

## Reviewed Preprint

v1 • April 2, 2026

Not revised

## ✉ For correspondence:

[leonie.glitz@bndu.ox.ac.uk](mailto:leonie.glitz@bndu.ox.ac.uk)[christopher.summerfield@psy.ox.ac.uk](mailto:christopher.summerfield@psy.ox.ac.uk)

**Competing interests:** Declaration of Interests The authors declare no competing interests.

**Funding:** See [page 22](#)

**Reviewing editor:** Anna C Schapiro, University of Pennsylvania, United States

© 2026, Glitz et al. This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

# How does curriculum learning influence cognitive map formation?

Leonie Glitz<sup>1,2</sup> ✉, Zilu Liang<sup>2</sup>, Helen C Barron<sup>1</sup>, Christopher Summerfield<sup>2</sup> ✉

<sup>1</sup>Medical Research Council Centre of Research Excellence in Restorative Neural Dynamics, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, United Kingdom • <sup>2</sup>Department of Experimental Psychology, University of Oxford, Oxford, United Kingdom

## eLife Assessment

This **valuable** paper provides **convincing** evidence that humans can navigate better through maps whose local transitions were learned in an intermixed order than maps whose local transitions were learned in neighboring groups. The authors put forward a potential mechanism in which the grouped learning resulted in mental fragmentation, though evidence for this mechanism is **incomplete**. The work will be of interest to researchers studying cognitive maps and curriculum learning.

<https://doi.org/10.7554/eLife.111035.1.sa3>

## Abstract

Representing relational knowledge in a mental map is useful for adaptive behaviour. Here, we studied how curriculum learning (the use of different training regimes) influences the way that mental (or ‘cognitive’) maps are formed. We investigated how the order in which people learned one-step transitions impacted their subsequent ability to perform multi-step inference using the cognitive maps they had formed. Participants learned one map with a curriculum that grouped information by rows and columns (‘grouped curriculum’) and another that selected transitions for training in a spatially disjoint way (‘disjoint curriculum’). Subsequent multi-step navigation was performed better after training using the disjoint curriculum. We hypothesise that learning under a grouped curriculum results in mental fragmentation of the learned maps, which impairs subsequent navigation across rows and columns. By contrast, learning one-step transitions in a pseudo-random order facilitates subsequent integration of learned information into a mental map to enable adaptive behaviour.

## Introduction

Cognitive maps are mental representations that encode how states of the world relate to each other (Behrens et al., 2018 [↗](#)). They support goal-directed behaviours by allowing us to mentally explore the connections between places (Bellmund, Deuker, Navarro Schröder, & Doeller, 2016 [↗](#); Tolman, 1948 [↗](#)), objects (Constantinescu, O’Reilly, & Behrens, 2016 [↗](#)) and people (Park, Miller, & Boorman, 2021 [↗](#)), and thereby make inferences about potentially new combinations of stimuli or events (Bellmund, Gärdenfors, Moser, & Doeller, 2018 [↗](#)). How cognitive maps are used to make decisions has been the subject of intense research over recent years, with a focus on the neurophysiology of the mammalian hippocampal-entorhinal system (Momennejad, 2020 [↗](#); Stachenfeld, Botvinick, & Gershman, 2017 [↗](#)). However, we know little about how cognitive maps are formed by experience, as people learn about the relations among stimuli. Here, we studied the impact of the sequence of training stimuli (or curriculum) on how cognitive maps are learned.

For almost all learners, the sequence of training stimuli (or curriculum) is an important driver of how knowledge is acquired. One salient feature of a curriculum is whether related materials are grouped or separated in time. Very often, temporal grouping is found to help learning. In educational settings, materials are naturally organised into thematically related lessons, so that

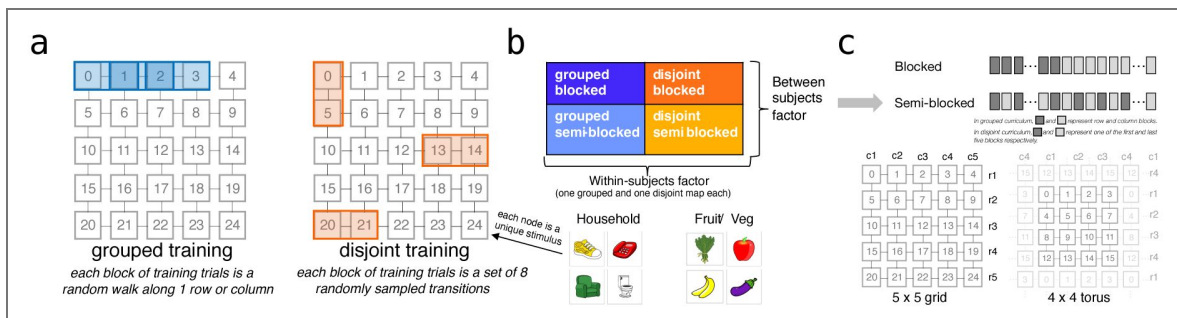
(for example) students are not required to learn French and Spanish at once. Temporal grouping has also been found to facilitate learning in laboratory studies. During rote memorisation, a technique by which to-be-learned stimuli are organised into small groups ('chunking') can benefit both short- and long-term retention (Mitchell & Martin, 1997 [↗](#); Xu & Padilla, 2013 [↗](#)). When learning different categorisation tasks in succession, grouping the tasks in temporally correlated blocks ('blocking') facilitates learning compared to a regime in which trials are randomly intermixed (Flesch, Balaguer, Dekker, Nili, & Summerfield, 2018 [↗](#); Noh, Yan, Bjork, & Maddox, 2016 [↗](#)). Similarly, in a paradigm where multiple visual features of an object predict a corresponding spatial location, generalisation is improved when participants learn separately about the predictive value of each feature (e.g. colour separately from shape), compared to when they are intermixed in time (Dekker, Otto, & Summerfield, 2022 [↗](#)).

Grouping information in time may help learning by allowing the student to 'divide and conquer' – to decompose a task into parts which are more manageable to learn (Krueger & Dayan, 2009 [↗](#); Mi & Summerfield, 2025 [↗](#)). In biological systems, separation in time may help the learning system partition new knowledge, so that learning of one part does not interfere with another. Place cell codes in the hippocampus, for example, have been shown to overlap more for distinct contexts experienced on the same day than on subsequent days and conditioned fear responses are transferred across contexts experienced on the same day, even if the actual fear conditioning only occurred in one of the contexts (Cai et al., 2016 [↗](#)), supporting the idea that separation in time can prevent interference.

However, it remains unclear how different learning curricula might support the formation of an integrated cognitive map to permit adaptive behaviour. Here, to test the effect of learning curriculum on cognitive map formation, participants learned to navigate grid or torus environments. Participants were always trained on 1-step transitions but were later tested on their ability to perform 2-step navigation which provided a readout for integration across the cognitive map. Across multiple experiments, we found that training on 1-step transitions was more successful at facilitating 2-step navigation when transitions were sampled randomly (disjoint condition), rather than being drawn from rows or columns in a temporally correlated fashion (grouped condition; Figure 1 [↗](#)). More fine-grained examination of the behavioural data hinted that relatively impaired inference following grouped training may have occurred because grouped training encouraged participants to learn rows and columns as a series of unconnected 'fragments', and thus to learn a partially fragmented version of the cognitive map.

