not hold for the community graph structures that were both non-spatial and non-euclidean. We discuss some possible experimental shortcomings that might have led to this null result.

Theory – "subspace generalisation"

How can we probe the neural correlates of generalisation of abstract tasks in the human brain? Popular representational analysis methods such as Representational Similarity Analysis (RSA) (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte 2017) and Repetition Suppression (Grill-Spector, Henson and Martin 2006; Barron, Garvert and Behrens 2016) have afforded some opportunities in this respect (Baram *et al.* 2021). However, because these methods rely on similarity measures between task states, they require labeling of a hypothesized similarity between each pair of states across tasks. Such labeling is not possible when we do not know which states in one task align with which states in another task. In the spatial example where states are locations, the mapping of each location in room A to locations in room B doesn't necessarily exist - particularly when the rooms differ in size or shape. This makes labeling of hypothesized similarity between each pair of locations impossible. How can we look for shared activity patterns in such a case?

We have recently proposed this can be achieved by studying the correlation of different neurons across states ((Samborska *et al.* 2022) (as opposed to RSA - which relies on the correlation of different states across neurons). If two tasks contain similar patterns of neural activity (regardless of when these occurred in each task), then the neuronXneuron correlation matrix (across time within-task) will look similar in both tasks. This correlation matrix can be summarised by its eigenvectors, which are patterns across neurons - akin to "cell assemblies" - and their eigenvalues, which indicate how much each pattern contributes to the overall variance in the data. If representations generalise across tasks, then patterns that explain a lot of variance in task 1 will also explain a lot of variance in task 2. We can compute the task 2 variance explained by each of the eigenvectors of task 1:

$$\Sigma_{12} = diag(U_1^T B_2 B_2^T U_1)$$

Where U_1 is a matrix with all task 1 eigenvectors as its columns, ordered by their eigenvalues, and B_2 is the neurons X states task 2 data. These eigenvectors are ordered according to the variance explained in task 1. Hence, if the same eigenvectors explain variance across tasks, early eigenvectors will explain more variance in task 2 than late

eigenvectors. The cumulative sum of Σ_{12} will be a concave function and the area under this concave function is a measure of how well neuronal patterns generalise across tasks (see results of grid cells data Fig1). We refer to this measure as subspace generalisation.

We first use this measure to illustrate the known generalisation property of grid cells in the rodent brain. We then test whether this measure can also be used in fMRI data to look for generalisation of the covariance between voxels rather than neurons. Finally, we use the approach to test whether voxel representations in EC generalise across complex non-spatial graphs with common regularities – analogous to the generalisation of grid cells in physical space.

Results

generalisation of grid cells

Grid cells generlise their representation across environments. Of relevance here: 1) Within an environment, grid cells representation lay on a low dimensional manifold (Yoon *et al.* 2013; Dordek *et al.* 2016; Gardner *et al.* 2022). 2) When a mouse moves from one environment to another the hexagonal grid activity pattern remains, grid cells do not remap**.** Further, the correlation structure between grid cells is preserved across environments, such that pairs of grid cells (within the same module) that have neighboring fields in environment 1 will also have neighboring fields in environment 2 (Fyhn *et al.* 2007). This is true even though each grid cell shifts and rotates its firing fields across environments - the grid cell population within a module realigns in unison (Gardner *et al.* 2022; Waaga *et al.* 2022). Crucially, the angle and phase of this realignment can't be predicted in advance, meaning it is not possible to create hypotheses to test regarding the similarity between representations at a given location in environment 1 and a given location in environment 2 - a requirement for methods such as RSA or repetition suppression.

We computed subspace generalisation for grid cells recorded in a previous study (Barry *et al.* 2012)), in which rats freely-foraged in two arenas, distinguished by their sensory properties. Crucially, we compared subspace generalisation between two different situations. One where "task 1" and "task 2" were actually from the same spatial environment (Fig1 - solid line) and one where "task 1" and "task 2" were from different environments (Fig1 - dot-dash line).

As predicted, within-environment, grid cell activity had a low dimensionality. That is – the first eigenvectors explain the majority of the activity variance (Fig 1). Across environments, grid cells' subspaces generalised: eigenvectors that were calculated using activity in one environment explained the variance in the other environment just as well as the within-environment baseline (Fig 1, compare solid and dot-dash lines, the difference between the area under the curve of the two line is significantly smaller than chance $p<0.001$ permutation test, see methods). In summary, our measure of subspace realignment

captured the predicted similarity in grid cell assemblies while animals forage in two different arenas.

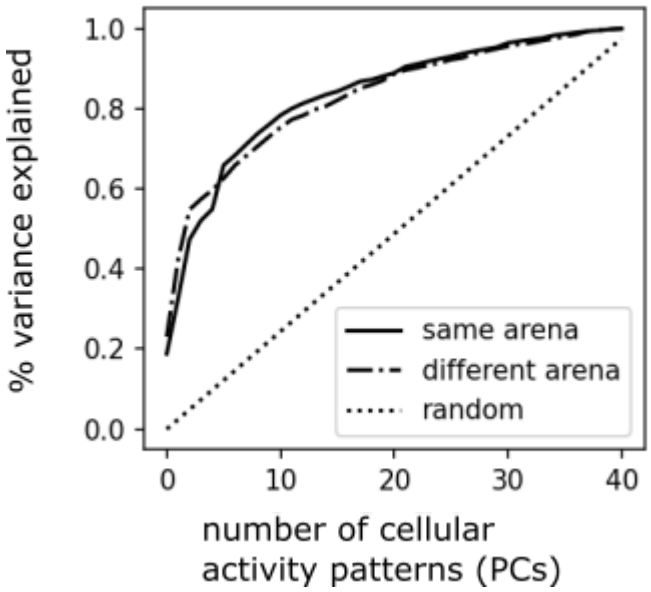

FIgure 1. The cumulative variance explained as a function of grid cells eigenvectors

Solid line - the cumulative activity variance explained while the animal free-forage in the same arena (different block, i.e cross validated), dot-dash - the cumulative activity variance explained while the animal forage in a different arena, dot - the cumulative variance plot that is expected by chance. It can be seen that the difference between the two arenas plots is very small, i.e grid cells activity patterns generalise over the different arenas ($p<0.001$ permutation test, see method).

Probing generalisation across abstract tasks with shared statistical rules - design and behavioural results

In human neuroimaging, the success of multivariate pattern analysis (MVPA, (Haxby *et al.* 2001)) and RSA (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte 2017)) tells us that, as with cells, the covariance between fMRI voxel activity contains information about the external world. It is therefore conceivable that we can measure the generalisation of fMRI patterns across related tasks using the same measure of subspace generalisation, but now applied to voxels rather than to cells. This will give us a measure of representational similarity in humans that can be used across tasks with no state-to-state mapping – e.g. when the size of the state space is different across tasks. In this section, we first describe the experimental paradigm we used to test whether, as in physical space, EC 1) generalises over abstract tasks governed by the same statistical rules; and 2) does so in a manner that is flexible to the size of the environment. In the next section we use the predicted properties

of visual encoding as a sanity check for the use of subspace generalisation in this task. Finally, we describe how the fMRI subspace generalisation results in EC depend on the statistical rules (structural forms) of tasks.

