

1 **Flexible neural representations of abstract structural knowledge in the
2 human Entorhinal Cortex**

3 ***Mark S², Schwartenbeck P.^{1,2}, Hahamy A¹, Samborska V², **Baram AB.¹, **Behrens
4 TEJ^{1,2,3}.**

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

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

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

8 *:corresponding author

9 **: Equal contribution

10

11 **Abstract**

12 Humans' ability for generalisation is outstanding. It is flexible enough to identify cases where
13 knowledge from prior tasks is relevant, even when many features of the current task are different,
14 such as the sensory stimuli or the size of the task state space. We have previously shown that in
15 abstract tasks, humans can generalise knowledge in cases where the only cross-task shared feature
16 is the statistical rules that govern the task's state-state relationships. Here, we hypothesized that this
17 capacity is associated with generalisable representations in the entorhinal cortex (EC). This
18 hypothesis was based on the EC's generalisable representations in spatial tasks and recent
19 discoveries about its role in the representation of abstract tasks. We first develop an analysis
20 method capable of testing for such representations in fMRI data, explain why other common
21 methods would have failed for our task, and validate our method through a combination of
22 electrophysiological data analysis, simulations and fMRI sanity checks. We then show with fMRI that
23 EC representations generalise across complex non-spatial tasks that share a hexagonal grid
24 structural form but differ in their size and sensory stimuli, i.e. their only shared feature is the
25 rules governing their statistical structure. There was no clear evidence for such generalisation in EC
26 for non-spatial tasks with clustered, as opposed to planar, structure.

27

28 **Introduction**

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

42 The most studied examples of such environments are spatial 2D tasks. In all spatial
43 environments, regardless of their size or shape, the relations between states (in this case
44 locations) are subject to the same Euclidean statistical constraints. The spatial example is
45 particularly useful because neural spatial representations are well-characterised. Indeed,
46 one of the most celebrated of these - grid cells in the entorhinal cortex (EC) - has been
47 suggested as (part of) a neural substrate for spatial generalisation (Behrens *et al.* 2018;
48 Whittington *et al.* 2022). This is because (within a grid module) grid cells maintain their
49 coactivation structure across different spatial environments (Fyhn *et al.* 2007; Yoon *et al.*
50 2013). In other words, the information embedded in grid cells generalises across 2D spatial
51 environments (including environments of different shapes and sizes). Following a surge of
52 studies showing that EC spatial coding principles are also used in non-spatial domains
53 (Constantinescu, O'Reilly and Behrens 2016; Garvert, Dolan and Behrens 2017; Bao *et al.*
54 2019; Park *et al.* 2020), we have recently shown that EC also generalises over non-spatial
55 environments that share the same statistical structure (Baram *et al.* 2021). Importantly, in
56 that work the graphs that described the same-structured environments were isomorphic -
57 i.e. there was a one-to-one mapping between states across same-structure environments.

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

70 We have recently developed an analysis method, referred to as "subspace generalisation",
71 which allows for the quantification of the similarities between linear neural subspaces, and
72 used it to probe generalisation in cell data (Samborska *et al.* 2022). Unlike other
73 representational methods for quantifying the similarity between activity patterns (like RSA,
74 used in Baram *et al.* (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte
75 2017)), this method has the ability to isolate the shared features underlying tasks that do
76 not necessarily have a straightforward cross-task mapping between states, such as when the
77 sizes of tasks underlying graphs are different. Here, we use it to quantify generalisation in
78 such a case, but on fMRI data of humans solving complex abstract tasks rather than on cell
79 data. We designed an abstract associative-learning task in which visual images were
80 assigned to nodes on a graph and were presented sequentially, according to their relative
81 ordering on the graph. The graphs belonged to two different families of graphs, each
82 governed by a different set of statistical regularity rules (structural forms (Kemp and
83 Tenenbaum 2008)) – hexagonal (triangular) lattice graphs, and community structure graphs.

84 There were two graphs of each structural form. Crucially, the graph size and embedded
85 images differed within a pair of graphs with the same structural form (Figure 3b), allowing
86 us to test generalisation due to structural form across both environment size and sensory
87 information.

88 We first validate our approach by showing that subspace generalisation detects the known
89 generalisation properties of entorhinal grid cells and hippocampal place cells when rodents
90 free-forage in two different spatial environments – properties that have inspired our study's
91 hypothesis. Next, we propose that our method can capture these properties even in low-
92 resolution data such as fMRI. We provide twofold support for this conjecture: through
93 sampling and averaging of the rodent data to create low resolution version of the data, and
94 through simulations of grid cells grouped into simulated voxels to account for the very low
95 resolution of the BOLD signal. We use these simulations to discuss how the sensitivity of our
96 method depends on various characteristics of the signal. Next, we validate the method for
97 real fMRI signals by showing it detects known properties of visual encoding in the visual
98 cortex in our task. Finally, and most importantly, we show that EC generalises its voxelwise
99 covariance patterns over abstract, discrete hexagonal graphs of different size and stimuli,
100 exactly as grid cells do in space. This result, however, did not hold for the community graph
101 structures. We discuss some possible experimental shortcomings that might have led to this
102 null result.

103 **Theory – “subspace generalisation”**

104 How can we probe the neural correlates of generalisation of abstract tasks in the human
105 brain? Popular representational analysis methods such as Representational Similarity
106 Analysis (RSA) (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte 2017)
107 and Repetition Suppression (Grill-Spector, Henson and Martin 2006; Barron, Garvert and
108 Behrens 2016) have afforded some opportunities in this respect (Baram *et al.* 2021).

109 However, because these methods rely on similarity measures between task states, they
110 require labeling of a hypothesized similarity between each pair of states across tasks. Such
111 labeling is not possible when we do not know which states in one task align with which
112 states in another task. In the spatial example where states are locations, the mapping of
113 each location in room A to locations in room B doesn't necessarily exist - particularly when
114 the rooms differ in size or shape. This makes labeling of hypothesized similarity between
115 each pair of locations impossible. How can we look for shared activity patterns in such a
116 case?

117 We have recently proposed this can be achieved by studying the covariance of different
118 neurons across states (Samborska *et al.* 2022) (as opposed to RSA - which relies on the
119 correlation of different states across neurons). If two tasks contain similar patterns of neural
120 activity (regardless of when these occurred in each task), then the *neuron X neuron*
121 covariance matrix (across states within-task) will look similar in both tasks. This covariance
122 matrix can be summarised by its principal components (PCs), which are patterns across
123 neurons - akin to “cell assemblies” - and their eigenvalues, which indicate how much each
124 pattern contributes to the overall variance in the data. If representations generalise across

125 tasks, then patterns that explain a lot of variance in task 1 will also explain a lot of variance
126 in task 2. We can compute the task 2 variance explained by each of the PCs of task 1:

$$V_{12} = \text{diag}(PC_1^T A_2 A_2^T PC_1)$$

127 Where PC_1 is a matrix with all task 1 PCs as its columns, ordered by their eigenvalues, and A_2
128 is the *neurons X states* task 2 data. These PCs are ordered according to the variance
129 explained in task 1. Hence, if the same PCs explain variance across tasks, early PCs will
130 explain more variance in task 2 than late PCs. The cumulative sum of V_{12} will be a concave
131 function and the area under this concave function is a measure of how well neuronal
132 patterns generalise across tasks (Figure 1a). We refer to this measure as subspace
133 generalisation.

134 As validation and demonstration of our method, we first use it to recover differences in
135 generalisation between grid cells and place cells in the rodent brain that have been shown
136 previously with other methods. Next, we demonstrate the feasibility of our method in
137 capturing this difference in generalization properties even after we manipulate the data and
138 reduce its resolution. To complete the logical bridge from cells to voxels, we address the
139 limitation of this demonstration: the low number of cells recorded. We simulate voxels from
140 synthetic grid cells and show how our method's power depends on various characteristics of
141 the signal. These analyses show that theoretically (and under reasonable conditions) our
142 method could still detect medial temporal lobe generalisation properties in fMRI BOLD
143 signal. Finally, and most importantly, we use our method to analyse fMRI data, testing for
144 generalisation of the covariance between voxel representations in human EC across
145 complex non-spatial graphs with common regularities – analogous to the generalisation of
146 grid cells in physical space. Crucially, in this task other representational methods common in
147 fMRI analysis such as RSA or repetition suppression would not be applicable (due to lack of
148 one-to-one mapping between states across graphs), highlighting the usefulness of our
149 method.

150

151 **Results**

152 **Subspace generalization captures known generalisation properties of grid and place cells**

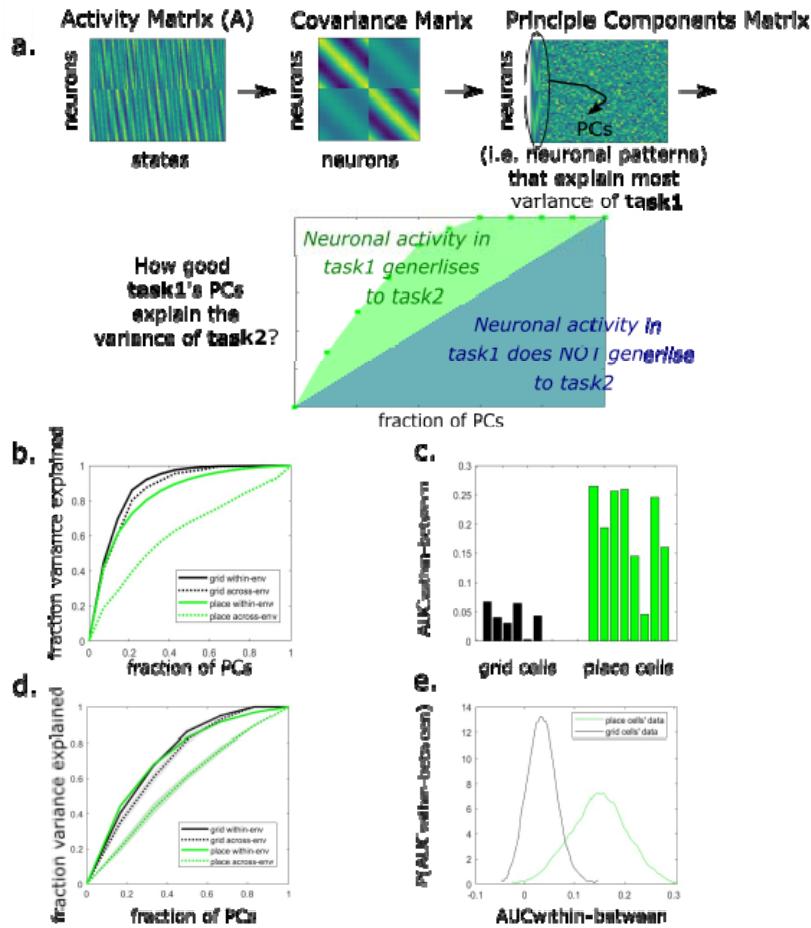
153 Grid cells and place cells differ in their generalisation property. When an animal moves from
154 one environment to another, place cells “remap”: they change their correlation structure
155 such that place cells that are neighbours in environment 1 need not be neighbors in
156 environment 2. By contrast grid cells do not remap: the correlation structure between grid
157 cells is preserved across environments, such that pairs of grid cells (within the same module)
158 that have neighboring fields in environment 1 will also have neighboring fields in
159 environment 2 (Fyhn *et al.* 2007). This is true even though each grid cell shifts and rotates
160 its firing fields across environments - the grid cell population within a module realigns in
161 unison (Gardner *et al.* 2022; Waaga *et al.* 2022). Crucially, the angle and phase of this
162 realignment can't be predicted in advance, meaning it is not possible to create hypotheses
163 to test regarding the similarity between representations at a given location in environment

164 1 and a given location in environment 2 - a requirement for fMRI-compatible methods such
165 as RSA or repetition suppression. In this section we demonstrate how subspace
166 generalisation - which can also be useful in fMRI - captures the generalisation properties of
167 grid and place cells that have previously been shown only with traditional analysis methods
168 that require access to firing maps of single cells.