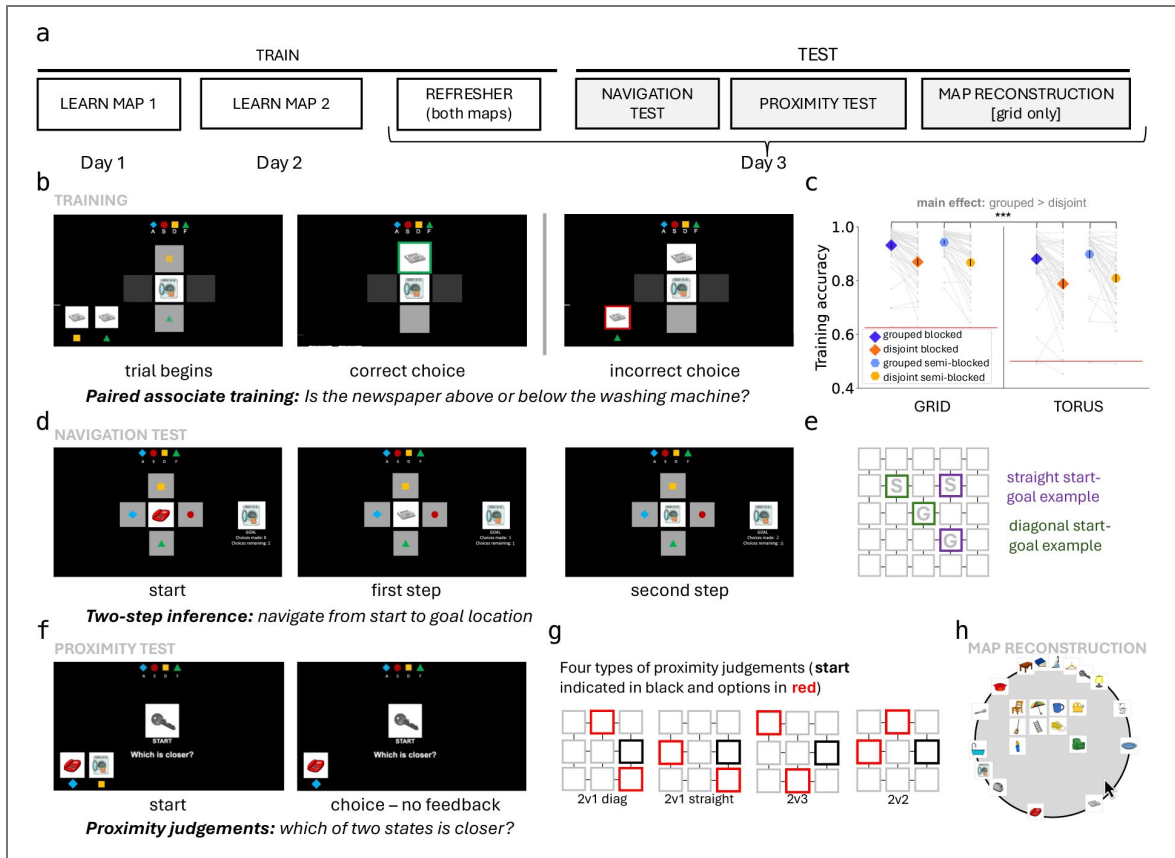
## Results

Our experiments used maps in the form of quadrilateral lattices arranged into either a grid (5×5) or a torus (4×4; manipulated between-subjects), in which each state corresponded to a unique image of either a household item (map 1) or a fruit / vegetable (map 2), assigned randomly between participants (Figure 1b [↗](#)). Participants never viewed these full maps at once (e.g. they were never shown a bird's eye view). To test the effect of learning curriculum on participants' ability to perform multi-step inference, over three days we trained participants on an exhaustive set of one-step associations between states before then testing their ability to complete two-step inference. During training, they navigated from one state (e.g. a washing machine) to one of its neighbours (e.g. a newspaper) by selecting whether the newspaper was up, down, left or right of the washing machine and receiving feedback as to whether their choice was correct or not (Figure 2b [↗](#)). On the second day, they repeated the learning task for the other map. On the third day, following a short refresher on both maps, they were tested on their ability to navigate between states that were 2 steps apart, either along a single axis (straight or I trials) or crossing the horizontal and vertical axes (diagonal or L trials, Figure 2c [↗](#)). They were also tested on their ability to make proximity judgements between different states. Both tasks required participants to have integrated the one-step adjacency associations they had learned on the preceding two days into a mental map. We tested each of the two maps they had learned separately on day 3.



**Figure 1. Illustration of the study design.**

**a** illustration of three successive example trials within a block of trials in the grouped (right, blue) and disjoint (left, orange) training curriculum. Numbers represent nodes on the map, coloured boxes represent transitions between adjacent nodes learned. Notably, the successive trials in grouped training are overlapping (pseudo-random walk), whereas the ones in disjoint training are not **b** *Top*. Illustration of the mixed study design, note that background colours for design graphic correspond to plot marker colours in subsequent figures (Figure 2 [↗](#) and Figure 3 [↗](#)). *Bottom*. Examples of the two stimulus categories, each was randomly assigned to one of the two within-subjects maps for each participant **c** *top* illustration of the between-subjects manipulation, filled squares correspond to blocks of training trials *bottom* illustration of the second between-subjects factor, map shape.



**Figure 2. Task sequence and training accuracy.**

**a** multi-day task schedule, text in boxes symbolises the sequence of events (from left to right), *test* events depicted with a grey background, *training* with a white background. Day 1, 2 and 3 refer to the day of the experiment on which the task in question occurred. **b** Trial sequence for a select-action training trial (for a select-state trial example, see [Supplementary Results 1](#)) including illustrations of feedback for a correct and incorrect choice. Washing machine is the starting node, the two shapes (here: yellow square and green triangle) indicate the two positions the adjacent node (here: newspaper) could be in relative to the starting node. Shapes at the top indicate which key participants need to press to select either of the options. **c** training accuracy (y-axis) averaged across initial training and refresher (left); diamond/hexagon markers indicate average accuracy for the different conditions, black vertical lines are standard error of the mean bars and grey horizontal lines are accuracies for individual participants, red horizontal line is chance accuracy; \*\*\*  $p < .0001$ . **d** trial sequence of a two-step navigation trial, as participants make a choice, the central node (current location) changes accordingly and the number of steps remaining decreases; **e** grid map layout with examples of the spatial relations between two-step starts and goals, example straight 2 step pair is indicated in purple and diagonal pair is indicated in green **f** illustration of the layout of a proximity test trial **g** as well as the types of proximity comparisons; central image examples shown in black and options for proximity judgements shown in red; **h** illustration of the task layout of the reconstruction task employed for the grid groups.