We designed an associative-learning task (Fig2, similar to the task in (Mark *et al.* 2020) fig1) where participants learned pairwise associations between images. The images can be thought of as nodes in a graph (unseen by participants), where the existence of an edge between nodes translates to an association between their corresponding images. There were two kinds of statistical regularities governing graph structures: a hexagonal/triangular structure form (which can be thought of as a discretized 2D structure) and a community structure. There were also two mutually exclusive image sets that could be used as nodes for a graph, meaning that each structural form had two different graphs with different image sets, resulting in a total of four graphs per participant. Importantly, two graphs of the same structural form were also of different sizes (36 and 42 nodes for the hexagonal structure; 35 and 42 nodes for the community structure - 5 or 6 communities of 7 nodes per community, respectively), meaning states could not be aligned even between graphs of the same structural form. The pairs of graphs with the (approximately) same sizes across structural forms used the same visual stimuli set (Fig 1). This design allowed us to test for subspace generalisation between tasks with the same underlying statistical regularities, controlling for the tasks stimuli and size.

Participants were trained on the graphs for four days and graph knowledge was assessed in each of the days using a battery of tests described previously. Some tests probed knowledge of pairwise (neighbouring) associations (Figure 2 C-D) and others probed “a sense of direction” in the graph, beyond the learned pairwise associations of neighboring nodes (Figure 2 E-F). In all tests, the performance of participants improved with learning and was significantly better than chance by the end of training (Fig 2 C-F), suggesting that participants were able to learn the graphs and developed a sense of direction even though they were never exposed to the graphs beyond pairwise neighbors. Note that while all participants performed well on tests of neighboring associations, the variance across participants for tests of non-neighboring nodes was high, with some participants performing almost perfectly and others close to chance (compare panels C-D to panels E-F).

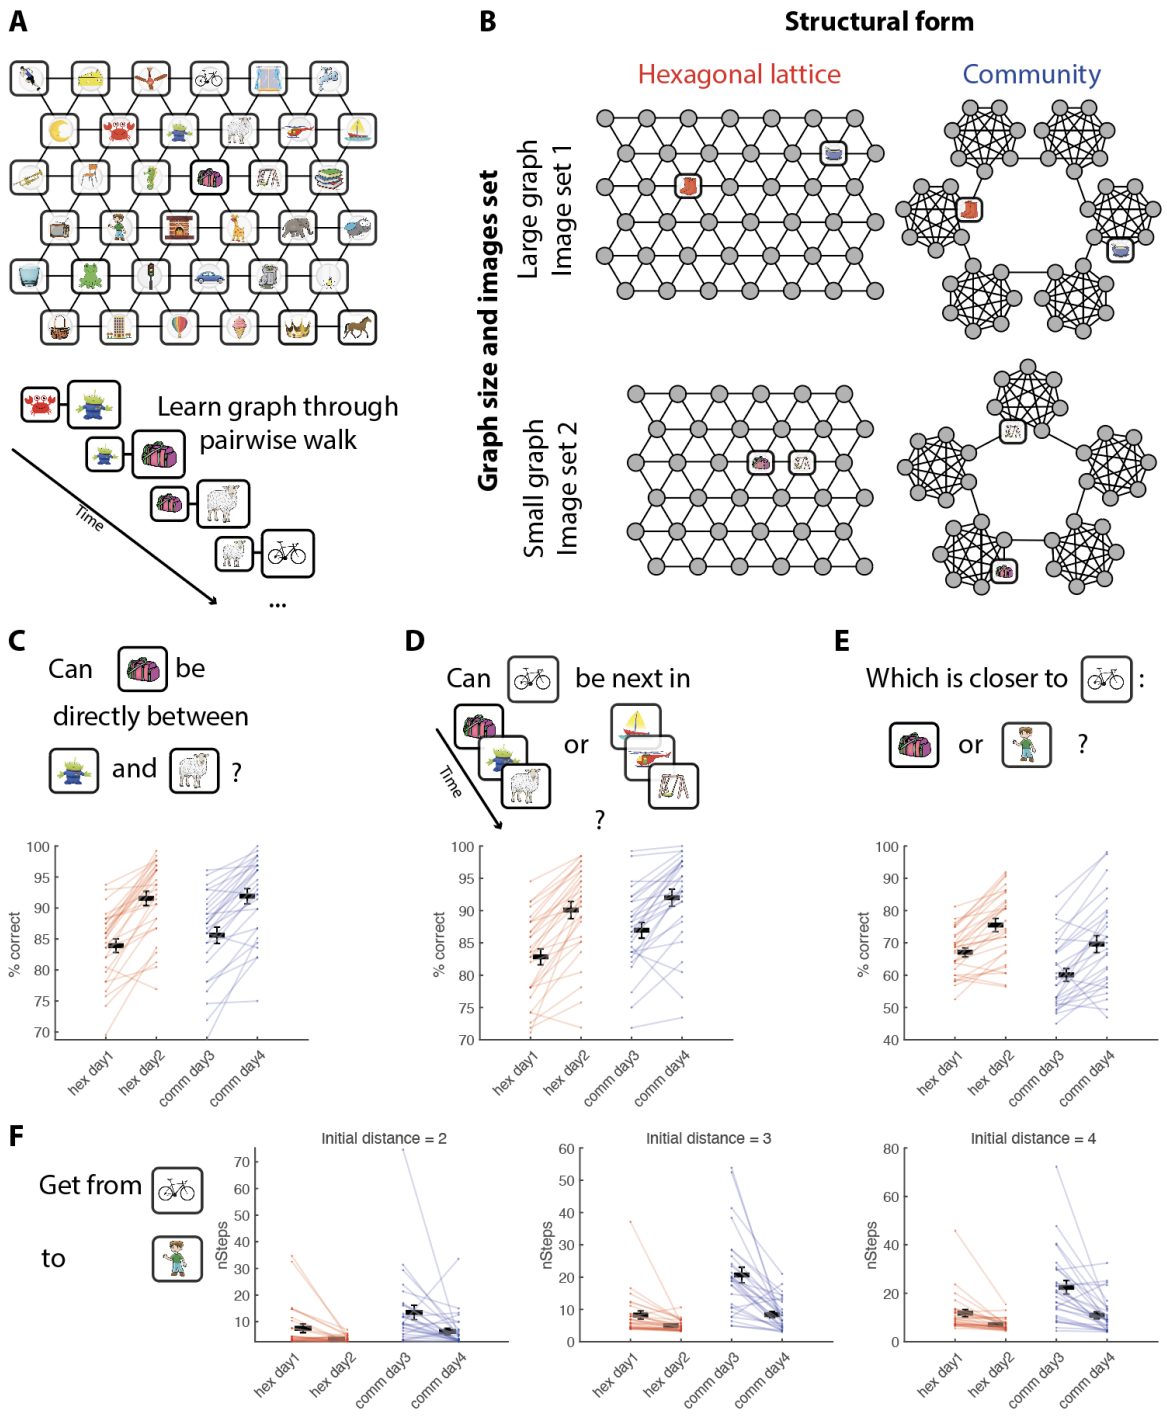

Figure 2. Experimental design and behavior. A. Example of an associative graph. Participants were never exposed to this top-down view of the graph - they learned the graph by viewing a series of pairs of neighboring images, corresponding to a walk on the graph. To aid memorisation, we asked participants to internally invent stories that connect the images. **B.** Each participant learned 4 graphs: two with an hexagonal lattice structure (both learned on days 1 and 2) and two with a community structure (both learned on days 3 and 4). For each structural form, there was one larger graph and one smaller graph. The nodes of graphs with