169 We computed subspace generalisation for grid and place cells recorded with
170 electrophysiology in a previous study (Chen *et al.* 2018), in which mice freely-foraged in two
171 square environments: a real physical and a virtual reality (VR) (see Methods for more
172 details). For our purposes, this dataset is useful because large numbers of both place cells
173 and grid cells were recorded (concurrently within a cell type) in two different environments
174 - rather than because of the use of a VR environment.

175 We compared two different situations: one where “task 1” and “task 2” were actually from
176 the same environment, Figure 1a - solid line, within-environment) and one where “task 1”
177 and “task 2” were from different environments (Figure 1a - dotted line, across-
178 environments).

179 As predicted, across environments grid cells’ subspaces generalised: PCs that were
180 calculated using activity in one environment explained the activity variance in the other
181 environment just as well as the within-environment baseline (Figure 1a, compare dotted
182 and solid black lines, plots show the average of the projections of activity from one
183 environment on EVs from the other environment and vice versa). The difference between
184 the area under the curve (AUC) of the two lines was significantly smaller than chance
185 ($p < 0.001$ using a permutation test, see Methods and supplementary Figure S1). Importantly,
186 grid cells generalized much better between the environments than place cells; the
187 difference in AUCs between the solid and dotted lines is significantly smaller for grid cells
188 compared to place cells (Figure 1b, $p < 0.001$, for both permutation test and 2 sample t-test,
189 see Methods and supplementary material). Interestingly, the difference in AUCs was also
190 significantly smaller than chance for place cells (Figure 1a, compare dotted and solid green
191 lines, $p < 0.05$ using permutation tests, see statistics and further examples in supplementary
192 material Figure S2), consistent with recent models predicting hippocampal remapping that is
193 not fully random (Whittington *et al.* 2020).



194

195 **Figure 1. Subspace generalisation across environments in grid and place cells in data from**
196 **Chen et al. 2018.**

197 a. Illustration of the subspace generalization method. The Principal Components (PCs) are
198 calculated using the covariance matrix of the neuronal activity matrix. Then the activity
199 matrix is projected on each PC (recorded when the animal was in the same or different
200 environment/task) and the variance explained along each PC dimension is calculated.
201 We calculate the Area Under the Curve (AUC) of the cumulative sum of the variance
202 explained on each PC's dimension as our similarity measure. When the similarity in
203 neuronal patterns during the two different tasks is higher the area under the curve is
204 higher (green AUC is added to the blue AUC)

205 b. The cumulative variance explained by the PCs calculated using the activity of grid (black)
206 or place (green) cells, within (solid lines) and across (dotted lines) environments.
207 Subspace generalization is calculated as the difference between the area under the
208 curve (AUC) of two lines. The difference between the black lines is small, indicating
209 generalisation of grid cells across environments. The difference between the green lines
210 is larger, indicating remapping of place cells ($p < 0.001$, permutation test, see Methods).

211 c. The difference between the within and across (solid and dashed lines in a., respectively)
212 environments AUCs of the cumulative variance explained by grid or place cells (black or
213 green lines in a., respectively). Data shown for all mice with enough grid or place cells
214 (>10 recorded cells of the same type, each bar is a mouse and a specific projection (i.e.
215 projecting on environment one or two)). The differences between the grid cells AUCs are
216 significantly smaller than the place cells ($p < 0.001$ permutation test, see supplementary
217 for more statistical analyses and specific examples).

218 d. An example of the cumulative variance explained by the PCs, calculated using the
219 constructed low-resolution version of grid and place cells data. The solid and dotted
220 lines are average over 10 samples and the shaded areas represent the standard error of
221 the mean across samples. Here, as above, the solid lines are projection within
222 environment and the dotted lines are projections between environments.

223 e. Subspace generalization in the low resolution version of the data captures the same
224 generalization properties of grid vs place cells. The distributions were created via
225 bootstrapping over cells from the same animal, averaging their activity, concatenating
226 the samples across all animals and calculating the AUC difference between within and
227 across environments projections ($p < 0.001$ Kolmogorov-Smirnov test).

228

229

230 **From neurons to voxels**

231 So far, we have validated our method when applied to neurons. However, our primary
232 interest in this manuscript is to apply it to fMRI data. To illustrate the efficacy of this
233 approach in revealing generalisable neuronal subspaces within low resolution data like
234 fMRI, we applied our method to such data – both from manipulated electrophysiology and
235 simulations. We first examined our method on low-resolution versions of the Chen et al.
236 rodent MTL data, obtained by grouping and averaging cells. We show that our method can
237 still detect subspace generalization even on the supra-cellular level. However, due to the
238 small number of recorded cells, this analysis does not fully replicate a voxel's BOLD signal,
239 which corresponds to the average activity of thousands of cells. To address this, we
240 simulated many grid cells and grouped them into voxels, with each voxel's activity
241 corresponding to the average activity of its cells. We then applied subspace generalisation
242 to the simulated pseudo-voxels, and examined how the results depend on various signal
243 characteristics.

244 Using Chen et al electrophysiology dataset, we first normalised each cell's firing rate maps,
245 and then created bootstrapped low-resolution data: for each sampling iteration we sampled
246 7 cells (with repeats) into 2 groups within each animal and averaged the activities of cells
247 within each group. This results in a 2-long vector for each animal. We then concatenate
248 these vectors across animals. Note that for grid cells, this pooling over independent groups
249 of neurons is reminiscent of pooling over different grid modules in a single subject. For each
250 sample we calculated the difference in the area under the curve (AUC) between within and

251 across environments projections as above (averaged over the projections on both
252 environments, Figure 1c). We repeat this bootstrapping step to create a distribution of the
253 differences in AUC for place cells and grid cells (Figure 1d). The difference in AUC was
254 smaller for grid cells than for place cells ($p < 0.001$ Kolmogorov-Smirnov test), as is expected
255 from the single cells' analysis above.

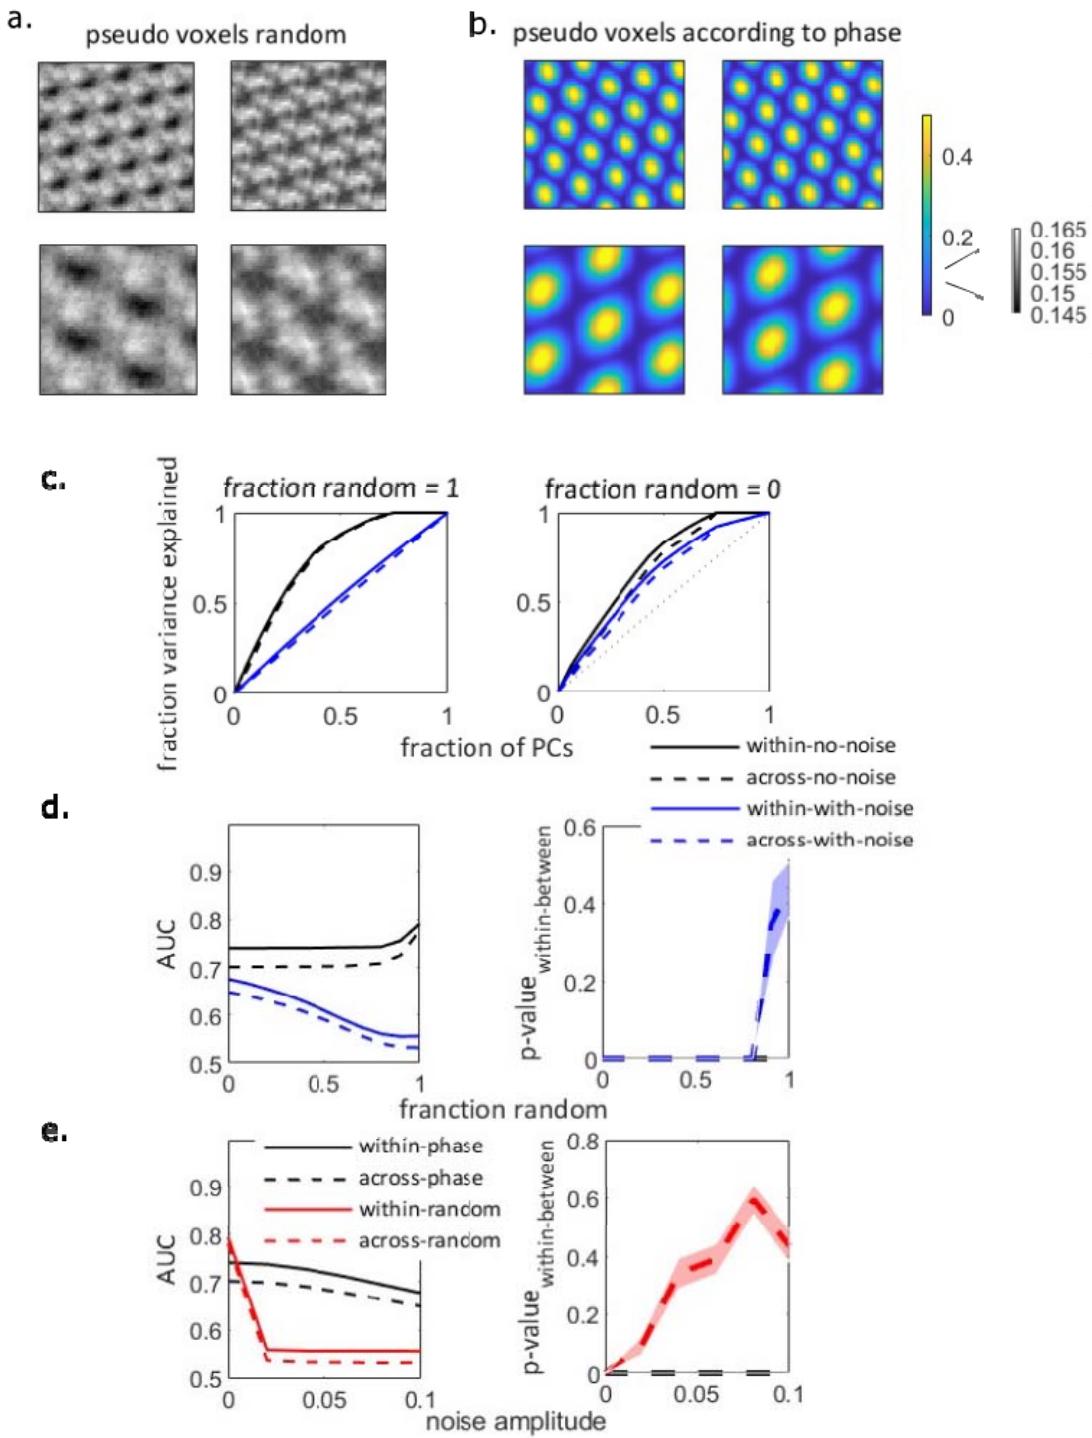
256

257 The required number of cells to simulate a voxel's activity (let alone multiple voxels) far
258 exceeds the number of cells in the Chen et al. dataset. To overcome this limitation and
259 support our conjecture that our method can detect subspace-generalization even in fMRI
260 BOLD signal, we next used simulated data. We simulated grid cells (see methods) organized
261 into four grid modules, each composed of more than 10000 cells. We organized the cells in
262 each module into four groups (pseudo-voxels) and averaged the activity within each group
263 (see supplementary info for an example of our analysis using different number of groups
264 within each module, and how our results are affected by the number of voxels per module,
265 Figure S3). We concatenated the pseudo-voxels from all modules into one vector and
266 calculated the difference in subspace-generalization measure (i.e. the AUC of within and
267 between environments). We explored how two characteristics of the data affect subspace
268 generalization: whether the grouping into voxels (within each module) was organized
269 according to grid phase, and the level of noise in the data.