To compare curricula for cognitive map learning, during training we manipulated the order in which participants experienced the 1-step transitions, as shown in [Figure 1](#). We varied both the content of the individual blocks of one-step transitions (within-subjects) and the order in which participants experienced the blocks of transitions (between-subjects). For one map, the one-step transitions learned in a block of training trials encompassed all transitions within a single row or column of the lattice (*grouped condition*) in the form of a semi-random walk, whereas for the other map, the trials in a single block of training were drawn randomly from across the lattice structure (*disjoint condition*; [Figure 1a](#)), thereby manipulating the temporal proximity with which participants learned about adjacent nodes. Across participants, we introduced a blocking manipulation to further temporally separate trials that may cause interference, more specifically the rows and columns of the lattice on the grouped curriculum map. To do this, we manipulated the order of the blocks of trials, with some participants experiencing half of the transitions twice (in grouped: e.g. all row blocks twice), followed by the second half of the transitions twice (in grouped: e.g. all column blocks twice; *blocked condition*) and other participants experiencing the blocks of trials in an interleaved way (*semi-blocked condition*; [Figure 1c](#)).

In total, we conducted 4 experiments (two groups per map shape; one blocked and one semiblocked) involving a total of 182 participants, of whom 92 learned grid maps (49 blocked, 44 semi-blocked) and 90 learned torus maps (46 blocked, 44 semi-blocked). Each cohort started with 50 participants and final group numbers depended on the participant attrition rate across the three days of the experiment.

We model data from all four studies in a single statistical model, entering map shape (grid vs torus; between subjects, [Figure 1c](#)), our main between-subjects factor (fully blocked or semiblocked, [Figure 1c](#)) and our within-subjects factor (disjoint or grouped, [Figure 1a](#)) as predictors in linear mixed effects models. For a table of all model results, as well as all results split up by map shape and group, please see [Supplementary Tables 1–3](#). Note that participants in the torus groups had lower accuracy on almost all cognitive map learning measures than participants in the grid groups, even when the necessarily lower training accuracy due to the lack of edges was taken into account ([Supplementary Table 1](#)).

## Training accuracy was higher in the grouped condition

During one-step training, participants were significantly more accurate on the map with the grouped curriculum than the disjoint curriculum (see [Figure 1](#), [Figure 2](#) and Methods). Overall training accuracy for all 1-step transitions on days 1, 2 and 3 (refresher) was higher for the map learned via a grouped than a disjoint curriculum ( $\beta = -.08$ ,  $z = -9.32$ ,  $SE = .01$ ,  $p_{\text{perm}} < .0001$ ; [Figure 2e](#);  $M_{\text{grouped}} = .91$  (SD .10),  $M_{\text{disjoint}} = .83$  (SD .11)). The same effect was observed when the data were split into initial training (days 1 or 2:  $\beta = -.09$ ,  $z = -9.69$ ,  $SE = .01$ ,  $p_{\text{perm}} < .0001$ ,  $M_{\text{grouped}} = .91$  (SD .10),  $M_{\text{disjoint}} = .82$  (SD .11)) and refresher (day 3:  $\beta = -.05$ ,  $z = -5.72$ ,  $SE = .01$ ,  $p_{\text{perm}} < .0001$ ,  $M_{\text{grouped}} = .93$  (SD .09),  $M_{\text{disjoint}} = .87$  (SD .12)). This accuracy advantage for the grouped curriculum is perhaps not surprising, as learning spatially proximal transitions (within rows and columns) together may be mutually reinforcing.

One possibility is that higher training accuracy on the grouped map occurs because subsequent trials were presented as a semi-random walk along a row or column; meaning that participants would frequently have to repeat the same action (e.g. go right) multiple times in a row (see Methods). To test whether the higher training accuracy could be attributed to purely response repetition, requiring less attention from the participant, we conducted a supplementary control experiment. In this, we compared grouped curricula in which rows and columns were learned strictly in order (e.g., A-B; B-C; C-D) or where they were shuffled within a row or column (e.g., B-C; A-B; D-C). In this additional study, we observed no difference between the ordered and shuffled grouped curriculum for training or test accuracy, suggesting that the grouping rather than response coherence was driving the effect ([Supplementary Results 2](#)).

There was no effect of blocking (blocked vs. semi-blocked condition) on training accuracy and no significant interaction (both  $p > .05$ , see [Supplementary Table 1](#) for details), suggesting that one-step training accuracy did not vary as a function of the order of the blocks of training trials and

that the effect of our within-subjects manipulation (disjoint vs grouped) on training accuracy was not dependent on the order of the blocks of training trials.

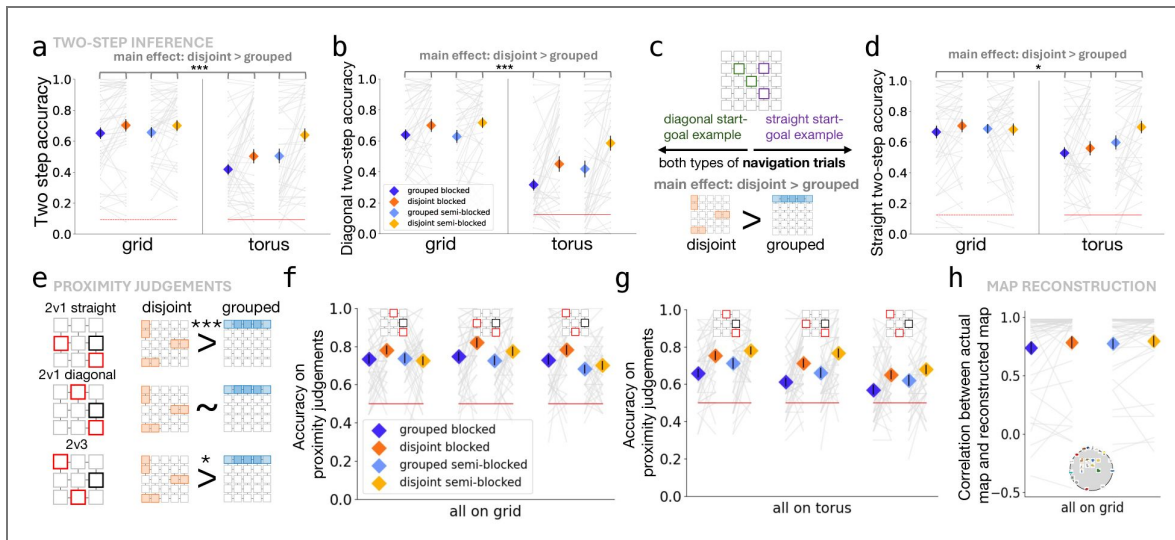
## Multi-step navigation was more accurate after training in the disjoint condition

On day 3, following the refresher on the one-step training, participants navigated to a goal state that was 2 steps away from the start state, so that two directions had to be chosen in succession to successfully complete the journey. During the two-step navigation, participants were given feedback on their choices in the form of state-transitions, i.e. after they had made their first choice, the current node would change. They were not able to make any more than two choices, however, meaning participants needed to make two correct successive choices to successfully reach the goal node. This test phase was identical for both maps, so any differences must be attributable to the curriculum used to learn each map. Two-step navigation was significantly more accurate for the map learned using the disjoint curriculum, relative to the grouped curriculum ( $\beta = .09$ ,  $z = 4.39$ ,  $SE = .02$ ,  $p < .0001$ ;  $M_{\text{grouped}} = .56$  (.28),  $M_{\text{disjoint}} = .64$  (.29); [Figure 3a](#)). This effect was driven by a difference in accuracy on both the first ( $\beta = .04$ ,  $z = 3.94$ ,  $SE = .01$ ,  $p < .0001$ ,  $M_{\text{grouped}} = .74$  (.19),  $M_{\text{disjoint}} = .78$  (.19)) and second (conditional on first choice correct:  $\beta = .06$ ,  $z = 4.43$ ,  $SE = .01$ ,  $p < .0001$ ,  $M_{\text{grouped}} = .71$  (.22),  $M_{\text{disjoint}} = .76$  (.24)) step of two-step trials. As in training, we observed no effect of blocking on two-step accuracy (fully vs. semi-blocked,  $\beta = -.05$ ,  $z = -1.10$ ,  $SE = .04$ ,  $p = .272$ ,  $M_{\text{blocked}} = .57$  (.29),  $M_{\text{semi-blocked}} = .62$  (.28)) and no interaction between grouping and blocking ( $\beta = -.02$ ,  $z = -0.80$ ,  $SE = .03$ ,  $p = .426$ ).