approximately the same size were drawn from the same set of images. **C-F**. In each day of training we used four tests to probe the knowledge of the graphs, as well as to promote further learning. In all tests, participants performed above chance level on all days and improved their performance between the first and second days of learning a graph. **C.** Participants were asked whether an image X can appear between images Y and Z (one sided t-test against chance level (50%): hex day1 $t(27) = 31.2$, $p < 10^{-22}$; hex day2 $t(27) = 35.5$, $p < 10^{-23}$; comm day3 $t(27) = 26.9$, $p < 10^{-20}$; comm day4 $t(27) = 34.2$, $p < 10^{-23}$; paired one sided t-test between first and second day for each structural form: hex $t(27) = 4.78$, $p < 10^{-5}$; comm $t(27) = 3.49$, $p < 10^{-3}$). **D.** Participants were shown two 3-long image sequences, and were asked whether a target image can be the fourth image in the first, second or both of the sequences (one sided t-test against chance level (33.33%): hex day1 $t(27) = 39.9$, $p < 10^{-25}$; hex day2 $t(27) = 42.3$, $p < 10^{-25}$; comm day3 $t(27) = 44.8$, $p < 10^{-26}$; comm day4 $t(27) = 44.2$, $p < 10^{-26}$; paired one sided t-test between first and second day for each structural form: hex $t(27) = 3.97$, $p < 10^{-3}$; comm $t(27) = 2.81$, $p < 10^{-2}$). **E.** Participants were asked whether an image X is closer to image Y or image Z (one sided t-test against chance level (50%): hex day1 $t(27) = 12.6$, $p < 10^{-12}$; hex day2 $t(27) = 12.5$, $p < 10^{-12}$; comm day3 $t(27) = 5.06$, $p < 10^{-4}$; comm day4 $t(27) = 7.42$, $p < 10^{-07}$; paired one sided t-test between first and second day for each structural form: hex $t(27) = 3.44$, $p < 10^{-3}$; comm $t(27) = 2.88$, $p < 10^{-2}$). **F.** Participants were asked to navigate from a start image X to a target image Y. In each step, the participant had to choose between two (randomly selected) neighbors of the current image. The participant repeatedly made these choices until they arrived at the target image (paired one sided t-test between number of steps taken to reach the target in first and second day for each structural form. Left: trials with initial distance of 2 edges between start and target images: hex $t(27) = 2.57$, $p < 10^{-2}$; comm $t(27) = 2.41$, $p < 10^{-2}$; MIddle: initial distance of 3 edges: hex $t(27) = 2.58$, $p < 10^{-2}$; comm $t(27) = 4.67$, $p < 10^{-2}$; Right: trials with initial distance of 4 edges: hex $t(27) = 3.02$, $p < 10^{-2}$; comm $t(27) = 3.69$, $p < 10^{-3}$). Note that while feedback was given for the local tests in panels C and D, no feedback was given for the tests in panels E-F to ensure that participants were not directly exposed to any non-local relations. The location of different options on the screen was randomised for all tests. of Hex: hexagonal lattice graphs. Comm: community structure graphs.

FMRI task and analysis

On the fifth day participants performed a task in the fMRI scanner. Each block of the scan included one of the four graphs the participant has learned, and started with a self-paced image-by-image random walk on the graph to allow inference of the currently relevant graph (Fig 3a). The second part of the block had two crucial differences. First, images were arranged into sequences of 3 images that were presented in rapid succession (470ms between successive images), corresponding to a walk of length 3 on the graph (Fig 3a). Once all three images appeared on the screen, they remained there for 470ms, and then the next

3-long image sequence started (Fig 3a – timings in scanner). Second, while the order within each 3-images sequence was dictated by the graph, the order across the sequences was pseudo-random. We needed this second manipulation to ensure coverage of the graph in every block and to eliminate the possibility of spurious temporal correlations between neighbouring sequences. However, if we had presented images individually in this random order, graphs with the same stimuli set would have been identical, making it difficult for subjects to maintain a representation of the current graph across the block. Whilst the images were the same across 2 graphs, the sequences of neighboring images uniquely identified each graph, inducing a sensation of "moving" through the graph. To encourage attention to the neighborhood of the sequence in the graph, in 12.5% of trials the sequence was followed by a single image, and participants had to indicate whether it was associated with the last image in the sequence (Fig 3a). In all analyses we used only the data from the second part.

To analyze this data, we used the subspace generalisation method as described for the rodent data, but replacing the firing of neurons at different spatial locations with the activity of fMRI voxels for different 3-images sequences. To do this, we first performed a voxelwise GLM where each regressor modeled all appearances of a particular 3-images sequence in a given run, together with several nuisance regressors (GLM1, see Methods). This gave us the activity of each voxel for each sequence. For each voxel, in each run, we extracted the 100 nearest voxels and formed a matrix of voxels X sequence. These are analogous to the data matrices, **B,** in equation 1. We then computed subspace generalisation using the eigenvectors of the voxel X voxel covariance matrix instead of the cell X cell covariance matrix (Fig 3b).

We then employed a leave-one-out cross-validation by repeatedly averaging the activation matrices from three runs of graph X, calculating the eigenvectors from this average representation, and then projecting the activation matrix of the held out run of control graph X (or a test graph Y) on these eigenvectors. This ensures that the "eigenvector" and "data" graphs are always from different runs. We then calculated the subspace generalisation between each pair of graphs resulting in a 4x4 matrix at each voxel of the brain (Fig 3b).

We refer to the elements of this 4x4 matrix in the following notation: we denote by H/C graphs of either hexagonal or community structure, and by s/l either small or large stimuli sets (matched across graphs of different structures). For example, HsCs denotes the element of the matrix corresponding to activity from the small hexagonal graph projected on eigenvectors calculated from the small (same image-set) community-structure graph.