270 We first grouped the cells into voxels randomly, i.e. without any a-priori assumption on the
271 relationship between the physical proximity of cells within the cortical layer and their firing
272 rate maps. Examples of the resulted "pseudo-voxels" activity maps can be seen in Figure 2a.
273 However, recent work has suggested there is a relationship between grid cells' physical
274 proximity and their grid phases (Gu et al. 2018). We therefore also simulated "pseudo
275 voxels" by grouping grid cells, within each module, according to their grid phase (Figure 2b).
276 The pseudo-voxel's signal in the latter case is substantially stronger (compare color bar
277 scales a between Figure 2a and 2b).

278 How does the difference between the signal variances affect the subspace generalization
279 measure? If the BOLD signal had no noise and all the cells within a voxel were indeed grid
280 cells, the actual variance of the signal would not affect our measure (Figure 2c, the solid and
281 dashed black lines are similar in both panels; i.e. the PCs that explain the activity variance
282 while the agent is in environment one explain the activity variance of environment two
283 similarly well, no matter how the cells are sampled into voxels). However, this is, of course,
284 unrealistic; the BOLD signal is noisy, and it is likely that voxel activity reflects non-grid cells
285 activity as well. To address this, we incorporated noise into our simulated voxel's activity
286 map. Figure 2c shows that increasing signal variance by grouping according to the grid
287 phase, leads to higher subspace generalization measure (AUC) compared to random
288 sampling; random sampling results in small AUC ($AUC \approx 0.5$) which is close to the expected
289 AUC following projections on random vectors (solid and dash blue lines in Figure 2c, left, see
290 supplementary info Figure S3 for further analysis). Predictably, as the fraction of randomly
291 sampled grid cells increases the ability to detect subspace generalization in the presence of

292 noise decreases (Figure 2d, Figure S3). Furthermore, sampling of grid cells according to
293 phase increases the statistical power of the subspace generalization method when the
294 amplitude of the noise increases (Figure 2e, Figure S3). To conclude, this shows under noisy
295 conditions, if nearby grid cells have similar phase tuning, as has been shown (Gu et al. 2018),
296 our method can in principle detect the generalization properties of grid cells, even in a very
297 low-resolution data, akin to the fMRI BOLD signal. It can in principle work to detect
298 generalization properties of any representation where nearby cells have similar tuning (such
299 as orientation tuning in V1).



300

301 *Figure 2: simulated voxels from simulated grid modules*

302

303 a. Examples of simulated voxels activity map in the two environments, without noise. upper:
304 higher frequency module, lower: lower frequency module. Cells are grouped into voxels
305 randomly.
306 b. Same as a. but with cells grouped into voxels according to the grid phase. Note the
307 different scale of the color-bar between a. and b.
308 c. Subspace generalization plot for the 16 simulated voxels, where the grouping into voxels
309 is either random (left) or according to phase (right). Legend as in d, noise amplitude = 0.1.
310 d. Left: AUCs of the subspace generalisation plots in c. as a function of the ratio of random
311 vs phase-organised cells in the voxels, with no noise (black) or with high amplitude of noise
312 (blue, *noise amplitude = 0.1*). Without noise (black lines), the subspace generalization
313 measure (AUC) remains high even when the fraction of randomly sampled cells increases.
314 However, in the presence of noise, the subspace generalization measure decreases with the
315 fraction of randomly sampled cells. Right: p-value of the effect according to the permutation
316 distribution (see methods, shaded area: standard error of the mean). In the presence of
317 noise and when the cells are sampled randomly, *AUCwithin-between* becomes non-significant,
318 see supplementary info Figure S3 for the dependency of the permutation distributions on
319 the presence of noise and sampling.
320 e. Same as d., except the continuous X-axis variable is the noise amplitude, for either of
321 phase-organized (black) or randomly organized voxels (red). AUC decreases sharply with
322 noise amplitude when the cells are sampled randomly, while it decreases more slowly when
323 the cells are sampled according to phase. The decrease in AUC to chance level (i.e. AUC =
324 0.5) with the increase in noise amplitude results in insignificant difference in subspace
325 generalization measure (*AUCwithin-between*). See supplementary info Figure S3 for the
326 permutation distributions.

327

328 **Probing generalisation across abstract tasks with shared statistical rules – task design and** 329 **behaviour**

330 In human neuroimaging, the success of multivariate pattern analysis (MVPA, (Haxby *et al.*
331 2001)) and RSA (Kriegeskorte, Mur and Bandettini 2008; Diedrichsen and Kriegeskorte
332 2017)) tells us that, as with cells, the covariance between fMRI voxel activity contains
333 information about the external world. It is therefore conceivable that we can measure the
334 generalisation of fMRI patterns across related tasks using the same measure of subspace
335 generalisation, but now applied to voxels rather than to cells. This will give us a measure of
336 generalisation in humans that can be used across tasks with no state-to-state mapping – e.g.
337 when the size of the state space is different across tasks. In this section, we first describe
338 the experimental paradigm we used to test whether, as in physical space, EC 1) generalises
339 over abstract tasks governed by the same statistical rules; and 2) does so in a manner that is
340 flexible to the size of the environment. In the next section we use known properties of visual
341 encoding as a sanity check for the use of subspace generalisation on fMRI data in this task.
342 Finally, we describe how the fMRI subspace generalisation results in EC depend on the
343 statistical rules (structural forms) of tasks.

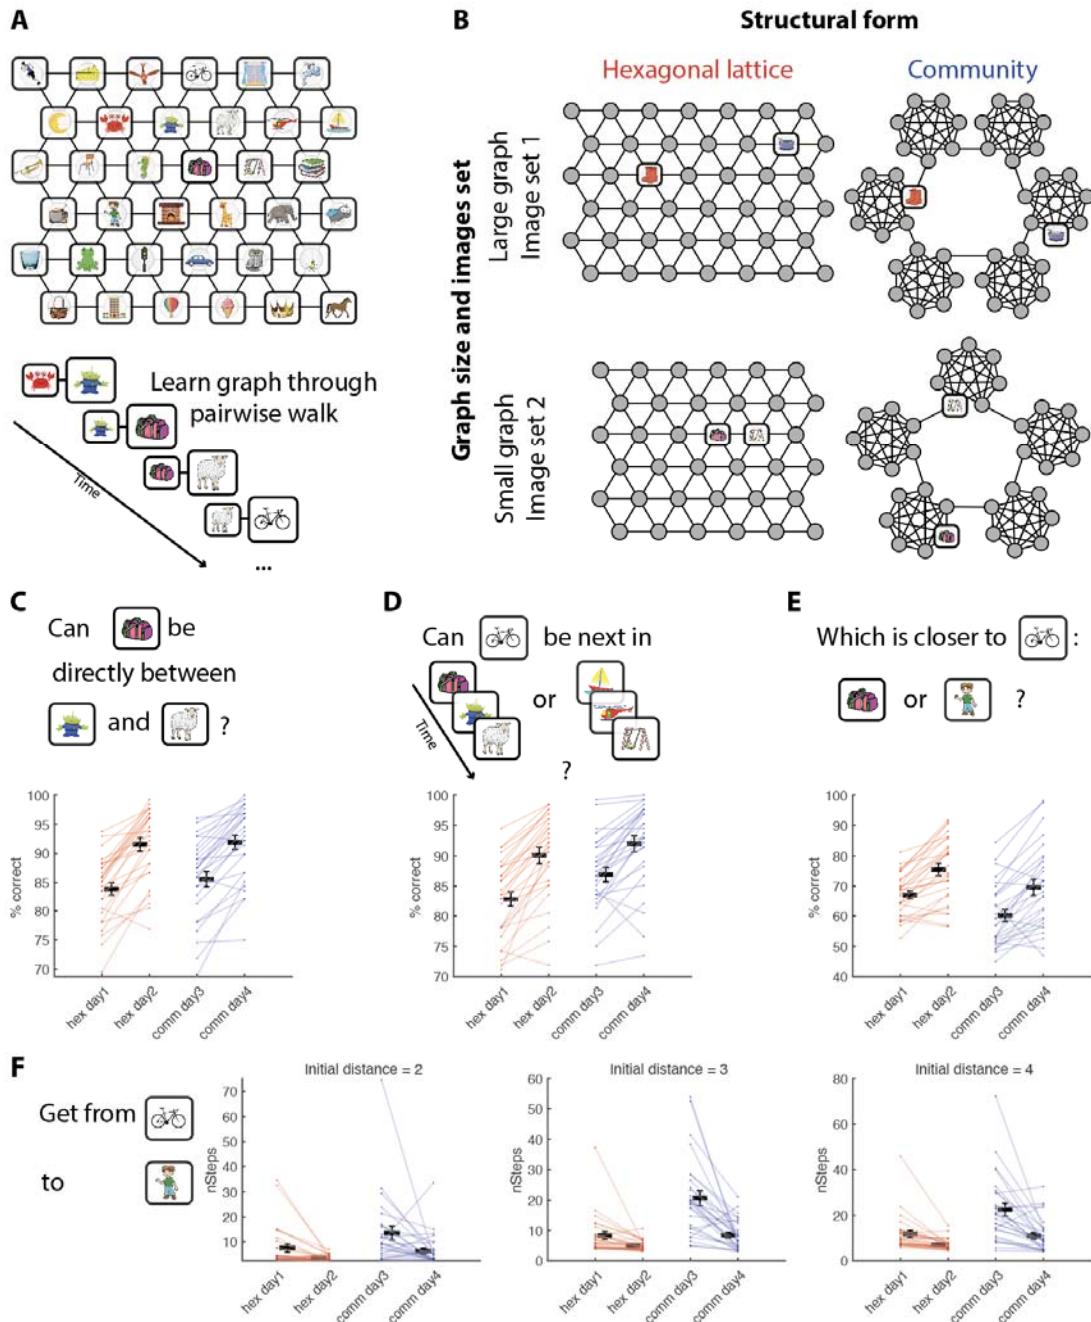
344 We designed an associative-learning task (Figure 3A and 3B, similar to the task in (Mark *et*
345 *al.* 2020)) where participants learned pairwise associations between images. The images can
346 be thought of as nodes in a graph (unseen by participants), where the existence of an edge

347 between nodes translates to an association between their corresponding images (Figure
348 3A). There were two kinds of statistical regularities governing graph structures: a
349 hexagonal/triangular structural form and a community structure. There were also two
350 mutually exclusive image sets that could be used as nodes for a graph, meaning that each
351 structural form had two different graphs with different image sets, resulting in a total of
352 four graphs per participant. Importantly, two graphs of the same structural form were also
353 of different sizes (36 and 42 nodes for the hexagonal structure; 35 and 42 nodes for the
354 community structure - 5 or 6 communities of 7 nodes per community, respectively),
355 meaning states could not be aligned even between graphs of the same structural form. The
356 pairs of graphs with the (approximately) same sizes across structural forms used the same
357 visual stimuli set (Figure 3B). This design allowed us to test for subspace generalisation
358 between tasks with the same underlying statistical regularities, controlling for the tasks'
359 stimuli and size.