This advantage for the map trained with the disjoint curriculum was observed for diagonal (or “L”) trials (where the goal was diagonal to the start, and navigation involved integrating knowledge from temporally separate row and column training;  $\beta = .13$ ,  $z = 5.07$ ,  $SE = .03$ ,  $p < .0001$ ,  $M_{\text{grouped}} = .50$  (.32),  $M_{\text{disjoint}} = .61$  (.32), [Figure 3c](#)). Perhaps surprisingly, it was also observed for straight (or “I”) trials, where the two steps required were in the same row or column;  $\beta = .05$ ,  $z = 2.20$ ,  $SE = .02$ ,  $p = .03$ ,  $M_{\text{grouped}} = .62$  (.27),  $M_{\text{disjoint}} = .66$  (.29) ([Figure 3d](#); see also [Supplementary Table 1](#)). However, using an expanded mixed linear effects model where type of two-step trial (diagonal vs. straight) was an additional predictor, we found a significant interaction between two-step trial type and our within-subjects manipulation on two step navigation accuracy ( $\beta = -.08$ ,  $z = 5.03$ ,  $SE = .02$ ,  $p = .02$ ), with the difference between disjoint and grouped map accuracy being greater for diagonal than straight trials during 2-step navigation (see [Supplementary Table 1](#) for full model results).

To further investigate why the disjoint training curriculum resulted in better multi-step navigation, we also analysed the nature of the errors made. Specifically, we investigated whether participants were more likely to perform two steps in the same direction (in line with staying within a row or column) or take a corner on trials where they failed to reach the intended goal. We observed a weak trend for participants’ errors to consist of a greater proportion of ‘straight line errors’ after grouped than disjoint training ( $\beta = -.11$ ,  $z = -1.95$ ,  $SE = .06$ ,  $p = .05$ ,  $M_{\text{grouped}} = .37$  (.25),  $M_{\text{disjoint}} = .32$  (.25)). Coupled with the interaction between curriculum (grouped vs. disjoint) and trial type (diagonal vs. straight) on accuracy, this may imply that after grouped training, participants struggle to integrate information across rows and columns, i.e. across groups of transitions that were trained separately in time.

Although we did not provide overt feedback during 2-step navigation, participants were shown the state that resulted from their first transition, and so could in theory have learned during the test phase. However, we observed no significant differences in accuracy between the first and second block of 2-step navigation, and no significant interaction between block and either our grouping or blocking curriculum interventions (all p-values above an alpha level of 5%).



**Figure 3. Participants are significantly better at multi-step planning and inference after disjoint compared to grouped training.**

**a** average two-step navigation accuracy across both types of two-step trials (y-axis), markers represent average accuracies, vertical black lines are standard error of the mean (SEM) bars, horizontal grey lines are accuracies for individual participants, red horizontal lines represent accuracy at chance; N.B. the same applies to the graphs in b, d, f, g and h **b** participants' accuracy on diagonal two-step trials; **c** illustration of the two types of two-step trials (*top*) as well as the direction of difference for both types (higher accuracy after disjoint than grouped training, *bottom*) **d** average two-step navigation accuracy for straight two-step trials **e** illustration of the significance of difference in accuracy on the different proximity judgement types, \*\*\*  $p < .0001$ , \*\*  $p < .01$ , \*  $p < .05$ , ~ n.s.; black frame on grid indicates example central node, red frame on grid indicates two example comparisons for proximity. These icons are repeated in f and g **f** accuracy on proximity judgement task in grid groups; grid icons above the average markers indicate the type of proximity trial **g** accuracy on proximity judgement task in torus groups; grid icons above the average markers indicate the type of proximity trial **h** map reconstruction accuracy at the end of testing in the grid groups, y-axis shows Spearman correlation between actual grid maps and reconstructed maps. *Red line indicates accuracy when performing at chance.*

## Proximity judgements were more accurate after training in the disjoint condition

We also measured how well participants were able to mentally reconstruct the physical layout of the map after each of the training curricula. After having completed the navigation part of the test phase, we asked participants to perform an additional task that involved making proximity judgements. Participants were shown an image (corresponding to a node) centrally on the screen and then asked to choose which of two other images (nodes) was closer to the central image on the map they had learned (Figure 3e [↗](#)). As with the two-step navigation, participants were more accurate at proximity judgements for the map learned with the disjoint (rather than grouped) curriculum ( $\beta = .05$ ,  $z = 3.79$ ,  $SE = .01$ ,  $p_{\text{perm}} < .0001$ ;  $M_{\text{grouped}} = .68$  (.15),  $M_{\text{disjoint}} = .74$  (.15)).

Between proximity trials, we varied the distance of the two nodes participants chose between relative to the starting node. There was a significant advantage of disjoint over grouped training for the 2-away (diagonal) vs 3-away (diagonal; Figure 3 e-g [↗](#),  $\beta = .04$ ,  $z = 2.09$ ,  $SE = .02$ ,  $p = .04$ ,  $M_{\text{grouped}} = .65$  (.19),  $M_{\text{disjoint}} = .70$  (.18)) and 2-away (straight) vs 1-away ( $\beta = .08$ ,  $z = 3.98$ ,  $SE = .02$ ,  $p_{\text{perm}} < .0001$ ,  $M_{\text{grouped}} = .69$  (.18),  $M_{\text{disjoint}} = .77$  (.17)), but not for the 2-away (diagonal) vs 1-away judgements ( $p > .05$ ,  $M_{\text{grouped}} = .71$  (.18),  $M_{\text{disjoint}} = .76$  (.16); all Figure 3 e-g [↗](#)). There was no effect of whether participants learned each map with disjoint or grouped training on their preference for judging straight or diagonal two-steps as closer (2 (diagonal) vs 2 (straight);  $\beta = .001$ ,  $z = 0.03$ ,  $SE = .02$ ,  $p = .98$ ,  $M_{\text{grouped}} = .47$  (.18),  $M_{\text{disjoint}} = .47$  (.19), see also [Supplementary Table 1](#) [↗](#)). Overall, this is consistent with the results from multi-step planning, suggesting that participants had integrated the learned paired associations into a cognitive map more successfully after disjoint than grouped training.