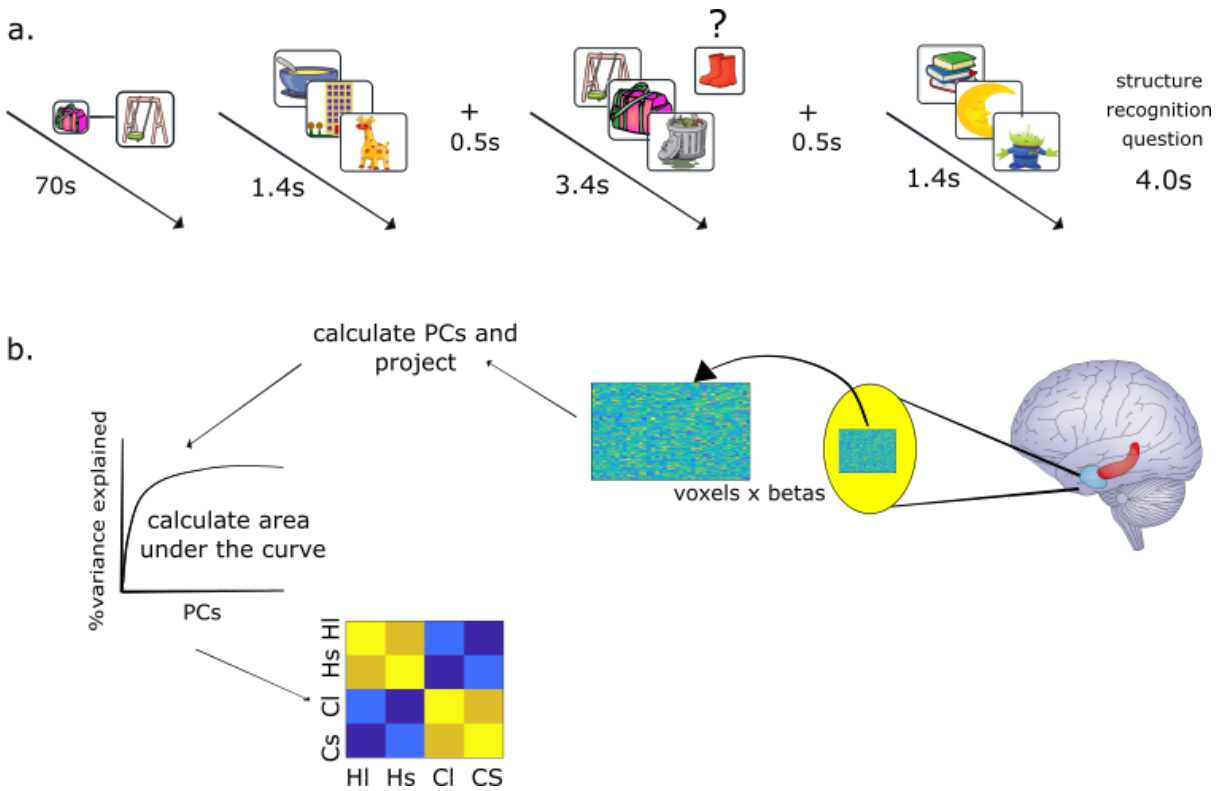

Figure 3. fMRI experiment and analysis method (subspace generalisation)

a. One block of the fMRI experiment. Each block starts with subjects observing a sequence of pictures from the graph at their own pace (on each step, two pictures appear on the screen). Then sequences of three pictures appear on the screen, while once a while there are catch trials in which we ask subjects to determine whether the questioned picture can appear next on the sequence.
b. Subspace generalisation method on fMRI voxels. Each searchlight extracts a voxels x betas matrix for each block. Then the PCs of each graph are calculated and the matrices of all other graphs are projected on the PCs. Following the projection we calculated the percentage of variance explained and the area under this curve. This leads to a 4 x 4 subspace generalization matrix (see main test and methods for more details).

Testing subspace generalisation on visual representations

To verify our analysis approach is indeed valid when used on our fMRI data, we first tested it on the heavily studied object encoding representations in lateral occipital cortex (LOC, (Malach *et al.* 1995; Grill-Spector, Henson and Martin 2006)). Recall that our stimuli in the scanner were concurrently presented sequences of three images of objects. We reasoned that these repeated sequences would induce correlations between object representations that should be observable in the fMRI data and detectable by our method. This would allow us to identify visual representations of the objects without ever specifying when the stimuli (i.e. 3-images sequences) were presented.

To this end we compared subspace generalization computed between different runs that included the same stimuli (3-images sequences, with different order across sequences between runs) with subspace generalization computed between runs of different stimuli while controlling for the graph structure. This led to the contrast [HlHl + ClCl + HsHs +CsCs] - [HlHs + HsHl + ClCs + CsCl], which had a significant effect in LOC (peak MNI [-44,-86,-8], t(27)_peak = 4.3, P_tfce = 0.002 based on a FWE-corrected nonparametric permutation test (Smith and Nichols 2009), corrected in bilateral LOC mask (Harvard-Oxford atlas, (Desikan *et al.* 2006)). In an additional exploratory analysis we tested the significance of the same contrast in a whole-brain searchlight. While this analysis did not reach significance once corrected for multiple comparisons, the strongest effect was found in LOC (fig. 4a). Note that in this contrast we intentionally ignored the elements of the 4x4 matrix where the data and the eigenvectors came from graphs with the same images set and a different structure (HlCl, HsCs, ClHl, CsHs), because they did not share the exact same visual stimuli (the 3-images sequence). In these cases we did not have a hypothesis about the subspace generalization in LOC. This suggests that we can detect the correlation structure induced by stimuli without specifying when each stimulus was presented.

EC generalizes a low-dimensional representation across discretized 2D graphs of different stimuli and sizes

Having established that the subspace generalization method can detect meaningful correlations between fMRI voxels, we next aimed to test whether EC will represent the statistical structure of abstract discretized 2D (hexagonal) graphs with generalisable low-dimensional representations. We therefore tested whether the EC subspaces from hexagonal graphs blocks were better aligned with the eigenvectors of other hexagonal blocks, than with the eigenvectors from community graphs blocks, i.e. ([HlHl + HlHs + HsHl + HsHs] - [HlCl + HlCs + HsCl + HsCs]). This contrast was significant in the right EC (peak MNI [28,-10,-40], t(27)_peak = 4.64, P_tfce = 0.01 based on a FWE-corrected nonparametric permutation test (Smith and Nichols 2009), corrected in a bilateral EC mask (Fig 4b) (Julich atlas,(Eickhoff *et al.* 2007)). An additional whole-brain exploratory analysis revealed that this right EC effect was in fact a brainwide maximum, but this effect did not survive correction for multiple comparisons in the whole brain.

This effect is actually made of two parts. The first is where the structures and stimuli are both aligned, compared to when only the stimuli are aligned ([HlHl + HsHs] - [HlCl + HsCs]). The second is where the structures are aligned, but the stimuli differ, compared to when both the structure and stimuli differ ([HlHs + HsHl] - [HlCs + HsCl]). To ensure robustness, we divided the analysis into these two parts and used one to define a functional ROI in which we tested the effect of the (orthogonal) other part. In each case, we took the peak voxel in the entorhinal mask from one contrast and tested the other (orthogonal) contrast at this voxel. Both effects were significant ($t(27)=4.01, p<0.001$ for [HlHs + HsHl] - [HlCs + HsCl] at peak of [HlHl + HsHs] - [HlCl + HsCs],[x =28, y = -10, ,z = -38], and $t=2.14$, $p<0.05$ vice versa, [x=34, y=-6, z=-40]).

Taken together, these results suggest that as in physical space, different abstract discretized 2D graphs are being represented on the same EC low-dimensional subspace. This is consistent with a view where the same EC cell assembly represents both hexagonal graphs, and that these cells covary together - even when the underlying size of the graph is different.