360 Participants were trained on the graphs for four days and graph knowledge was assessed in
361 each of the days using a battery of tests described previously (Mark et al. 2020 and
362 methods). Some tests probed knowledge of pairwise (neighboring) associations (Figure 3C-
363 D) and others probed “a sense of direction” in the graph, beyond the learned pairwise
364 associations of neighboring nodes (Figure 3 E-F). In all tests, the performance of participants
365 improved with learning and was significantly better than chance by the end of training
366 (Figure 3 C-F), suggesting that participants were able to learn the graphs and developed a
367 sense of direction even though they were never exposed to the graphs beyond pairwise
368 neighbors. Note that while all participants performed well on tests of neighboring
369 associations, the variance across participants for tests of non-neighboring nodes was high,
370 with some participants performing almost perfectly and others close to chance (compare
371 panels C-D to panels E-F). At the end of the training days, we asked participants whether
372 they noticed how the images are associated with each other, 26 out of 28 participants
373 recognized that in two sets, the pictures were grouped.

374

375



376

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

385 **C-F.** In each day of training we used four tests to probe the knowledge of the graphs, as well
386 as to promote further learning. In all tests, participants performed above chance level on all
387 days and improved their performance between the first and second days of learning a
388 graph. **C.** Participants were asked whether an image X can appear between images Y and Z
389 (one sided t-test against chance level (50%): hex day1 $t(27) = 31.2$, $p < 10^{-22}$; hex day2
390 $t(27) = 35.5$, $p < 10^{-23}$; comm day3 $t(27) = 26.9$, $p < 10^{-20}$; comm day4 $t(27) = 34.2$, $p <$
391 10^{-23} ; paired one sided t-test between first and second day for each structural form: hex
392 $t(27) = 4.78$, $p < 10^{-5}$; comm $t(27) = 3.49$, $p < 10^{-3}$). **D.** Participants were shown two 3-
393 long image sequences, and were asked whether a target image can be the fourth image in
394 the first, second or both of the sequences (one sided t-test against chance level (33.33%):
395 hex day1 $t(27) = 39.9$, $p < 10^{-25}$; hex day2 $t(27) = 42.3$, $p < 10^{-25}$; comm day3 $t(27) =$
396 44.8 , $p < 10^{-26}$; comm day4 $t(27) = 44.2$, $p < 10^{-26}$; paired one sided t-test between
397 first and second day for each structural form: hex $t(27) = 3.97$, $p < 10^{-3}$; comm $t(27) = 2.81$
398 , $p < 10^{-2}$). **E.** Participants were asked whether an image X is closer to image Y or image Z, Y
399 and Z are not neighbors of X on the graph (one sided t-test against chance level (50%): hex
400 day1 $t(27) = 12.6$, $p < 10^{-12}$; hex day2 $t(27) = 12.5$, $p < 10^{-12}$; comm day3 $t(27) = 5.06$,
401 $p < 10^{-4}$; comm day4 $t(27) = 7.42$, $p < 10^{-07}$; paired one sided t-test between first and
402 second day for each structural form: hex $t(27) = 3.44$, $p < 10^{-3}$; comm $t(27) = 2.88$, $p <$
403 10^{-2}). **F.** Participants were asked to navigate from a start image X to a target image Y. In
404 each step, the participant had to choose between two (randomly selected) neighbors of the
405 current image. The participant repeatedly made these choices until they arrived at the
406 target image (paired one sided t-test between number of steps taken to reach the target in
407 first and second day for each structural form. Left: trials with initial distance of 2 edges
408 between start and target images: hex $t(27) = 2.57$, $p < 10^{-2}$; comm $t(27) = 2.41$, $p < 10^{-2}$;
409 Middle: initial distance of 3 edges: hex $t(27) = 2.58$, $p < 10^{-2}$; comm $t(27) = 4.67$, $p < 10^{-2}$ -
410 2; Right: trials with initial distance of 4 edges: hex $t(27) = 3.02$, $p < 10^{-2}$; comm $t(27) = 3.69$
411 , $p < 10^{-3}$). Note that while feedback was given for the local tests in panels C and D, no
412 feedback was given for the tests in panels E-F to ensure that participants were not directly
413 exposed to any non-local relations. The location of different options on the screen was
414 randomised for all tests. Hex: hexagonal lattice graphs. Comm: community structure graphs.

415

416 **FMRI task and analysis**

417 On the fifth day participants performed a task in the fMRI scanner. Each block of the scan
418 included one of the four graphs the participant has learned and started with a self-paced
419 image-by-image random walk on the graph to allow inference of the currently relevant
420 graph (Figure 4a, data not used in this manuscript). The second part of the block had two
421 crucial differences. First, images were arranged into sequences of 3 images that were
422 presented in rapid succession, corresponding to a walk of length 3 on the graph (Figure 4b
423 and Figure S5 for the partitioning the graphs into 3 images sequences). The time between

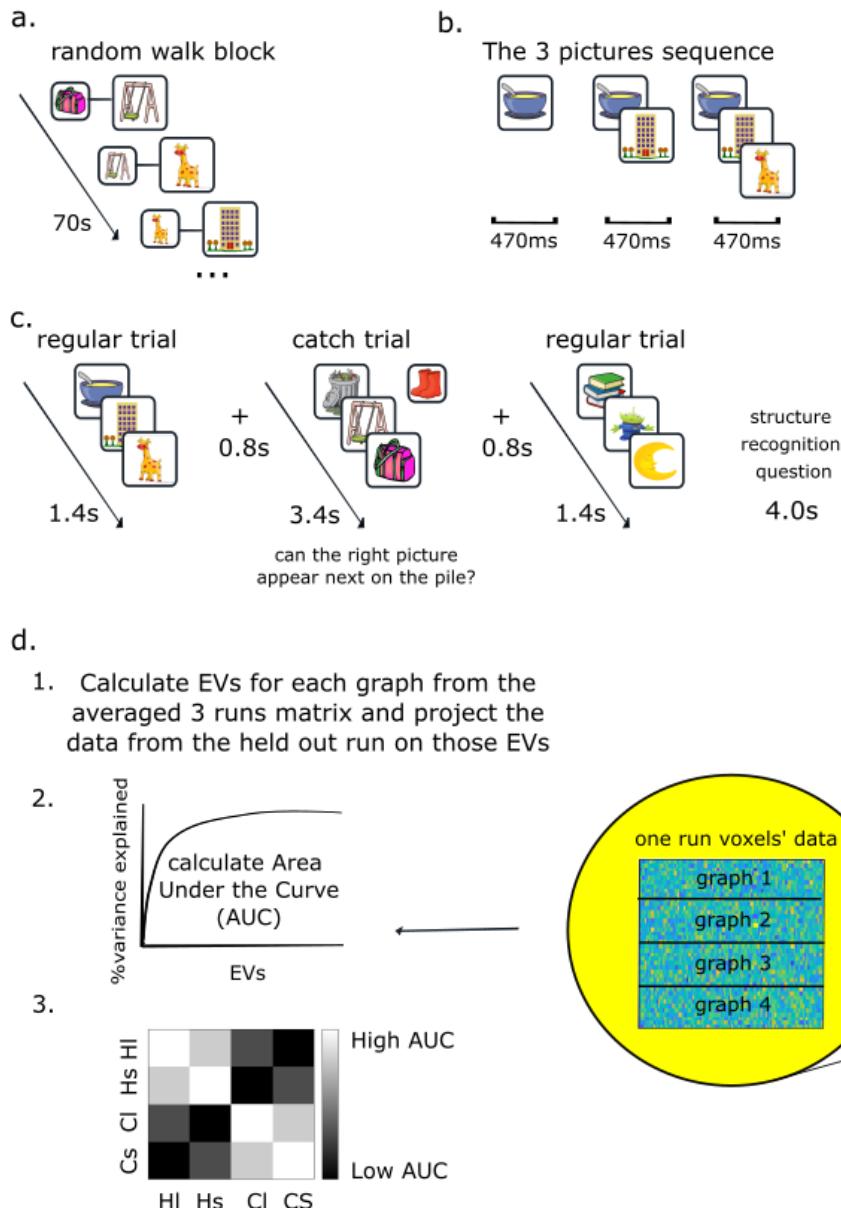
424 two successive sequences was 800ms (Figure 4c). Second, while the order within each 3-
425 images sequence was dictated by the graph, the order across the sequences was pseudo-
426 random. We needed this second manipulation to ensure coverage of the graph in every
427 block and to eliminate the possibility of spurious temporal correlations between
428 neighboring sequences. However, if we had presented images individually in this random
429 order, graphs with the same stimuli set would have been identical, making it difficult for
430 subjects to maintain a representation of the current graph across the block. Whilst the
431 images were the same across 2 graphs, the sequences of neighboring images uniquely
432 identified each graph, inducing a sensation of “moving” through the graph. To encourage
433 attention to the neighborhood of the sequence in the graph, in 12.5% of trials the sequence
434 was followed by a single image (“catch trial” in Figure 4c), and participants had to indicate
435 whether it was associated with the last image in the sequence (Figure 4c). Participants
436 answered these questions significantly better than chance (Figure S6), indicating that they
437 indeed recognize the correct graph and maintain the correct representation during the
438 block (t-test, $p < 0.001$ for both structures, $t[27]_{\text{hex}} = 11.3$, $t[27]_{\text{comm}} = 10.6$). At the end of
439 each block participants were asked whether they recognised which images set they
440 currently observed (see Method and supplementary for more details). Participants
441 answered these questions significantly better than chance (t-test, $p < 0.001$ for both
442 structures, $t[27]_{\text{hex}} = 3.8$, $t[27]_{\text{comm}} = 9.96$, see supplementary Figure S6), again indicating
443 that they correctly recognised the current graph in the scanner.