To test participants' ability to reconstruct the maps based on different training curricula, we also conducted an additional test in which we asked participants to drag and drop the items into an arena in a way that matched the organisation of the map (in the grid cohorts). On average, participants were able to perform this task above chance for both maps and in all groups, as assessed by comparing the Fisher z-transformed correlation of their reconstructed map to the actual map to zero (blocked:  $M_{\text{grouped}} = .74$  (.36),  $t(48) = 11.96$ ,  $p < .0001$ ,  $M_{\text{disjoint}} = .79$  (.37),  $t(48) = 13.53$ ,  $p < .0001$ ; semi-blocked:  $M_{\text{grouped}} = .77$  (.29),  $t(42) = 12.29$ ,  $p < .0001$ ,  $M_{\text{disjoint}} = .80$  (.31),  $t(42) = 13.14$ ,  $p < .0001$ ). However, while participants were numerically better at reconstructing the maps learned with disjoint training ( $M_{\text{diff}} = .04$ ,  $SD_{\text{diff}} = .19$ ), this difference was not statistically significant ( $\beta = .02$ ,  $z = .75$ ,  $SE = .03$ ,  $p = .46$ ).

## Discussion

How humans and other animals learn mental maps has been intensively studied. However, training conditions that might facilitate the acquisition and use of relational knowledge are not well understood. We tested how curricula that grouped adjacent transitions in time (grouped vs. disjoint curriculum) impact human learning, as indexed by subsequent multi-step navigation performance. We found that disjoint training was most effective. This is perhaps surprising, because other recent studies have emphasised that correlating information in time can be helpful for the acquisition of associative and categorical information.

The most interesting explanation is that grouped training is detrimental because it encourages participants to learn “fragments” of knowledge that are encoded distinct from each other (corresponding to individual rows and columns of the map), whereas disjoint training encourages the learning of a more integrated relational map. Importantly, we controlled how trained associations varied within grouped and disjoint conditions, with disjoint conditions similarly organised into mini-blocks in which a small set of transitions were repeated in temporal proximity. Thus, the only difference between the conditions was how the states that were associated during training were arranged geometrically in the map. For example, in the grouped condition, participants might learn associations A-B and B-C in block 1, and F-G and G-H in block 2, encouraging them to form two “fragments” (A-B-C and F-G-H). By contrast, in the disjoint condition, participants could in theory learn A-B and F-G in block 1, and B-C and G-H in block 2. In

other words, by integrating across blocks participants could form a “big picture” (A-B-C) indicating how the map was organised. In particular, it is widely accepted that paired associate inference can occur over time when a new association (e.g. B paired with C) elicits memory for a previous association (A-B), and the recall process provides a window for neural plasticity to allow A to be paired with C via B (Ghandour et al., 2019 [↗](#); Josselyn & Tonegawa, 2020 [↗](#); Liu et al., 2012 [↗](#); Schacter & Loftus, 2013 [↗](#); Yokose et al., 2017 [↗](#)). In our disjoint condition, this process could have allowed for associative links to be formed between temporally distant memories, facilitating an integrated map representation and ultimately leading to superior 2-step navigation.

That being said, we report two observations that might be considered incompatible with this explanation. Firstly, if grouped training encouraged the formation of disconnected “islands” of knowledge corresponding to grouped states within rows and columns, then we might expect test trials that required navigation within a single row or column to be more effective in this case. Surprisingly, however, even the straight navigation trials (which involve movement along two successive states of a single row or column) were performed better in the disjoint condition, even if this effect was significantly weaker than for diagonal trials. Secondly, if rows and columns forming unconnected fragments is the culprit, then we might expect performance to be disproportionately worse on grouped training in the blocked condition (where they are learned in two entirely distinct halves of the training phase) relative to the semi-blocked condition (where they could in theory be encountered in adjacent mini-blocks). However, we saw no interaction between curriculum and blocking schedule in our study. Thus, we advance the “fragments” hypothesis somewhat tentatively.

A less interesting explanation for the advantage of disjoint training is that training in the grouped condition was simply easier as responses on successive trials built on each other, causing participants to try less hard or disengage during initial learning. We cannot rule out this explanation entirely, but point out that we did not observe any differences in training or test accuracy between a control version of this curriculum where response repetition was possible as opposed to impossible (while keeping adjacent transitions close together in time, **Supplementary Results 3**) as this explanation would predict if it were acting in isolation.

The results reported here, in which grouping in space and time seems to hurt (rather than help) learning, are consistent with a long tradition showing that in applied domains (e.g. educational settings, memorisation or motor learning), mixing different exemplars over time (“interleaving”) is beneficial to knowledge or skill acquisition (Goode & Magill, 1986 [↗](#); Rohrer & Taylor, 2007 [↗](#); Simon & Bjork, 2001 [↗](#); Taylor & Rohrer, 2010 [↗](#)). These reports are consistent with established theories, such as complementary learning systems theory, which propose that mixing associations through mental rehearsal (or replay) allows us to form integrated knowledge structures. This benefit of mixing is in the spirit of the explanation provided above. However, the findings do contradict recent reports from the category learning literature, where blocking seems to help learning and generalisation (Dekker et al., 2022 [↗](#); Flesch et al., 2018 [↗](#); Noh et al., 2016 [↗](#)). It may be that where the goal is not to learn a complex knowledge structure – like a map – but simply to compress exemplars by mapping them onto a smaller number of labels – the benefits of blocking emerge. Indeed, one recent paper has argued (in simulation) that blocking should help learn generalisation more than memorisation (Russin, Pavlick, & Frank, 2025 [↗](#)).

We did not find any evidence for a significant effect of curriculum on participants’ ability to reconstruct the 2D maps using a drag and drop interface. This may have been because seeing all stimuli on screen allowed participants to use their knowledge of directly trained associations to integrate stimuli into a unified map based on logical reasoning and knowledge of the desired map shape. Unfortunately, we did not record the process of map reconstruction, otherwise the sequence in which participants reconstructed the map by moving the icons may have allowed us to test this directly.

In this paper, we have studied the impact that different orderings of training stimuli have on how cognitive maps are learned and found that training curriculum seems to have a large impact on the resulting cognitive map. Unlike in other learning tasks such as categorisation, not factorising associations into components such as rows or columns during initial exposure seems to be

beneficial for subsequent map integration and adaptive behaviour. Our findings have important implications for both the design of future research and for how we interpret existing research. When learning new concepts, maps, or relations, we need to think carefully about how we can optimise our learning and how our chosen curriculum might affect what we - or our participants - actually learn rather than what we want them to learn.

## Methods

### Design

In this section, we will provide an overview of the task design and manipulations. More detailed information such as trial numbers, trial sequences and task details will be provided in the task and procedure section below. We conducted a three-day experiment, over the course of which each participant had to learn two novel maps of image associations each (Figure 1b [↗](#), Figure 2a [↗](#)). The maps were grids (5×5) or tori (4×4) of states in which each state corresponded to a unique image with transitions between states in the up, down, left and right direction. Each of the two maps had a unique set of images assigned to it which either consisted of household items or fruit and vegetables (assigned randomly, Figure 1b [↗](#)). Participants learned these maps through trial-and-error from paired associate image associations on the first two days of the experiment (one map per day). On the third day, they completed a refreshed phase on both maps before being tested on their map knowledge. To assess the cognitive maps they had formed, we asked them to perform one and two-step navigation on the maps and make proximity judgements about sets of images. Both of these tasks required mental integration of the learned paired associates into a spatial cognitive map.