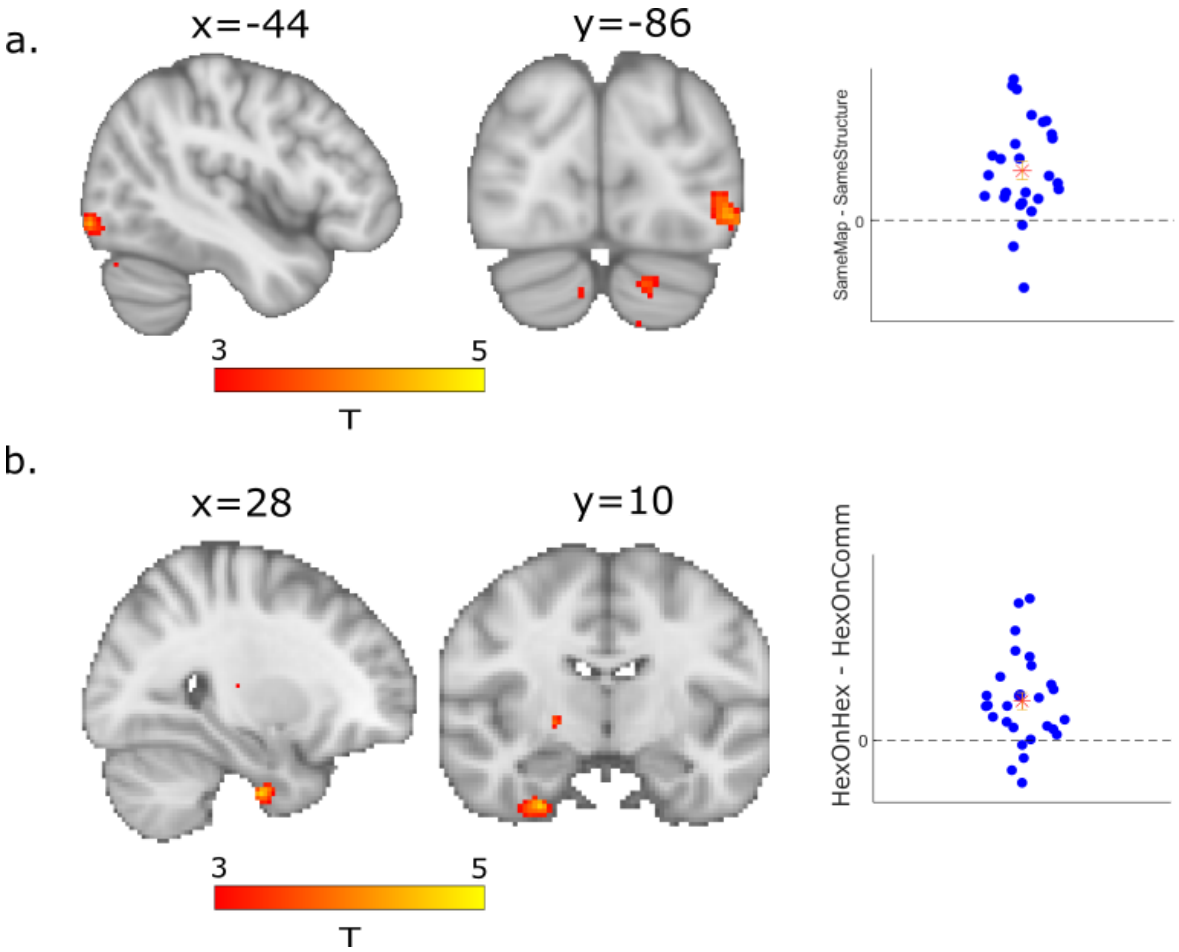

Figure 4: subspace generalisation in visual and structural representations.

A. Subspace generalisation of visual representations in LOC. Difference in subspace generalization was computed between different runs that included the same stimuli with subspace generalization computed between runs of different stimuli while controlling for the graph structure, i.e ([HlHl + ClCl + HsHs +CsCs] - [HlHs + HsHl + ClCs + CsCl]). t(27)_peak = 4.3, P_tfce = 0.002 over LOC.

B. EC generalises over the structure of Hexagonal graphs. The same voxels' subspaces represent both Hexagonal graphs even though the size of the graphs is different. The difference in subspace generalization was computed between subspace generalization of different Hexagonal graphs runs and subspace generalization of Hexagonal graph and Community graphs , i.e. ([HlHl + HlHs + HsHl + HsHs] - [HlCl + HlCs + HsCl + HsCs]). t(27)_peak = 4.64, P_tfce = 0.01 over EC.

Discussion

The contributions of this manuscript are two-fold: first, we show that EC representations generalize over hexagonal grid graphs of different sizes, highlighting the importance of the statistical properties of the environment to generalization. This expands our previous work (both experimental (Baram *et al.* 2021) and theoretical (Whittington *et al.* 2020)), suggesting EC plays an important role in generalization over abstract tasks, to the case where the tasks are governed by the same statistical rules but are not governed by the exact underlying graph (transition structure). This view builds on the known generalization properties of EC in physical space (Fyhn *et al.* 2007; Gardner *et al.* 2022) and on recent literature highlighting parallels between medial temporal lobe representations in spatial and non-spatial environments (Behrens *et al.* 2018; Whittington *et al.* 2022). Second, we present an fMRI analysis method ("subspace generalization"), adapted from related work in electrophysiology analysis (Samborska *et al.* 2022), to quantify generalization in cases where a mapping between states across environments is not available (though see (Hahamy and Behrens 2019) for our previous fMRI application of this method in the visual domain).

Exploiting previous knowledge while making decisions in new environments is a hard challenge that humans and animals face regularly. To enable generalization from loosely related previous experiences, knowledge should be represented in an abstract and flexible manner that does not depend on the particularities of the current task. Understanding the brain's solution to this computational problem requires a definition of a "generalisable representation", and a way of quantifying it. Here, we define generalization as sharing of cell assemblies, with similar activity patterns, across representations of related environments. Here, the assumption is that such cell assemblies encode the relevant information for a tasks' (environment) family, each belonging to a particular structural form.

An example of such generalization has previously been observed in the spatial domain, in grid cells recordings across different physical environments, regardless of shape or size (Fyhn *et al.* 2007; Gardner *et al.* 2022). This was usually done through direct comparison of the pairwise activity patterns of cells (Fyhn *et al.* 2007; Yoon *et al.* 2013; Gardner *et al.* 2022). However, this is not possible to do in fMRI, rendering the examination of EC generalization in complex abstract tasks difficult. "Subspace generalization" relies on the idea that similarity in activity patterns across tasks implies similarity of the within-task correlations between neurons. These are summarized in the similarity between the (low dimensional) linear subspaces where the activity of the neurons/voxels representing the two tasks lies. For fMRI purposes, this similarity between within-task neuronal correlations should be reflected in the similarity between within-task correlations across voxels, as long as the relevant neurons anatomically reside across a large enough number of voxels. Importantly, comparing similarity in neuronal correlations structures rather than similarity in states representations patterns (as in RSA (Kriegeskorte, Mur and Bandettini 2008)) allows us to examine flexible knowledge representations when a mapping between states in the two tasks does not exist. We present two validations of this method: in cells, we show it captures expected properties

of grids. In fMRI, we show it captures the expected correlations induced by the visual properties of a task in LOC.