444 To analyze this data, we used the subspace generalisation method as described for the
445 rodent data but replacing the firing of neurons at different spatial locations with the activity
446 of fMRI voxels for different 3-images sequences. To do this, we first performed a voxelwise
447 GLM where each regressor modeled all appearances of a particular 3-images sequence in a
448 given run, together with several nuisance regressors (see Methods). This gave us the activity
449 of each voxel for each sequence. For each voxel, in each run, we extracted the 100 nearest
450 voxels and formed a matrix of sequence X voxels. These are analogous to the data matrices,
451 \mathbf{B} , in equation 1. We then computed subspace generalisation using the PCs of the voxel X
452 voxel covariance matrix instead of the cell X cell covariance matrix (Figure 4d).

453 We then employed a leave-one-out cross-validation by repeatedly averaging the activation
454 matrices from three runs of graph X, calculating the PCs from this average representation,
455 and then projecting the activation matrix of the held out run of control graph X (or a test
456 graph Y) on these PCs. This ensures that the “PC” and “data” graphs are always from
457 different runs. We then calculated the subspace generalisation between each pair of graphs
458 resulting in a 4x4 matrix at each voxel of the brain (Figure 4d).

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

464



465

466 **Figure 4. fMRI experiment and analysis method (subspace generalisation)**

467 a. Each fMRI block starts with 70s of random walk on the graph: a pair of pictures
468 appears on the screen, each time a participant presses enter a new picture appears
469 on the screen and the previous picture appears behind (similar to the three pictures
470 sequence, see below). During this phase participants are instructed to infer which
471 "pictures set" (i.e graph) they are currently playing with. Note that fMRI data from
472 this phase of the task is not included in the current manuscript.

473 b. The three pictures sequence: three pictures appear one after the other, while
474 previous picture/s still appear on the screen.

475 c. Each block starts with the random walk (panel a). Following the random walk,
476 sequences of three pictures appear on the screen. Every few sequences there was a
477 catch trial in which we ask participants to determine whether the questioned picture
478 can appear next on the sequence.

479 d. Subspace generalisation method on fMRI voxels. Each searchlight extracts a beta X
480 voxels' coefficients (of 3-images sequences) matrix for each graph in each run
481 THEREFORE, there are four such matrices). Then, using cross-validation across runs,
482 the left out run matrix of one graph is projected on the EVs from the (average of 3
483 runs of the) other graph. Following the projections, we calculate the cumulative
484 percentage of variance explained and the area under this curve for each pair of
485 graphs. This leads to a 4 X 4 subspace generalization matrix that is then being
486 averaged over the four runs (see main text and methods for more details). The colors
487 of this matrix indicate our original hypothesis for the study: that in EC, graphs with
488 the same structure would have larger (brighter) AUCs than graphs with different
489 structures (darker).

490

491 **Testing subspace generalisation on visual representations**

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

500 To this end we compared subspace generalization computed between different runs that
501 included the same stimuli (3-images sequences, with different order across sequences
502 between runs) with subspace generalization computed between runs of different stimuli
503 while controlling for the graph structure. This led to the contrast [H|H| + C|C| + Hs|Hs| + Cs|Cs|] -
504 [H|Hs| + Hs|H| + C|Cs| + Cs|C|], which had a significant effect in LOC (Figure 5a, peak MNI [-44,-
505 86,-8], $t(27)$ _peak = 4.96, $P_{tfce} < 0.05$ based on a FWE-corrected nonparametric
506 permutation test, corrected in bilateral LOC mask (Harvard-Oxford atlas, Desikan 2006,
507 Neuroimage). In an additional exploratory analysis, we tested the significance of the same
508 contrast in a whole-brain searchlight. While this analysis did not reach significance once
509 corrected for multiple comparisons, the strongest effect was found in LOC (Figure 5a). Note
510 that in this contrast we intentionally ignored the elements of the 4x4 matrix where the data
511 and the PCs came from graphs with the same images set and a different structure (H|C|,
512 Hs|Cs|, C|H|, Cs|Hs|), because they did not share the exact same visual stimuli (the 3-images
513 sequence). In these cases, we did not have a hypothesis about the subspace generalization
514 in LOC. These results suggest that we can detect the correlation structure induced by stimuli
515 without specifying when each stimulus was presented.

516 **EC generalizes a low-dimensional representation across hexagonal graphs of different**
517 **stimuli and sizes**

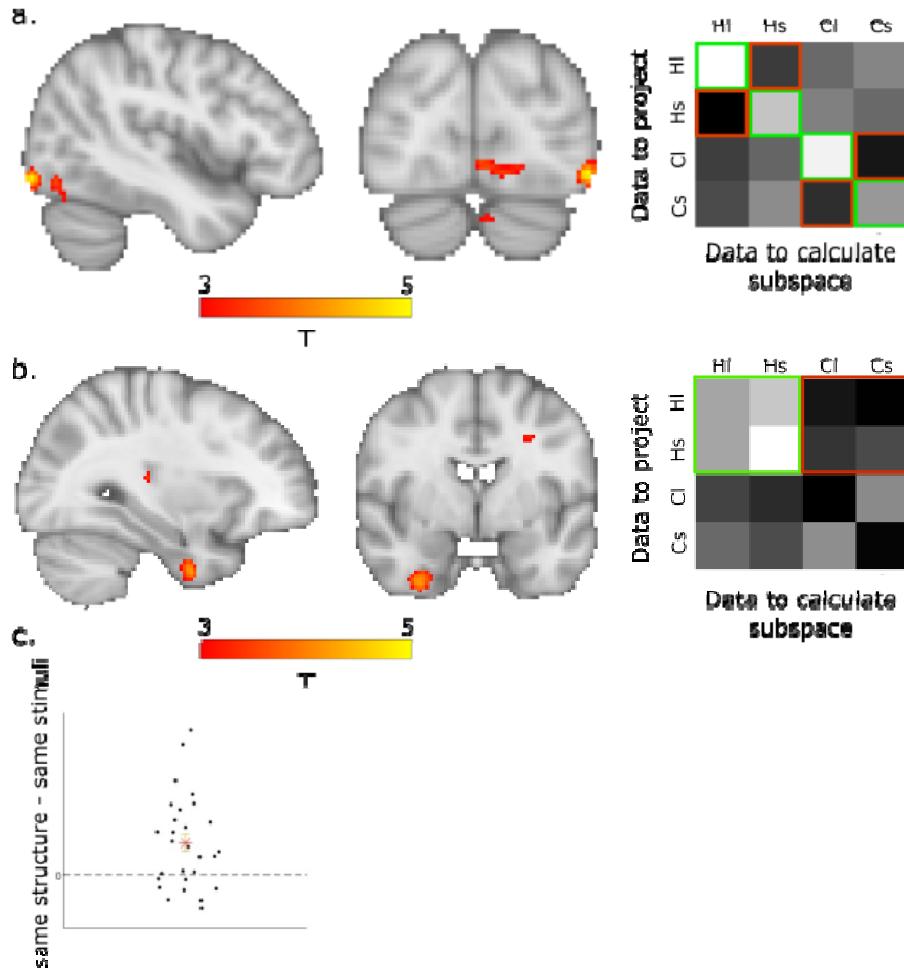
518 Having established that the subspace generalization method can detect meaningful
519 correlations between fMRI voxels, we next aimed to test whether EC will represent the
520 statistical structure of abstract graphs with generalisable low-dimensional representations.
521 We first tested this for discretized 2D (hexagonal) graphs, using the community structure
522 graphs as controls: We tested whether the EC subspaces from hexagonal graphs blocks were
523 better aligned with the PCs of other hexagonal blocks, than with the PCs from community
524 graphs blocks, i.e. $[[\text{HIHI} + \text{HIHs} + \text{HsHI} + \text{HsHs}] - [\text{HICl} + \text{HICs} + \text{HsCl} + \text{HsCs}]]$, Figure 5b). This
525 contrast was significant in the right EC (peak MNI [28, -10, -40], $t(27)$ _peak = 4.2, P_{tfce}
526 <0.01 based on a FWE-corrected nonparametric permutation test, corrected in a bilateral EC
527 mask (Figure 5b) (Julich atlas, Eickhoff 2007). We obtained a null result for the equivalent
528 analysis for community structure graphs $[[\text{ClCl} + \text{ClCs} + \text{CsCl} + \text{CsCs}] - [\text{ClHI} + \text{ClHs} + \text{CsHI} +$
529 $\text{CsHs}]]$. This was particularly due to low subspace generalization across different runs of the
530 same community structure graphs (bottom two diagonal elements in Figure 5b right,
531 compare to our original hypothesis subspace generalization matrix in Figure 4d). See the
532 Discussion for possible interpretations of this null result.

533 To ensure the robustness of the hexagonal graphs result we next tested the same effect in
534 an orthogonal ROI from our previous study. In (Baram *et al.* 2021) we have shown that EC
535 generalises over different reinforcement learning tasks with the exact same structure. We
536 therefore tested the same effect in that ROI (all voxels in the green cluster in Figure 3d in
537 Baram 2021 *et al.*, peak MNI: [25, -5, -28]), and indeed the $[[\text{HIHI} + \text{HIHs} + \text{HsHI} + \text{HsHs}] - [\text{HICl} + \text{HICs} + \text{HsCl} + \text{HsCs}]]$ contrast was significant (one sided t -test, $t(27) = 3.6$, $p < 0.001$, Figure 5c).

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

544

545



546

547 **Figure 5: subspace generalisation in visual and structural representations.**

548 a. Subspace generalisation of visual representations in LOC. Left: difference in subspace
549 generalization was computed between different blocks that included the same stimuli with
550 subspace generalization computed between blocks of different stimuli while controlling for
551 the graph structure, i.e. $[HIHI + CICl + HsHs + CsCs] - [HlHs + HsHI + CICs + CsCl]$. $t(27)_{peak} =$
552 4.96 , $P_{tfce} < 0.05$ over LOC. Right: visualization of the subspace generalisation matrix
553 (averaged over all LOC voxels with $t > 2$ for the $[HIHI + CICl + HsHs + CsCs] - [HlHs + HsHI + CICs$
554 $+ CsCl]$ contrast, i.e. green minus red entries.

555 b. EC generalises over the structure of hexagonal graphs. Left: the effect for the contrast
556 $[HIHI + HlHs + HsHI + HsHs] - [HlCl + HlCs + HsCl + HsCs]$, i.e. the difference between
557 subspace generalisation of hexagonal graphs data, when projected on PCs calculated from
558 (cross-validated) hexagonal graphs (green elements in right panel) vs community structure
559 graphs (red elements). $t(27)_{peak} = 4.2$, $P_{tfce} < 0.01$ over EC. Right: Same as in a. right but
560 for the $[HIHI + HlHs + HsHI + HsHs] - [HlCl + HlCs + HsCl + HsCs]$ contrast in EC.