We conducted our experiment both using a 5×5 grid and a 4×4 torus to test whether any effects of our training manipulation (outlined below) held true across different geometric shapes and was not dependent on border or wall effects (see Figure 1c [↗](#) for illustration of how the border concern does not apply to the torus maps). N.B. that the number of transitions in each row or column are identical for a 5×5 grid and a 4×4 torus as the torus wraps around.

We manipulated the order in which participants experience the one-step transitions that make up the new cognitive maps to investigate how this affects their map learning and retention. For each of the map shapes, we employed a mixed 2×2 design (Figure 1b [↗](#)).

Our within-subjects manipulation pertains to the spatial sampling of transitions during training. Each participant learned one map in what we call the grouped curriculum and one map in the disjoint curriculum. In the grouped curriculum, the transitions experienced within each block of training were a pseudo-random walk along one row or one column of the map space (Figure 1a [↗](#)). The other map was learned using the disjoint curriculum, in which all transitions that made up the map were randomly shuffled and then divided into blocks of 8 transitions, meaning that each block of training subsequently consisted of 8 randomly sampled unique transitions from across the map (Figure 1a [↗](#)). Once assigned, which 8 transitions made up each unique block of disjoint training trials stayed consistent throughout the experiment to match the grouped training.

The between-subjects manipulation pertained to the temporal spacing of repetitions of transitions across blocks of training (Figure 1c [↗](#)). One group of participants per shape (grid, torus) would experience the blocks of training for both maps in a blocked fashion. In the grouped condition, this meant that all row or column blocks were experienced twice over before encountering the other dimension blocks twice over. It was introduced to further separate blocks of trials that may cause interference in time (see *Results*). In the disjoint condition, it meant seeing the first five (torus: 4) blocks twice over before experiencing the second five (torus: 4) blocks twice over. The other group of participants would experience the blocks of training in an interleaved fashion, which meant in the grouped condition that row and column blocks were interleaved until all blocks had been repeated twice and in the disjoint condition that the first five (torus: 4) and last five (torus: 4) blocks were interleaved and presented twice.

## Task and procedure

We will explain the tasks making up the three-day experiment chronologically. The first two days and the first half of the third day consisted of the training phase and the second half of the third day was the test phase. Before starting the experiment on day 1, participants familiarised themselves with the controls of the task. Instead of using arrow keys, we asked participants to use the ASDF keys to indicate their responses. In this pre-experiment phase, participants had to learn their unique mapping of the ASDF keys onto the four directions left, right, up and down. This mapping stayed consistent throughout the task and was indicated using geometric shapes that remained on-screen throughout the task. In the familiarisation phase, one of the geometric shapes would appear on its associated square (left, right, up or down) and participants had to press the correct key out of ASDF to select the indicated direction. Participants performed this familiarisation phase until they had reached 90% accuracy on at least 26 trials.

### Training phase

After completing the familiarisation phase, participants started the training phase. Participants learned one map on day 1 and the other map on day 2. The learning order of the grouped and disjoint maps was randomized across participants. On each day of training, participants in the grid groups completed 640 (torus: 512) training trials, split into 20 blocks of 32 trials (torus: 16 blocks of 32 trials). Each block of 32 trials consisted of 8 unique paired associates, repeated four times each. Each block of training trials was repeated twice over the course of training.

There were two different types of training trials. The layout was similar across these trials, but participants were required to make different decisions. In both cases, one stimulus (referred to as the first stimulus hereafter) from a given paired associate was presented at the centre of the screen, surrounded by four boxes corresponding to the four possible adjacent locations (Figure 2b [↗](#)).

In the first type of trial, referred to as select-action trials (Figure 2b [↗](#)), participants were shown the other stimulus (referred to as the second stimulus hereafter) from the associated pair at the bottom of the screen and were required to indicate its relative position with respect to the centrally presented stimulus (left vs. right or up vs. down). These trials comprised three out of four repetitions of each paired associate within a block. Feedback was provided via a coloured frame: if the response was correct, the second stimulus appeared in the selected location with a green frame; if the response was incorrect, the second stimulus appeared in the correct location, while the incorrectly selected location was highlighted with a red frame (Figure 2b [↗](#)).

The second type of trial, referred to as the select-stimulus trials (Supplementary Results 1 [↗](#)), was introduced to increase participant engagement. Here, one location adjacent to the first stimulus (centrally presented) was highlighted, and participants were required to select which of two stimuli belonged in that location. If the response was correct, the correct stimulus appeared in the highlighted location with a green frame; if the response was incorrect, the correct stimulus appeared in the location, while the incorrectly selected stimulus was highlighted with a red frame (Supplementary Figure 1 [↗](#)). These trials comprised the remaining one out of four repetitions per paired associate within each block.

Participants had up to 10s to indicate their choice on both types of training trial. This long timeout was mostly implemented to ensure people did not leave their desk for prolonged periods of time whilst completing the task at home. Positive feedback was shown for 2s and negative feedback was shown for 2.7s. Participants were shown their accuracy on the previous block in-between blocks and were allowed to take breaks in-between blocks. In the grouped condition, the second stimulus on trial  $t-1$  became the central stimulus on trial  $t$ . In the disjoint condition, the central stimulus on trial  $t$  and second stimulus on trial  $t-1$  tended to be non-adjacent/unrelated. Days 1 and 2 took approximately 90 minutes each.

After having learned the two maps on days 1 and 2 respectively, participants completed a shorter refresher of both maps on day 3. The refresher was identical to day 1 and 2 training, but instead of experiencing each block twice as during training (e.g. all row blocks twice, followed by all column

blocks twice in the blocked condition), participants only experienced each block once (e.g. all row blocks once, followed by all column blocks once in the blocked condition), meaning each participant completed 10 (torus: 8) blocks of 32 trials per map. This refresher also took approximately 90 minutes (45 minutes per map). The order in which the maps were experienced in the refresher was the reverse of the initial training order, i.e. the map learned on day 2 was revisited first followed by the map learned on day 1.

## Test phase

For the second half of day 3, which we consider the test phase, we assessed participants' map knowledge using in three (torus: two) ways (Figure 2c-e [↗](#)): navigation test, proximity test and map reconstruction test. Participants completed these three tests in chronological order. The torus groups only completed the navigation test and the proximity test.

### Navigation test

Participants were asked to navigate to goals that were either one or two steps away from a starting location. The one-step trials served as retention checks, as they resembled the selectaction trials during training except that participant chose among all possible actions instead of two. There were two types of two-step trials: *straight* trials, in which the start and goal locations came from the same row or column, and *diagonal* trials, in which they did not. It was not communicated to participants whether a two-step trial was a straight or diagonal trial. Each navigation block consisted of 25 two-step trials (12 straight and 13 diagonal) and 5 one-step trials presented in random order. Participants completed six navigation blocks in total, alternating between maps (60 trials of navigation per map, split into 2 blocks of 30).