Our main finding of subspace generalization in EC across discretized 2D graphs with different sizes and stimuli significantly strengthens the suggestion that EC flexibly represents all 'spatial- like' tasks such as 2D hexagonal graphs. Recently, we presented a theoretical framework for this idea: a neural network trained to predict future states, that when trained on 2D graphs displayed known spatial EC representations (the Tolman Eichenbaum Machine (TEM) (Whittington *et al.* 2020)). However, 'spatial-like' structures are not the only prevalent structure in natural tasks. The relations between task states often follow other structural forms (such as periodicities, hierarchies or community structures), inference of which can aid behavior (Mark *et al.* 2020). Representations of non-Euclidian task structures have been found in EC (Garvert, Dolan and Behrens 2017; Baram *et al.* 2021) and these generalize over different reinforcement learning tasks that are exactly the same except for their sensory properties (Baram *et al.* 2021). Indeed, when TEM was trained on non-Euclidean structures like hierarchical trees, it learned representations that were generalisable to novel environments with the same structure (Whittington *et al.* 2020). Further, we have previously shown that representing each family of graphs of the same structural form with the relevant stable representation (i.e, basis set) allows flexible transfer of the graph structure and therefore inference of unobserved transitions (relations between task's states) (Mark *et al.* 2020). Together these studies suggest that flexible representation of structural knowledge may be encoded in the EC.

Based on these, we hypothesized that EC representations will also generalize over non-'spatial-like' (here, community-structure) tasks of different sizes. We could not find conclusive evidence for such a representation: the relevant contrast ([ClCl + ClCs + CsCl + CsCs] - [ClHl + ClHs + CsHl + CsHs]) did not yield a statistically significant effect in EC or elsewhere (in an exploratory analysis corrected across the whole brain, though see below). This is despite clear behavioral evidence that participants use the community structure of the graph to inform their behaviour: participants have a strong tendency to choose to move to the connecting nodes over non-connecting nodes ((Mark *et al.* 2020), Fig 6). Moreover, in the post-experiment debriefing, participants could verbally describe the community structure of the graphs (26 out of 28 participants). This was not true for the hexagonal graphs. Why, then, did we not detect any generalization signals for the community structure graphs? There are both technical and psychological differences between the community structure and the hexagonal graphs that might have contributed to the difference in the results between the two structures. First, we have chosen a particular nested structure in which communities are organized on a ring. Subspace generalisation may not be suitable for the detection of community structure: for example, a useful generalisable representation of such structure is composed of a binary 'within-community nodes' vs 'connecting nodes' (nodes that connect two different communities) representation. If this is the representation used by in the brain, it means all "community-encoding" voxels are similarly active in response to all stimuli (as all 3-images sequences contain at least two non-connecting node images), and only "connecting nodes encoding" voxels change their activation during stimuli

presentation. Therefore, there is very little variance to detect. Second, there was high variation in participants' performance and in understanding of the ring structure, which may prevent us from detecting its representation.

Though this manuscript has focused on EC, it is worth noting that there is evidence for structural representations in other brain areas. Perhaps the most prominent of these is mPFC, where structural representations have been found in many contexts (Klein-Flügge *et al.* 2019; Baram *et al.* 2021; Klein-Flügge, Bongioanni and Rushworth 2022). Indeed, the strongest grid-like signals in abstract 2D tasks are often found in mPFC (Constantinescu, O'Reilly and Behrens 2016; Bao *et al.* 2019; Park *et al.* 2020; Bongioanni *et al.* 2021) and task structure representations have been suggested to reside in mOFC (Wilson *et al.* 2014; Schuck *et al.* 2016; Xie and Padoa-Schioppa 2016). The difference and interaction between PFC and MTL representations is a very active topic of research. One such suggested dissociation that might be of relevance here is the preferential contribution of MTL and PFC to latent and explicit learning, respectively. A related way of discussing this dissociation is to think of mPFC signals as closer to the deliberate actions subjects are taking . Circumstantial evidence from previous studies in our lab (tentatively) suggest the existence of such dissociation also for structural representations: when participants learnt a graph structure without any awareness of it, this structure was represented in MTL but not mPFC (Garvert, Dolan and Behrens 2017). On the other hand, when participants had to navigate on a 2D abstract graph to locations they were able to articulate, we observed much stronger grid-like signals in mPFC than MTL (though a signal in EC was also observed,(Constantinescu, O'Reilly and Behrens 2016)). In addition, Baram et al. found that while the abstract structure of a reinforcement learning task was represented in EC, the structure-informed learning signals that inform trial-by-trial behaviour with generalisable information were found in mPFC. Taken together, these results suggest that here, it is reasonable to expect generalisation signals of community structure graphs (of which participants were aware) in PFC, as well as the signals reported in EC for hexagonal graphs (of which participants were unaware). Indeed, when we tested for subspace generalisation of community structure graphs in the same region in vmPFC where Baram et al. found generalisable learning signals, we obtained a significant result (see supplementary).

To summarize, we have extended the understanding of EC representations and showed that EC represents Euclidian structures of different sizes, similarly to grid cells representation of spatial environments. We did this by using an analysis method which we believe will prove very useful for the study of generalisable representations in different neural recording modalities. More work is needed to verify whether this principle of EC representations extends to other, non-''spatial-like' structural forms.

Methods

Grid cells analysis:

Grid cells data were taken from (Barry *et al.* 2012). In short, grid cells were recorded while the animals forage in two different arenas; there were five different blocks such that during blocks 1&5 the animals were in one arena while in blocks 2-4 at the other. This experimental design results in five different firing rate maps, one for each block. After preprocessing (imputation of missing values, and smoothing of the 50x50 bins arenas firing rate maps), we calculated the 'subspace generalisation' score:

a. calculating the neuronXneuron correlation matrix from the first firing rate map (the neuronal activity from block 1) and its principal components (PCs).
b. project the firing rate maps from the other blocks on these PCs
c. calculated the cumulative variance explained as a function of PCs (that are organized according to their corresponding eigenvalues)
d. calculated the area under the curve (AUC).

*permutation tes*t: Our hypothesis is that the neuronXneuron correlation structure is preserved while the animals forage in the two different arenas, i.e that the active cells' assemblies remain the same. Therefore, the null hypothesis is that the cells' assemblies are random and do not remain the same while animals forage in the two arenas. We therefore permuted the cells' identity of block 1 firing rate maps and calculated the 'subspace generalsation' score between blocks 2-4 and the permuted block 1 PCs (calculated the PCs, projected the other blocks firing rate maps and calculated the AUC of the cumulative variance explained graph). We then subtracted the 'subspace generalsation' score (AUC) of the block5 on block 1 by these values. This creates our NULL distribution. We then show that the difference between the 'subspace generalsation' score of the block 5 on block 1 (same arena) and the 'subspace generalsation' score of the blocks 2-4 on block 1 (different arenas) is significantly smaller than randomly expected.

Participants: All participants are UCL students. We have recruited 60 participants. As the training is long and hard, we have recruited two participants to the training sessions, when both participants finished the training, the best participant was chosen to be scanned. We have scanned 34 participants and excluded 6 participants from the analysis because of severe movement or too long sleeping period within the scanner.
The study was approved by the University College London Research Ethics Committee (Project ID 11235/001). Participants gave written informed consent before the experiment.

fMRI training task

First training day was performed in the lab while the other three training days were performed from home.