561 c. The average effect in an ROI from Baram et al. (green cluster in figure 3d of Baram et al.)
562 for each participant. Star denotes the mean, error bars are SEM.

563 **Discussion**

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

578 Exploiting previous knowledge while making decisions in new environments is a hard
579 challenge that humans and animals face regularly. To enable generalization from loosely
580 related previous experiences, knowledge should be represented in an abstract and flexible
581 manner that does not depend on the particularities of the current task. Understanding the
582 brain’s solution to this computational problem requires a definition of a “generalisable
583 representation”, and a way of quantifying it. Here, we define generalization as sharing of
584 neuronal manifold across representations of related tasks. The particular assumption here is
585 that in the EC, such manifolds encode the relevant information about the particular
586 structural form of the task.

587 An example of such generalization has previously been observed in the spatial domain, in
588 grid cells recordings across different physical environments, regardless of shape or size
589 (Fyhn *et al.* 2007; Gardner *et al.* 2022). This was usually done through direct comparison of
590 the pairwise activity patterns of cells (Fyhn *et al.* 2007; Yoon *et al.* 2013; Gardner *et al.*
591 2022). However, this is not possible to do in fMRI, rendering the examination of EC
592 generalization in complex abstract tasks difficult. “Subspace generalization” relies on the
593 idea that similarity in activity patterns across tasks implies similarity of the within-task
594 correlations between neurons. These are summarized in the similarity between the (low
595 dimensional) linear subspaces where the activity of the neurons/voxels representing the
596 two tasks lies. For fMRI purposes, this similarity between within-task neuronal correlations
597 should be reflected in the similarity between within-task correlations across voxels, as long
598 as the relevant neurons anatomically reside across a large enough number of voxels.
599 Importantly, comparing similarity in neuronal correlations structures rather than similarity
600 in states representations patterns (as in RSA) allows us to examine flexible knowledge
601 representations when a mapping between states in the two tasks does not exist. We
602 present three validations of this method: in cells, we show it captures all expected
603 properties of grid and place cells, even if we reduce the data resolution by averaging over
604 the activity of group of cells. In simulation, we show that calculating subspace generalization

605 using simulated voxels from simulated grid cells results in significant generalization effect
606 under realistic condition. In fMRI, we show it captures the expected correlations induced by
607 the visual properties of a task in LOC.

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

627 Based on these, we hypothesized that EC representations will also generalize over non-
628 'spatial-like' tasks (here, community-structure) of different sizes. However, we could not
629 find conclusive evidence for such a representation: the relevant contrast ([CICl + CICs + CsCl
630 + CsCs] - [CIHI + CIHs + CsHI + CsHs]) did not yield a statistically significant effect in EC (or
631 elsewhere, in an exploratory analysis corrected across the whole brain). This is despite clear
632 behavioral evidence that participants use the community structure of the graph to inform
633 their behavior: participants have a strong tendency to choose to move to the connecting
634 nodes (nodes that connect two different communities) over non-connecting nodes ((Mark *et*
635 *al.* 2020), and Figure S4a). Moreover, in the post-experiment debriefing, participants could
636 verbally describe the community structure of the graphs (26 out of 28 participants). This was
637 not true for the hexagonal graphs. Why, then, did we not detect any neural generalization
638 signals for the community structure graphs? There are both technical and psychological
639 differences between the community structure and the hexagonal graphs that might have
640 contributed to the difference in the results between the two structures. First, we have
641 chosen a particular nested structure in which communities are organized on a ring.
642 Subspace generalisation may not be suitable for the detection of community structure: for
643 example, a useful generalisable representation of such structure is composed of a binary
644 'within-community nodes' vs 'connecting nodes' representation. If this is the representation
645 used by the brain, it means all "community-encoding" voxels are similarly active in response
646 to all stimuli (as all 3-images sequences contain at least two non-connecting node images),

647 and only “connecting nodes encoding” voxels change their activation during stimuli
648 presentation. Therefore, there is very little variance to detect.

649 Though this manuscript has focused on EC, it is worth noting that there is evidence for
650 structural representations in other brain areas. Perhaps the most prominent of these is
651 mPFC, where structural representations have been found in many contexts (Klein-Flügge *et*
652 *al.* 2019; Baram *et al.* 2021; Klein-Flügge, Bongioanni and Rushworth 2022). Indeed, the
653 strongest grid-like signals in abstract 2D tasks are often found in mPFC (Constantinescu,
654 O'Reilly and Behrens 2016; Bao *et al.* 2019; Park *et al.* 2020; Bongioanni *et al.* 2021) and task
655 structure representations have been suggested to reside in mOFC (Wilson *et al.* 2014;
656 Schuck *et al.* 2016; Xie and Padoa-Schioppa 2016). The difference and interaction between
657 PFC and MTL representations is a very active topic of research. One such suggested
658 dissociation that might be of relevance here is the preferential contribution of MTL and PFC
659 to latent and explicit learning, respectively. A related way of discussing this dissociation is to
660 think of mPFC signals as closer to the deliberate actions subjects are taking. Circumstantial
661 evidence from previous studies in our lab (tentatively) suggest the existence of such
662 dissociation also for structural representations: when participants learnt a graph structure
663 without any awareness of it, this structure was represented in MTL but not mPFC (Garvert,
664 Dolan and Behrens 2017). On the other hand, when participants had to navigate on a 2D
665 abstract graph to locations they were able to articulate, we observed much stronger grid-
666 like signals in mPFC than MTL (though a signal in EC was also observed, (Constantinescu,
667 O'Reilly and Behrens 2016)). In addition, Baram *et al.* found that while the abstract structure
668 of a reinforcement learning task was represented in EC, the structure-informed learning
669 signals that inform trial-by-trial behavior with generalisable information were found in
670 mPFC. Taken together, these results suggest that here, it is reasonable to expect
671 generalisation signals of community structure graphs (of which participants were aware) in
672 PFC, as well as the signals reported in EC for hexagonal graphs (of which participants were
673 unaware). Indeed, when we tested for subspace generalisation of community structure
674 graphs in the same ROI in vmPFC where Baram *et al.* found generalisable learning signals,
675 we obtained a significant result (though this is a weak effect, and we hence report it with
676 caution in the supplementary material, Figure S4b).

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

683 **Methods**

684 **Rodent cells analysis**

685 Cells electrophysiology data were taken from (Chen *et al.* 2018). In short, cells (place cells from CA1
686 and grid cells from dmEC) were recorded while the animals foraged in two different square arenas;
687 one real arena and one virtual reality (VR) arena, real arena is 60x60 and the VR arena is 60x60 or

688 90x90 cm. The VR system restrained head-movements to horizontal rotations, and included an air-
689 suspended ball on which the mice could run and turn. A virtual environment reflecting the mouse's
690 movements on the ball was projected on screens in all horizontal directions and on the floor. Mice
691 were implanted with custom-made microdrives (Axona, UK), loaded with 17mm platinum-iridium
692 tetrodes, and providing buffer amplification. We analyzed grid cells data from three animals; two
693 animals had only grid cells data and one animal had both place cells and grid cells data. We analyzed
694 place cells data from three more animals that had only place cells data (mouse 1 had 14 grid cells,
695 mouse 2 and 3 had 21 grid cells, mouse 1, 4, 5 had 25 place cells). This experimental design results in
696 two different firing rate maps, one for each arena. After preprocessing (calculate the firing rate map
697 using on 64X64 bins matrix and smoothing of the firing rate maps with 5 bins boxcar), we calculated
698 the 'subspace generalisation' score, as follows:

699 a. Calculate the neuron X neuron correlation matrix from the first firing rate map (one
700 of the environments) and its principal components (PCs).
701 b. Project the firing rate maps from this environment and the other environment on
702 these PCs.
703 c. Calculate the cumulative variance explained as a function of PCs (that are organized
704 according to their corresponding eigenvalues)
705 d. Calculate the area under the curve (AUC).

706 *Permutation test 1 (within cell type):* Our hypothesis is that the neuron X neuron correlation
707 structure is preserved while the animals forage in the two different arenas, i.e. that the active cells'
708 assemblies remain the same. Therefore, the null hypothesis is that the cells' assemblies are random
709 and did not remain the same while animals forage in the two arenas. We therefore calculated the
710 PCs using the firing rate map while the animal foraged in one environment and permuted the cells'
711 identity of the firing rate maps correspond to the second environment. We then calculated the
712 difference between the 'subspace generalisation' score within and across environments. This creates
713 our null distribution, which we compare to the subspace generalisation score of the non-permuted
714 data.

715 *permutation test 2 (between cell types):* Our hypothesis is that grid cells generalise better than place
716 cells, i.e. that the difference between the AUC of within arena projection to across arenas projection
717 is smaller in grid cells compared to place cells. To this end, we created AUC-differences distribution
718 using place cells activity as our null distribution; we sample place cells from each animal, such that
719 the number of grid cells and place cells was equal (mouse 1 had 14 grid cells, mouse 2 and 3 had 21
720 grid cells, mouse 1, 4, 5 had 25 place cells). Then, for each sample, we calculated the difference in
721 AUC (same arena - different arenas), as before. We calculated the distribution of these AUC-
722 differences values from all three animals. We then checked whether the AUC-differences in grid
723 cells, for all three animals, is significantly smaller than those predicted by the sampled place cells
724 distribution (Figure S1).

725 Reducing the resolution of the electrophysiological data

726 We first normalized all firing rate maps. Then, for each animal we randomly sampled (with repeats)
727 seven cells into two groups and averaged the cells' activity within each group, separately for each
728 environment. We then concatenated the resulted size-2 vectors from all animals into one vector and
729 used this vector as above to calculate the AUC differences between within and across environments.

730 The number of bootstraps was 400, therefore we had 800 repetitions to calculate the distribution
731 (for each sample we project on both environments therefore getting two AUC - difference values).
732 The plots in Figure.1d were smoothed with smoothing window of 9, the number of bins to calculate
733 the distribution was 50.

734 **Simulating pseudo voxels**

735 Grid cells are simulated as a thresholded sum of three 2D cosines (Burgess et al. 2007). Each module
736 is simulated by shifting the grid cells within a grid that spans the rhombus of the hexagonal grid,
737 such that the average over all grid cells within a module is a constant across the box (note that due
738 to numerical issues this is almost constant).

739 We simulated 13456 cells per module (116*116 in the x-y plane, i.e. covering the grid's rhombus).
740 The box is simulated with 50*50 resolution (the size of the "box" is 10*10). We simulated four
741 different modules that differ in their grid spacing and phases. Each environment was simulated by a
742 different phase and shift of the grid fields such that the relationships between the cells remain the
743 same across environments.

744 Voxels were simulated by averaging cells within a module. Each module was segregated into four
745 groups of cells (therefore there are 3364 cells within each voxel, see supplementary for different
746 segregations). Each voxel is an average over the cells' firing rate map within the group. The
747 averaging was done in two stages:

- 748 a. sampling grid cells randomly - i.e. not related to their grid phase
- 749 b. The remaining cells were segregated into four groups according to their phase.