On each navigation trial, participants would first see the goal stimulus with text underneath it indicating how many steps away from the start it was. After 500ms, the starting location would appear in the centre of the screen, surrounded by the four choice options. In the grid groups, if the starting location was by the edge of the grid, a wall icon would appear in the location(s) where the starting stimulus was wall-adjacent, and the affected locations could not be selected. Participants then had to take either one or two steps to reach the goal. They had 12s to make a choice and used the same controls as during select-action trials in training. As during training, they also had to first select and then confirm a choice (repeat keypress). As participants made choices, the central stimulus (current location) changed according to their choices and the counter with choices made and choices remaining updated accordingly. The trial terminated once the minimum number of steps required had been taken and participants could see whether they had reached the goal correctly or not by comparing the final location to the goal. Their final choice remained on-screen for 1.5s.

### Proximity test

After completing the navigation test, participants commenced the second part of the test phase. In this part, they were asked to make proximity judgements to assess how well they had learned the two maps. On each trial, they were first shown a central stimulus (*target*) and then asked to indicate which of two other stimuli (*probes*) was closer to the central stimulus on the currently visible map (Figure 2f [↗](#)). Each probe had a defined distance (i.e., number of steps) from the target, resulting in a comparison between two distances on every trial. Participants were allowed a maximum of 12s to respond and no feedback was provided. Trials were split into four types according to the specific distance comparison involved, and each block had:

- 9 trials of 2-away (straight) vs 1-away
- 10 trials of 2-away (diagonal) vs 1-away
- 10 trials of 2-away (diagonal) vs 3-away (diagonal)
- 11 trials of 2-away (diagonal) vs 2-away (straight)

In the 2-way (straight) vs 1-away trials, for example, one probe was within the same row or column as the target and was two steps away, while the other was one step away from the target. All were sampled randomly from all possible comparisons and presented randomly.

### Map reconstruction test

The grid groups were asked to reconstruct the two maps by using their mouse and a drag and drop interface and we calculated a reconstruction accuracy score for each map.

Participants were incentivised with a bonus of 2.5p for each goal reached during the navigation phase and 2.5p for each correct proximity judgement (apart from the 2vs2 trials where there was no objectively correct answer), capped at a maximum bonus payment of £6.

## Participants

A total of 182 participants (grid: 49 blocked, 43 interleaved; torus: 46 blocked, 44 interleaved) were recruited using Prolific Academic (<https://www.prolific.com/>) in 2022 and early 2023. All participants self-reported that they were between 19 and 40 years of age and resided in the United Kingdom. The mean age of all participants was 30.03 (SD: 5.88) and 80 identified as female. To ensure only motivated individuals completed our main task, we only invited participants who had previously passed a performance criterion on a brief attention task. The attention task was a short 6 minute 2-back task in which participants were instructed to press the space bar whenever the symbol presented two trials ago matched the symbol currently on the screen. Participants who detected at least 21/30 repetitions and had less than 10 false alarms on the 2-back task were invited to participate in the multi-day map learning experiment. Participants were compensated for their time at a rate of £10/hour, both for the attention and the main tasks. They were able to earn a performance-based bonus of up to £6 on the main task. All online experiments were approved by the Medical Sciences Research Ethics Committee of the University of Oxford.

## Materials

25 clipart images of fruit and vegetables and 25 clipart images of common household items and furniture available under the common use license were taken from Google images and edited to have the same dimensions and background using Microsoft PowerPoint. 16 images were randomly selected from each category for each participant in the torus groups and all 25 images from each category were used for each participant in the grid groups. All stimuli used are displayed in [Supplementary Figure 2](#), examples are displayed in [Figure 1b](#).

## Software

All experiments in this chapter were implemented in JavaScript (supported by php and html code) by modifying a custom lab-internal toolbox for running online experiments. The training curricula were pre-generated using Matlab 2019b. All statistical analyses were carried out in Python using the Pandas 1.2.3, NumPy 1.19, SciPy 1.60, Statsmodels 0.13 and Scikit-Learn 0.24.1 packages. All data figures were generated using Matplotlib 3.3.2 and Seaborn 0.12.

## Statistical analyses

The majority of analyses presented in this paper were conducted using mixed linear models in the statsmodels.MixedLM package for python. All models had participants as the grouping variable. Most models used had the following format, with Y being the outcome variable of interest (usually accuracy)

$$Y \sim \text{temporal manipulation} * \text{spatial manipulation} + \text{shape} + (1 | \text{participant}) \quad (1)$$

Exceptions to this format were the following analyses.

The analysis conducted to determine whether there was an interaction between initial training or refresher (train feature) and our spatial manipulation was specified as

$$\text{accuracy} \sim \text{train feature} * \text{spatial manipulation} + \text{temporal manipulation} + \text{shape} + (1 | \text{participant}) \quad (2)$$

$$\text{accuracy} \sim \text{temporal manipulation} * \text{spatial manipulation} \\ * \text{trial type} + \text{shape} + (1 | \text{participant}) \quad (3)$$

The analysis querying whether participants' overall test accuracy increased from the first to the second test block was specified as

$$\text{accuracy} \sim \text{temporal manipulation} * \text{spatial manipulation} \\ * \text{block} + \text{shape} + (1 | \text{participant}) \quad (4)$$

The normality of residuals was assessed for all analyses using the `scipy.stats.normaltest` function and if the residuals were non-normal at an alpha level of .01, a permutation test was conducted to determine the appropriate p-value. Where this is the case, p-values were denoted as  $p_{\text{perm}}$ . We used 1000 permutations and ran separate permutations for the different factors in the affected models. Where interaction effects were significant, the permutation was only applied to one of the two factors making up the interaction.

Map reconstruction scores for the grid groups were computed by first calculating a representational dissimilarity matrix (RDM) consisting of the Euclidean distance of each stimulus to each other stimulus in the reconstruction and then computing the spearman rank correlation score between the participant RDM and the actual veridical distance RDM. Fisher's z transformed correlation scores were entered into a mixed effects model specified as

$$r \sim \text{temporal manipulation} * \text{spatial manipulation} + (1 | \text{participant}) \quad (5)$$

## Additional information

### Instructions

On days 1 and 2 and for the refresher on day 3, participants were instructed that they were travellers who had landed on an alien planet. To prove their goodwill, the aliens had tasked them with learning the layout of two of their planet's moons (doughnut shaped moons in the torus groups and square fields lined with walls in the grid groups). The atmosphere on the moons was foggy, and thus participants would have to learn the position of objects relative to each other. Participants were informed that they would have to remember the learned information as they would be tested on it by the aliens on day 3 (but not told how) and that good performance on day 3 would be rewarded. They were shown an example trial sequence for training and asked some control questions about a small example torus in the torus groups to ensure they understood the concept of a torus shaped space. Before commencing the test phase on day 3, participants were informed that the aliens now wanted to test their knowledge and shown example trial sequences (with different sets of stimuli) for both test tasks as well as informed of the bonus payment that could be gained.

### Post-task questionnaire

After completing each day of the task, participants were asked to indicate whether they had used any visual aids or written down any of the answers as well as which map(s) they had seen.

### Curriculum sampling

Training curricula were pre-generated to ensure equal coverage of all paired associates in a row or column in the grouped condition. For ease of implementation, disjoint curricula were also pre-generated. For each participant, a pre-generated trial sequence for training and refresher was selected randomly out of 50 pre-generated sequences in the grid groups and 100 pre-generated sequences in the torus groups.

## Supplementary Materials

### Supplementary Results 1: select action training trials

The variable dividing training trials into select-state and select-action trials was not saved due to a technical error for 17 of the torus group participants, which means that these results are based on observations from 165 participants.