Graphs' structures. One hexagonal graph consisted of 36 nodes and the other 42 nodes as shown in Fig. 1. One community structured graph consisted of five communities and the other six communities with seven nodes each. Within a community, each node was connected to all other nodes except for connecting nodes that were not connected to each other but were each connected to a connecting node of a neighboring community (Fig. 1). Therefore, all nodes had a degree of six, similarly to hexagonal graphs (the nodes on the hexagonal graphs border had degree less than six). Our community-structure graph had a hierarchical structure, wherein communities are organized on a ring.

Experimental procedures. In each of the training days, participants learned two graphs with the same underlying structure but different stimuli. During the first two days participants learned the hexagonal graphs, while during the third and fourth days participants learned the community structured graphs. During the fifth day, before the fMRI scan, participants learned all four graphs with two repetitions of the hexagonal graphs and one repetition of each community structured graph. Stimuli were selected randomly, for each participant, from a bank of stimuli (each pair of graphs, one hexagonal and one of a community structured graph shared the same bank). Each graph was learnt during four blocks (Fig. 1b; 4 blocks for graph 1 followed by 4 blocks for graph 2). Participants could take short resting breaks during the blocks. They were instructed to take a longer resting break after completing the four blocks of the first graph of each learning day.
Block structure. The structure of each block during each day is as follows: 1) Learning phase 2) Extending pictures sequences 3) Can it be in the middle 4) Navigation task 5) Distance estimation (see Figure 1). Next, we elaborate the various components of each block.
Learning phase: Participants learned associations between graph nodes by observing a sequence of pairs of pictures which were sampled from a random walk on the graph (successive pairs of pictures shared a common picture). Participants were instructed to 'say something in their head' in order to remember the associations. Hexagonal graphs included 120 steps of the random walk per block and community-structured graphs included 180 steps per block (we introduced more pictures in the community graph condition as random walks on such graphs result in high sampling of transitions within a certain community and low sampling of transitions between communities).
Extending pictures sequences: Given a target picture, which of two sequences of three pictures can be extended by that picture (a sequence can be extended by a picture only if it is a neighbor of the last picture in the sequence, the correct answer can be sequence 1/sequence 2/both sequences): Sixteen questions per block. (A picture could not appear twice in the same sequence, therefore, if the target picture is already in the sequence the correct answer was necessarily the other sequence).
Can it be in the middle: Determine whether a picture can appear between two other pictures, the answer is yes if and only if the picture is a neighbor of the two other pictures. Sixteen questions per block.
Navigation task: The aim—navigate to a target picture (appears at the right of the screen). The task was explained as a card game. Participants are informed that they currently have the picture that appears on the left of the screen. They were asked to choose between two pictures that are associated with that picture or skip and sample again two pictures that are associated with the current picture (skip is counted as a step). On each step participants were instructed to choose a picture that they think has a smaller number of steps to the target picture (according to their memory). Following their choice, the chosen picture appeared on the left and two new pictures, that correspond to states that are neighbors of the chosen picture, appear as new choices (Fig. 1b). Once a participant selects a neighbor of the target picture, the target picture itself can appear as a picture that can be chosen. The game terminated when either the target was reached or 200 steps were taken (without reaching the target). In the latter case a message 'too many steps' was displayed. On the first block, for each step, the number of links from the current picture to the target picture was shown on the screen. Participants played three games in each block. The starting distance (number of links) between the starting picture to the target was 2, 3 and 4.
Distance estimation: Which of two pictures has the smallest number of steps to a target picture: 45 questions per block (no feedback was given).

fMRI scanning task

The task consisted of four runs. Each run was divided into five blocks (four for each graph and one more repetition for one of the hexagonal graphs; the repetition was not used in the analysis). On each block participants observed pictures that belong to one of the graphs. A block started with one

minute in which participants observed a random walk on the graph at their own pace; two neighboring pictures appeared on the screen, when participants pressed 'enter' a new picture appeared on the screen (similar to the training learning phase). The new picture appeared in the middle of the screen and the old picture appeared on its left. Participants were instructed to infer which 'pictures set' they are currently observing. No information about the graph was given. Following this phase, three pictures' sequences appeared on the screen, one after the other. To keep participants engaged on the task, once in a while a fourth picture appeared and participants had to indicate whether this picture can appear next on the sequence ('catch trials'). Before starting the fMRI scan participants were asked whether they found any differences between the picture sets during the first two day and the last two days, most participants (26 out of 28) could indicate that there were groups of pictures during the last two days and that this is not the case during the first two days. Therefore, at the end of each block participants answered whether or not there are groups in the current picture set (participants that were not aware of the groups were asked whether this set belongs to the first two training days). Participants were given a bonus for answering correctly such that 100% correct results in a ten pounds bonus.

fMRI data acquisition

FMRI data was acquired on a 3T Siemens Prisma scanner using 32 channel head coil. Functional scans were collected using a T2*-weighted echo-planar imaging (EPI) sequence with a multi-band acceleration factor of 4 (TR = 1.450 s, TE = 35 ms, flip angle = 70 degrees, voxel resolution of 2x2x2mm). A field map with dual echo-time images (TE1 = 10ms, TE2 = 12.46ms, whole-brain coverage, voxel size 2x2x2mm) was acquired to correct for geometric distortions due to susceptibility-induced field inhomogeneities. Structural scans were acquired using a T1-weighted MPRAGE sequence with 1x1x1mm voxel resolution. We discarded the first six volumes to allow for scanner equilibration.

Pre-processing

Pre-processing was performed using tools from the fMRI Expert Analysis Tool (FEAT), part of FMRIB's Software Library (FSL (Jenkinson *et al.* 2012)). Data from each of the four scanner runs was preprocessed separately. Each run was aligned to the xxx (the .fsf says mean to highres) using the motion correction tool MCFLIRT. Brain extraction was performed using automated brain extraction tool BET (Smith, 2002). All data were high-pass temporally filtered with a cut-off of 100s. Registration of EPI images to high-resolution structural images and to standard (MNI) space was performed using FMRIB's Linear (FLIRT, (Jenkinson and Smith 2001; Jenkinson *et al.* 2002)). No spatial smoothing was performed during pre-processing (see below for different smoothing protocols for each analysis). Because of the notable breathing- and susceptibility-related artifacts in the entorhinal cortex, we pre-cleaned the data with FMRIB's ICA tool, FIX.

Univariate analysis

Due to incompatibility of FSL with the MATLAB RSA toolbox (Nili *et al.* 2014) used in subsequent analyses, we estimated all first-level GLMs and univariate group-level analyses using SPM12 (Wellcome Trust Centre for Neuroimaging, https://www.fil.ion.ucl.ac. uk/spm).

For estimating manifolds alignment, we constructed a GLM to estimate the activation as a result of each three pictures' sequence (a 'pile'). The GLM includes the following regressors: mean CSF

regressor and 6 motion parameters as nuisance regressors, bias term modeling the mean activity in each fMRI run, a regressor for the 'start' message (as a delta), a regressor for the self- paced random walk on each graph (a delta for each new picture that appears on the screen), a regressor for each pile in each graph (duration of a pile: 1.4sec), regressor for the catch onset (delta) and the pile that corresponds to the catch (pile duration). All regressors beside the 6 motion regressors and CSF regressor were convolved with the HRF.

Multivariate analysis

Quantifying manifold alignments:

We have calculated noised normalized GLM betas within each searchlight using the RSA toolbox. For each searchlight and each graph, we had a voxels by ten activation matrix ($B_{voxel \times pile}$) that describes the activation of a voxel as a result of a particular pile (three pictures' sequence). We exploited this matrix to quantify the manifold alignment within each searchlight.