750 The above process was repeated for different fractions of random/(according to phase) ratio
751 (*ratio_random* = [0,1], 0: only segregated according to phase, 1: only segregated randomly). We
752 further added spatial white noise to each voxel, noise std ranging from 0 to 0.1. When examining the
753 effect of random sampling, the noise std was 0 or 0.1.

754

755 **FMRI experiment**

756 **Participants:** 60 UCL students were originally recruited. As the training is long and hard, for each scan
757 we recruited two participants for the training sessions, and chose the better performing of the two
758 to be scanned. Overall, we scanned 34 participants and excluded 6 participants from the analysis
759 because of severe movement or sleepiness in the scanner.

760 The study was approved by the University College London Research Ethics Committee (Project ID
761 11235/001). Participants gave written informed consent before the experiment.

762

763 **Behavioural training for fMRI training task**

764 To ensure that participants understood the instructions, the first training day was performed in the
765 lab while the other three training days were performed from the participant's home.

766 **Graphs.** One hexagonal graph consisted of 36 nodes and the other 42 nodes as shown in Figure 3b.
767 One community structured graph consisted of 5 communities and the other 6 communities, with 7
768 nodes each. Within a community, each node was connected to all other nodes except for the two
769 connecting nodes that were not connected to each other but were each connected to a connecting
770 node of a neighboring community (Figure 3b). Therefore, all nodes had a degree of six, similarly to
771 the hexagonal graphs (except the nodes on the hexagonal graphs border, which had degree less than

772 six). Our community structure graph had a hierarchical structure, wherein communities were
773 organized on a ring.

774 Training procedures. In each of the training days, participants learned two graphs with the same
775 underlying structure but different stimuli. During the first two days participants learned the
776 hexagonal graphs, while during the third and fourth days participants learned the community
777 structured graphs. We chose to first teach the hexagonal graphs structure for all participants and not
778 randomize the order because learning community structure graph changes participants' learning
779 strategy (mark et al. 2020). During the fifth day, before the fMRI scan, participants were reminded of
780 all four graphs, with two repetitions of each hexagonal graph and one repetition of each community
781 structured graph. Stimuli were selected randomly, for each participant, from a bank of stimuli (each
782 pair of graphs, one hexagonal and one of a community structured graph shared the same bank).
783 Each graph was learnt during four blocks (Figure. 3b; 4 blocks for graph 1 followed by 4 blocks for
784 graph 2 in each training day). Participants could take short resting breaks during the blocks. They
785 were instructed to take a longer resting break after completing the four blocks of the first graph of
786 each learning day.

787

788 Block structure. Each block during training was made of the following tasks: 1) Learning phase 2)
789 Extending pictures sequences 3) Can it be in the middle 4) Navigation 5) Distance estimation (see
790 Figure 3). Next, we elaborate the various components of each block.

791 Learning phase (Figure 3a): Participants learned associations between graph nodes by observing a
792 sequence of pairs of pictures which were sampled from a random walk on the graph (successive
793 pairs of pictures shared a common picture). Participants were instructed to 'say something in their
794 head' in order to remember the associations. Hexagonal graphs included 120 steps of the random
795 walk per block and community-structured graphs included 180 steps per block (we introduced more
796 pictures in the community graph condition as random walks on such graphs result in high sampling
797 of transitions within a certain community and low sampling of transitions between communities).

798 Extending pictures sequences (Figure 3d): Given a target picture, which of two sequences of three
799 pictures can be extended by that picture (a sequence can be extended by a picture only if it is a
800 neighbor of the last picture in the sequence, the correct answer can be sequence 1/sequence 2/both
801 sequences): Sixteen questions per block. (note that a picture could not appear twice in the same
802 sequence, i.e. if the target picture is already in the sequence the correct answer was necessarily the
803 other sequence).

804 Can it be in the middle (Figure 3c): Determine whether a picture can appear between two other
805 pictures, the answer is yes if and only if the picture is a neighbor of the two other pictures. Sixteen
806 questions per block.

807 Navigation (Figure 3e): The aim—navigate to a target picture (appears at the right of the screen).
808 The task was explained as a card game. Participants are informed that they currently have the card
809 of the picture that appears on the left of the screen. They were asked to choose between two
810 pictures that are associated with their current picture. They could also skip and sample again two
811 pictures that are associated with the current picture, if they thought their two current options did
812 not get them closer to the target (skipping was counted as a step). In each step participants were
813 instructed to choose a picture that they thought had a smaller number of steps to the target picture
814 (according to their memory). Following choice, the chosen picture appeared on the left and two new
815 pictures, that correspond to states that are neighbors of the chosen picture, appear as new choices.

816 After a participant selected a neighbor of the target picture, that target picture itself could appear as
817 one of the new options for choice. The game terminated when either the target was reached or 200
818 steps were taken (without reaching the target). In the latter case a message 'too many steps' was
819 displayed. On the first block, for each step, the number of links from the current picture to the target
820 picture was shown on the screen. Participants played three games (i.e. navigation until the target
821 was reached or 200 steps passed) in each block, where the starting distance (number of links)
822 between the starting picture to the target was 2, 3 and 4.
823 Distance estimation: Which of two pictures has the smallest number of steps to a target picture: 45
824 questions per block (none of the 2 pictures was a direct neighbor on the graph, i.e. the minimal
825 distance was 2 and no feedback was given).

826

827 **fMRI scanning task**

828 The task consisted of four runs. Each run was divided into five blocks (one block for each graph and
829 one more repetition for one of the hexagonal graphs; the repetition was not used in the analyses in
830 this manuscript). On each block participants observed pictures that belong to one of the graphs. A
831 block started with 70sec in which participants observed, at their own pace, a random walk on the
832 graph; two neighboring pictures appeared on the screen and when participants pressed 'enter' a
833 new picture appeared on the screen (similar to the training learning phase). The new picture
834 appeared in the middle of the screen and the old picture appeared on its left. Participants were
835 instructed to infer which 'picture set' they are currently observing. No information about the graph
836 was given. This random walk phase was not used in any analyses in this manuscript.
837 Next, sequences of three pictures appeared on the screen, one after the other (note the first and
838 second pictures did not disappear from the screen until after the third picture in the sequence was
839 presented - all three pictures disappeared together, prior to the next trial, Figure 4b). To keep
840 participants engaged, once in a while (5 out of 45 sequences) a fourth picture appeared and
841 participants had to indicate whether this picture can appear next on the sequence ('catch trials',
842 Figure 4c). Before starting the fMRI scan participants were asked whether they found any
843 differences between the picture sets during the first two days (when the hexagonal graphs were
844 learnt) and the last two days (when the community graphs were learnt). Most participants (26 out of
845 28) could indicate that there were groups of pictures (i.e. communities) in the last two days, and that
846 this was not the case during the first two days. At the end of each block in the scanner participants
847 answered whether or not there are groups in the current picture set (participants that were not
848 aware of the groups were asked whether this set belongs to the first two training days or not).
849 Participants were given a bonus for answering correctly, such that 100% correct results in a ten
850 pounds bonus.

851

852 **fMRI data acquisition**

853 FMRI data was acquired on a 3T Siemens Prisma scanner using a 32 channels head coil. Functional
854 scans were collected using a T2*-weighted echo-planar imaging (EPI) sequence with a multi-band
855 acceleration factor of 4 (TR = 1.450 s, TE = 35ms, flip angle = 70 degrees, voxel resolution of
856 2x2x2mm). A field map with dual echo-time images (TE1 = 10ms, TE2 = 12.46ms, whole-brain
857 coverage, voxel size 2x2x2mm) was acquired to correct for geometric distortions due to
858 susceptibility-induced field inhomogeneities. Structural scans were acquired using a T1-weighted

859 MPRAGE sequence with 1x1x1mm voxel resolution. We discarded the first six volumes to allow for
860 scanner equilibration.

861 **Pre-processing**

862 Pre-processing was performed using tools from the fMRI Expert Analysis Tool (FEAT, Woolrich MW
863 et al. 2001; Woolrich MW et al. 2004), part of FMRIB's Software Library (FSL, Smith et al. 2004). Data
864 from each of the four scanner runs was preprocessed separately. Each run was aligned to a
865 reference image using the motion correction tool MCFLIRT. Brain extraction was performed using
866 the automated brain extraction tool BET (Smith, 2002). All data were temporally high-pass filtered
867 with a cut-off of 100s. Registration of EPI images to high-resolution structural images and to
868 standard (MNI) space was performed using FMRIB's Linear Registration Tool (FLIRT (Jenkinson et al.,
869 2002; Jenkinson and Smith, 2001)). No spatial smoothing was performed during pre-processing (see
870 below for different smoothing protocols for each analysis). Because of the notable breathing- and
871 susceptibility-related artifacts in the entorhinal cortex, we cleaned the data with FMRIB's ICA tool,
872 FIX (Griffanti *et al.* 2014; Salimi-Khorshidi *et al.* 2014).

873 **Univariate analysis**

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

877 For estimating subspace generalization, we constructed a GLM to estimate the activation as a result
878 of each three images' sequence (a 'pile' of pictures). The GLM includes the following regressors:
879 mean CSF regressor and 6 motion parameters as nuisance regressors, bias term modeling the mean
880 activity in each fMRI run, a regressor for the 'start' message (as a delta function), a regressor for the
881 self- paced random walk on each graph (a delta function for each new picture that appears on the
882 screen), a regressor for each pile in each graph (duration of a pile: 1.4sec), regressor for the catch
883 trial onset (delta) and the pile that corresponds to the catch (pile duration). All regressors beside the
884 6 motion regressors and CSF regressor were convolved with the HRF. The GLM was calculated using
885 non-normalized data.

886 **Multivariate analysis**

887 *Quantifying subspace generalization:*

888 We calculated noise normalized GLM betas within each searchlight using the RSA toolbox. For each
889 searchlight and each graph, we had a nVoxels (100) by nPiles (10) activation matrix ($B_{voxel \times pile}$) that
890 describes the activation of a voxel as a result of a particular pile (three pictures' sequence). We
891 exploited the (voxel x voxel) covariance matrix of this matrix to quantify the manifold alignment
892 within each searchlight.

893 To account for fMRI auto-correlation we used Leave One Out (LOO) approach; For each fMRI scanner
894 run and graph, we calculated the mean activation matrix over the three others scanner runs ($\hat{B}^{\sim j}$).
895 We then calculated the left Principal Component (PCs) of that matrix ($U_{voxel \times voxel}^{\sim j}$). To quantify the
896 alignment, we projected the excluded scanner run graph activation matrix (B^j) of each graph on

897 these PCs and calculated the accumulated variance explained as a function of PCs, normalized by the
898 total variance of each graph within each run. Therefore, for each run and graph we calculated:

$$P^{a,b} = U_a^{\sim j} \cdot B_b^j$$

$$M_k^{a,b} = \frac{\sum_{l=1}^{10} (P_{a,b}^{l,k})^2}{S^j}$$

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

900 Where $P^{a,b}$ is the projection matrix of dimensions $voxel \times pile$ of graph 'b' on the PCs of graph 'a',
901 $M_k^{a,b}$ is the normalized variance explained on the 'k' direction, S^j is the summation of the diagonal of
902 Σ^j , the total variance as a result of the graph piles (three images sequence). We then calculated the
903 cumulative variance explained over all 'k' PCs directions. As a summary statistic we calculated the
904 area under this curve. This gives us a 4x4 alignment matrix, for each run, such that each entry (a, b)
905 in this matrix is a measure of the alignment of voxels patterns as a result of the two graphs a&b
906 (Figure 4d). We then averaged over the four runs and calculated different contrasts over this matrix.