To account for fMRI auto-correlation we used Leave One Out (LOO) approach; For each fMRI scanner run and graph, we calculated the mean activation matrix over the three others scanner runs ($\hat{B}^{\sim j}$). We then calculated the left Principles Component (PCs) of that matrix ($U^{\sim j}_{voxel \times voxel}$). To quantify the alignment, we projected the excluded scanner run graph activation matrix (B^j) of each graph on these PCs and calculated the accumulated variance explained as a function of PCs, normalized by the total variance of each graph within each run. Therefore, for each run and graph we have calculated:

$$P^{a,b} = U^{\sim j}_{a} \bullet B^{j}_{b}$$

$$M^{a,b}_{k} = \frac{\sum_{l=1}^{10} \left(P^{l,k}_{a,b}\right)^2}{S^j}$$

$$\Sigma^j = U^{j^T} B^{j^T} B^j U^j$$

$P^{a,b}$ is the projection matrix $voxel \times pile$ of graph *'b'* on the PCs of graph *'a'*, $M^{a,b}_{k}$ is the normalized variance explained on the *'k'* direction, S^j is the summation of the diagonal of Σ^j, the total variance as a result of the graph piles. We then calculated the cumulative variance explained over all *'k'* PCs directions. As a summary statistic we calculated the area under this curve. This gives us a 4x4 alignment matrix, such that each entry (*a,b*) in this matrix is a measure of the alignment of voxels patterns as a result of the two graphs *a&b*. We then calculated different contrasts over this matrix.

The searchlight results were normalized and smoothed with a kernel of 6mm FWHM using SPM12 before performing group level statistics.

References

Bao X, Gjorgieva E, Shanahan LK *et al.* Grid-like Neural Representations Support Olfactory Navigation of a Two-Dimensional Odor Space. *Neuron* 2019;**102**:1066-1075.e5.

Baram AB, Muller TH, Nili H *et al.* Entorhinal and ventromedial prefrontal cortices abstract and generalize the structure of reinforcement learning problems. *Neuron* 2021;**109**:713–23.

Barron HC, Garvert MM, Behrens TE. Repetition suppression: a means to index neural representations using BOLD? *Philos Trans R Soc B Biol Sci* 2016;**371**:20150355.
Barry C, Ginzberg LL, O'Keefe J *et al.* Grid cell firing patterns signal environmental novelty by expansion. *Proc Natl Acad Sci* 2012;**109**:17687–92.
Behrens TE, Muller TH, Whittington JC *et al.* What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* 2018;**100**:490–509.
Bongioanni A, Folloni D, Verhagen L *et al.* Activation and disruption of a neural mechanism for novel choice in monkeys. *Nature* 2021;**591**:270–4.
Burak Y, Fiete IR. Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol* 2009;**5**:e1000291.
Constantinescu AO, O'Reilly JX, Behrens TE. Organizing conceptual knowledge in humans with a gridlike code. *Science* 2016;**352**:1464–8.
Desikan RS, Ségonne F, Fischl B *et al.* An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 2006;**31**:968–80.
Diedrichsen J, Kriegeskorte N. Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. *PLOS Comput Biol* 2017;**13**:e1005508.
Dordek Y, Soudry D, Meir R *et al.* Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife* 2016;**5**:e10094.
Eickhoff SB, Paus T, Caspers S *et al.* Assignment of functional activations to probabilistic cytoarchitectonic areas revisited. *NeuroImage* 2007;**36**:511–21.
Fyhn M, Hafting T, Treves A *et al.* Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 2007;**446**:190–4.
Gardner RJ, Hermansen E, Pachitariu M *et al.* Toroidal topology of population activity in grid cells. *Nature* 2022;**602**:123–8.
Gardner RJ, Lu L, Wernle T *et al.* Correlation structure of grid cells is preserved during sleep. *Nat Neurosci* 2019;**22**:598–608.
Garvert MM, Dolan RJ, Behrens TE. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *elife* 2017;**6**:e17086.
Grill-Spector K, Henson R, Martin A. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci* 2006;**10**:14–23.
Hahamy A, Behrens TE. Measuring the spatial scale of brain representations. *2019 Conference on Cognitive Computational Neuroscience*. 2019, 2019–1174.
Haxby JV, Gobbini MI, Furey ML *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 2001;**293**:2425–30.
Jenkinson M, Bannister P, Brady M *et al.* Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage* 2002;**17**:825–41.
Jenkinson M, Beckmann CF, Behrens TEJ *et al.* FSL. *NeuroImage* 2012;**62**:782–90.
Jenkinson M, Smith S. A global optimisation method for robust affine registration of brain images. *Med Image Anal* 2001;**5**:143–56.
Kemp C, Tenenbaum JB. The discovery of structural form. *Proc Natl Acad Sci* 2008;**105**:10687–92.
Klein-Flügge MC, Bongioanni A, Rushworth MFS. Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* 2022;**110**:2743–70.
Klein-Flügge MC, Wittmann MK, Shpektor A *et al.* Multiple associative structures created by reinforcement and incidental statistical learning mechanisms. *Nat Commun* 2019;**10**:4835.
Kriegeskorte N, Mur M, Bandettini P. Representational similarity analysis - connecting the branches of systems neuroscience. *Front Syst Neurosci* 2008;**2**.
Malach R, Reppas JB, Benson RR *et al.* Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc Natl Acad Sci* 1995;**92**:8135–9.
Mark S, Moran R, Parr T *et al.* Transferring structural knowledge across cognitive maps in humans and

models. *Nat Commun* 2020;**11**:4783.
Nili H, Wingfield C, Walther A *et al.* A Toolbox for Representational Similarity Analysis. *PLOS Comput Biol* 2014;**10**:e1003553.
Park SA, Miller DS, Nili H *et al.* Map Making: Constructing, Combining, and Inferring on Abstract Cognitive Maps. *Neuron* 2020;**107**:1226-1238.e8.
Samborska V, Butler JL, Walton ME *et al.* Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nat Neurosci* 2022;**25**:1314–26.
Schuck NW, Cai MB, Wilson RC *et al.* Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron* 2016;**91**:1402–12.
Smith SM, Nichols TE. Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage* 2009;**44**:83–98.
Tolman EC. Cognitive maps in rats and men. *Psychol Rev* 1948;**55**:189–208.
Waaga T, Agmon H, Normand VA *et al.* Grid-cell modules remain coordinated when neural activity is dissociated from external sensory cues. *Neuron* 2022;**110**:1843-1856.e6.
Whittington JC, McCaffary D, Bakermans JJ *et al.* How to build a cognitive map. *Nat Neurosci* 2022;**25**:1257–72.
Whittington JC, Muller TH, Mark S *et al.* The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* 2020;**183**:1249–63.
Wilson RC, Takahashi YK, Schoenbaum G *et al.* Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron* 2014;**81**:267–79.
Xie J, Padoa-Schioppa C. Neuronal remapping and circuit persistence in economic decisions. *Nat Neurosci* 2016;**19**:855–61.
Yoon K, Buice MA, Barry C *et al.* Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* 2013;**16**:1077–84.