907 The above calculations were performed in subject space, we therefore normalized the searchlight
908 results and then smoothed with a kernel of 6mm FWHM using FSL FLIRT and FNIRT before
909 performing group level statistics.

910 For group level we calculated the t-stat over participants of each contrast:

911 Visual contrast was $[H|H|H + C|C|C + Hs|Hs + Cs|Cs] - [H|H|Hs + Hs|H + C|Cs + Cs|C]$, i.e. same exact sequences
912 controlled by the same structure.

913 Structural contrast was $[H|H|H + H|H|Hs + Hs|H + Hs|Hs] - [H|C|C + H|C|Cs + Hs|C|C + Hs|Cs|Cs]$, i.e. the difference
914 between subspace generalisation of hexagonal graphs data, when projected on PCs calculated from
915 (cross-validated) hexagonal graphs (yellow elements in middle panel) vs community structure graphs
916 (red elements).

917 Multiple comparisons correction

918 Multiple comparison correction was performed using the permutation tests machinery (Nichols and
919 Holmes 2002) in PALM (Winkler *et al.* 2014): within the mask we used for multiple comparisons
920 correction (details in main text), we first measured the TFCE statistic for the current contrast. We
921 then repeated this procedure for each of the 10000 random sign-flip iterations (each participant's
922 contrast sign was randomly flipped and the statistic over participants was calculated). Using these
923 values we then created a null distribution of TFCE statistics by saving only the voxel with the highest
924 TFCE in each iteration. Comparing the true TFCE to the resulting null distributions results in FWE-
925 corrected TFCE P-values.

926 The code for the analysis and simulation is in:

927 https://github.com/ShirleyMgit/subspace_generalization_paper_code/tree/main

928 **Author contributions**

929 S.M and T.B conceive the research, S.M and P.S design and perform the experiment, S.M and A.B
930 analyzed the data. A.H and V.S provided consultation with analysis. S.M, A.B and T.B wrote the
931 manuscript.

932 **Acknowledgements**

933 T.B. is supported by a Wellcome Principal Research Fellowship (219525/Z/19/Z), a Wellcome
934 Collaborator award (214314/Z/18/Z), for the purpose of Open Access, the author has applied a CC BY
935 public copyright licence to any Author Accepted Manuscript version arising from this submission,
936 a JS McDonnell Foundation award (JSMF220020372), and by the Jean Francois and Marie-Laure de
937 Clermont Tonerre Foundation. The Wellcome Centre for Integrative Neuroimaging and Wellcome
938 Centre for Human Neuroimaging are each supported by core funding from the Wellcome Trust
939 (203139/Z/16/Z, 203147/Z/16/Z). The Sainsbury-Wellcome centre is supported by core funding from
940 the Wellcome Trust (219627/Z/19/Z) and the Gatsby Charitable Foundation (GAT3755).

941 A.H. was supported by the European Molecular Biology Organization nonstipendiary Long-Term
942 Fellowship (848–2017), Human Frontier Science Program (LT000444/2018), Israeli National
943 Postdoctoral Award Program for Advancing Women in Science, and the European Union’s Horizon
944 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No.
945 789040.

946 **References**

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

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

951 Barron HC, Garvert MM, Behrens TE. Repetition suppression: a means to index neural
952 representations using BOLD? *Philos Trans R Soc B Biol Sci* 2016;371:20150355.

953 Behrens TE, Muller TH, Whittington JC *et al.* What is a cognitive map? Organizing knowledge for
954 flexible behavior. *Neuron* 2018;100:490–509.

955 Bongioanni A, Folloni D, Verhagen L *et al.* Activation and disruption of a neural mechanism for novel
956 choice in monkeys. *Nature* 2021;591:270–4.

957 Burak Y, Fiete IR. Accurate path integration in continuous attractor network models of grid cells.
958 *PLoS Comput Biol* 2009;5:e1000291.

959 Burgess N, Barry C, O’keefe J. An oscillatory interference model of grid cell firing. *Hippocampus*
960 2007; 801-812, 17(9).

961 Chen G, King JA, Lu Y *et al.* Spatial cell firing during virtual navigation of open arenas by head-
962 restrained mice. *Elife* 2018;7:e34789.

963 Constantinescu AO, O’Reilly JX, Behrens TE. Organizing conceptual knowledge in humans with a
964 gridlike code. *Science* 2016;352:1464–8.

965 Diedrichsen J, Kriegeskorte N. Representational models: A common framework for understanding
966 encoding, pattern-component, and representational-similarity analysis. *PLOS Comput Biol*
967 2017;13:e1005508.

968 Fyhn M, Hafting T, Treves A *et al.* Hippocampal remapping and grid realignment in entorhinal cortex.
969 *Nature* 2007;446:190–4.

970 Gardner RJ, Hermansen E, Pachitariu M *et al.* Toroidal topology of population activity in grid cells.
971 *Nature* 2022;602:123–8.

972 Gardner RJ, Lu L, Wernle T *et al.* Correlation structure of grid cells is preserved during sleep. *Nat
973 Neurosci* 2019;22:598–608.

974 Garvert MM, Dolan RJ, Behrens TE. A map of abstract relational knowledge in the human
975 hippocampal–entorhinal cortex. *elife* 2017;6:e17086.

976 Griffanti L, Salimi-Khorshidi G, Beckmann CF *et al.* ICA-based artefact removal and accelerated fMRI

977 acquisition for improved resting state network imaging. *NeuroImage* 2014;**95**:232–47.

978 Grill-Spector K, Henson R, Martin A. Repetition and the brain: neural models of stimulus-specific
979 effects. *Trends Cogn Sci* 2006;**10**:14–23.

980 Gu Yi, Lewallen S, Kinkhabwala A, Domnisoru C, Yoon K, Gauthier JL, Fiete IR, and Tank DW. Cell
981 2018; 175, 735-750. Hahamy A, Behrens TE. Measuring the spatial scale of brain
982 representations. *2019 Conference on Cognitive Computational Neuroscience*. 2019, 2019–
983 1174.

984 Haxby JV, Gobbini MI, Furey ML *et al.* Distributed and overlapping representations of faces and
985 objects in ventral temporal cortex. *Science* 2001;**293**:2425–30.

986 Kemp C, Tenenbaum JB. The discovery of structural form. *Proc Natl Acad Sci* 2008;**105**:10687–92.

987 Klein-Flügge MC, Bongioanni A, Rushworth MFS. Medial and orbital frontal cortex in decision-making
988 and flexible behavior. *Neuron* 2022;**110**:2743–70.

989 Klein-Flügge MC, Wittmann MK, Shpektor A *et al.* Multiple associative structures created by
990 reinforcement and incidental statistical learning mechanisms. *Nat Commun* 2019;**10**:4835.

991 Kriegeskorte N, Mur M, Bandettini P. Representational similarity analysis - connecting the branches
992 of systems neuroscience. *Front Syst Neurosci* 2008;**2**.

993 Mark S, Moran R, Parr T *et al.* Transferring structural knowledge across cognitive maps in humans
994 and models. *Nat Commun* 2020;**11**:4783.

995 Nichols TE, Holmes AP. Nonparametric permutation tests for functional neuroimaging: A primer with
996 examples. *Hum Brain Mapp* 2002;**15**:1–25.

997 Nili H, Wingfield C, Walther A *et al.* A Toolbox for Representational Similarity Analysis. *PLOS Comput
998 Biol* 2014;**10**:e1003553.

999 Park SA, Miller DS, Nili H *et al.* Map Making: Constructing, Combining, and Inferring on Abstract
1000 Cognitive Maps. *Neuron* 2020;**107**:1226-1238.e8.

1001 Salimi-Khorshidi G, Douaud G, Beckmann CF *et al.* Automatic denoising of functional MRI data:
1002 Combining independent component analysis and hierarchical fusion of classifiers.
1003 *NeuroImage* 2014;**90**:449–68.

1004 Samborska V, Butler JL, Walton ME *et al.* Complementary task representations in hippocampus and
1005 prefrontal cortex for generalizing the structure of problems. *Nat Neurosci* 2022;**25**:1314–26.

1006 Schuck NW, Cai MB, Wilson RC *et al.* Human Orbitofrontal Cortex Represents a Cognitive Map of
1007 State Space. *Neuron* 2016;**91**:1402–12.

1008 Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, Bannister
1009 PR, De Luca M, Drobniak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady
1010 JM, Matthews PM. Advances in functional and structural MR image analysis and implementation as
1011 FSL. *NeuroImage* 2004; 23: S208-S219.

1012 Tolman EC. Cognitive maps in rats and men. *Psychol Rev* 1948;**55**:189–208.

1013 Trettel SG, Trimper JB, Hwaun E, Fiete IR, Colgin LL. Grid cell co-activity patterns during sleep reflect
1014 spatial overlap of grid fields during active behaviors. *Nature Neuroscience* 2019; 22: 609-617.

1015 Waaga T, Agmon H, Normand VA *et al.* Grid-cell modules remain coordinated when neural activity is
1016 dissociated from external sensory cues. *Neuron* 2022;**110**:1843-1856.e6.

1017 Whittington JC, McCaffery D, Bakermans JJ *et al.* How to build a cognitive map. *Nat Neurosci*
1018 2022;**25**:1257–72.

1019 Whittington JC, Muller TH, Mark S *et al.* The Tolman-Eichenbaum machine: unifying space and
1020 relational memory through generalization in the hippocampal formation. *Cell*
1021 2020;**183**:1249–63.

1022 Wilson RC, Takahashi YK, Schoenbaum G *et al.* Orbitofrontal Cortex as a Cognitive Map of Task
1023 Space. *Neuron* 2014;**81**:267–79.

1024 Winkler AM, Ridgway GR, Webster MA *et al.* Permutation inference for the general linear model.
1025 *NeuroImage* 2014;**92**:381–97.

1026 Woolrich M W, Ripley BD, Brady M, & Smith, SM. (2001). Temporal Autocorrelation in Univariate
1027 Linear Modeling of fMRI Data. *NeuroImage*, 14(6), 1370–1386

1028 Woolrich MW, Behrens TEJ, Beckmann CF, Jenkinson M, Smith SM. Multilevel linear
1029 modelling for fMRI group analysis using Bayesian inference. *NeuroImage* 2004; 21:
1030 1732-1747.

1031 Xie J, Padoa-Schioppa C. Neuronal remapping and circuit persistence in economic decisions. *Nat
1032 Neurosci* 2016;19:855–61.

1033 Yoon K, Buice MA, Barry C *et al.* Specific evidence of low-dimensional continuous attractor dynamics
1034 in grid cells. *Nat Neurosci* 2013;16:1077–84.