#### Select-state trials

Participants' training accuracy was significantly greater on the map that they learned using grouped than that learned with disjoint training ( $\beta = -.09$ ,  $p < .0001$ ). There was neither significant effect of the blocking manipulation on training accuracy nor a significant interaction of the blocking and spatial manipulation (both  $p > .05$ ). When splitting this up into initial training and refresher, we found the same pattern of results, with participants' training accuracy being significantly higher on their grouped than disjoint map (train:  $\beta = -.10$ ,  $p < .0001$ ; refresh:  $\beta = -.06$ ,  $p < .0001$ ) and no significant effect of the blocked manipulation nor an interaction (all  $ps > .05$ ).

#### Select-action trials


Participants' training accuracy was significantly greater on the map that they learned using grouped than that learned with disjoint training ( $\beta = -.08$ ,  $p < .0001$ ). There was neither a significant effect of the temporal manipulation on training accuracy nor a significant interaction of the temporal and spatial manipulation (both  $p > .05$ ). When splitting this up into initial training and refresher, we found the same pattern of results, with participants' training accuracy being significantly higher on their grouped than disjoint map (train:  $\beta = -.09$ ,  $p < .0001$ ; refresh:  $\beta = -.05$ ,  $p < .0001$ ) and no significant effect of our temporal manipulation nor an interaction (all  $ps > .05$ ).

### Supplementary Results 2: Results of the control experiment shuffled grouped vs continuous grouped

This control experiment was preregistered on OSF under this link [https://osf.io/x8ju3/?view\\_only=df19ad57d03f40f098b5724bbfafb7fa](https://osf.io/x8ju3/?view_only=df19ad57d03f40f098b5724bbfafb7fa). In brief, we preregistered and conducted a control experiment to assess whether differences in training accuracy (and the inference abilities) were indeed simply due to a reduced need for attention. Participants ( $N = 63$ ) in the control condition also learned two maps, but this time both of them were spatially grouped. The difference between the two maps was that in one condition the transitions in a row or column were experienced as a pseudo-random walk (like in the main experiment), and in the other they were experienced in a shuffled manner. The latter meant that temporally adjacent transitions were not necessarily spatially adjacent. We hypothesised that we should see lower training accuracy in the shuffled grouped compared to the pseudo-random walk grouped map if the reason for the high training accuracy was the reduced need for attention rather than our grouping manipulation.

We did not find evidence for this, with there being no difference in training accuracy between the two curricula ( $\beta = -.01$ ,  $SE = .01$ ,  $z = -1.34$ ,  $p = .18$ ). To make sure that this was not due to cohort differences, we compared the training accuracy on both shuffled and continuous grouped to the grouped map training accuracy in the original blocked grouped cohort ( $N = 49$ ). There was no difference between the continuous grouped condition accuracy in the control experiment and the original blocked grouped grid training accuracy ( $\beta = -.29$ ,  $z = -1.39$ ,  $SE = .20$ ,  $p = .16$ ). When compared to the whole behaviour-only cohort, there was also a significant interaction, with the difference between disjoint and grouped being significantly greater than between grouped continuous and grouped shuffled ( $\beta = .07$ ,  $z = 5.92$ ,  $SE = .01$ ,  $p < .001$ ).

If the disadvantage of the original grouped cohort was due to a reduced need for attention due to the trajectoriness of the condition, we would also expect to see higher inference accuracy in the shuffled condition. We do not see evidence for this in either 2-step navigation (diagonal:  $\beta = .03$ ,  $z = 1.5$ ,  $SE = .02$ ,  $p = .14$ ; straight:  $\beta = -.00$ ,  $z = -.17$ ,  $SE = .02$ ,  $p = .87$ ) and neither 2-step navigation accuracy being significantly different from the matched blocked grouped grid-shape cohort



**Supplementary Figure 1. Example of a select-stimulus trial**

sequence in which the incorrect option was selected. Had the correct option been selected, the washing machine would have appeared with a green frame around it and the kettle would have disappeared at the time of feedback.

(straight:  $\beta = -.23$ ,  $z = -.85$ ,  $SE = .27$ ,  $p = .40$ ; diag:  $\beta = -.22$ ,  $z = -.78$ ,  $SE = .29$ ,  $p = .44$ ). Comparing to the whole original cohort, we also find a significant interaction for diagonal two-steps, with the difference between disjoint and grouped accuracy being significantly greater than between shuffled grouped and grouped accuracy ( $\beta = -.08$ ,  $z = -2.63$ ,  $SE = .03$ ,  $p = .008$ ). We also do not find evidence for a significant difference between the shuffled grouped and original grouped condition in terms of proximity judgements (2v1 straight:  $\beta = .02$ ,  $z = .73$ ,  $SE = .02$ ,  $p = .46$ ; 2v1 diagonal:  $\beta = .01$ ,  $z = .28$ ,  $SE = .02$ ,  $p = .78$ ; 2v3:  $\beta = .03$ ,  $z = 1.19$ ,  $SE = .23$ ,  $p = .23$ ), nor are the continuous grouped proximity judgement accuracies significantly different to the original blocked grouped grid accuracy (2v1 straight:  $\beta = -.28$ ,  $z = -1.11$ ,  $SE = .25$ ,  $p = .27$ ; 2v1 diagonal:  $\beta = -.27$ ,  $z = -1.06$ ,  $SE = .26$ ,  $p = .29$ ; 2v3:  $\beta = -.24$ ,  $z = -.90$ ,  $SE = .26$ ,  $p = .37$ ). When comparing overall average proximity judgement accuracy in the control experiment to proximity judgement accuracy in the whole behaviour-only dataset, we also see a significant interaction, with the difference between disjoint and grouped being significantly greater than between shuffled grouped and original grouped ( $\beta = -.05$ ,  $z = -2.67$ ,  $SE = .02$ ,  $p = .008$ ).

While it is challenging to draw inferences based on null results, the combination of no difference between the shuffled and original grouped curriculum (with its accuracy being the same as the matched original data) and a significant interaction effect with our our main experiment curriculum effect suggests that a lack of a need to pay attention is not the main driver behind the worse inference performance in the grouped than disjoint condition

**Table 1. Full table of results for mixed effects model results reported in the main text + training accuracy split into initial training and refresher**

All models in this table were specified as follows, with Y being the outcome variable of interest, cohort our between-subjects manipulation (blocked (1) vs semi-blocked (0)), condition the within-subjects manipulation (grouped 0 vs disjoint 1) and shape being torus (1) or grid (0).  $Y \sim \text{cohort} * \text{condition} + \text{shape} + (1 | \text{participant})$

Y	Training accuracy overall			Training accuracy initial training			Training accuracy refresher		
	Coef.	z	p> z	Coef.	z	p> z	Coef.	z	p> z
Intercept	.95	75.66	0.000	.94	73.78	0.000	.96	70.87	0.000
Cohort	-.01	-.98	0.327	-.02	-1.01	0.311	-.01	-.91	0.362
Condition	-.08	-9.32	0.000	-.09	-9.69	0.000	-.05	-5.72	0.000
Cohort:condition	.01	0.46	0.647	.01	0.79	0.430	-.01	-.55	0.584
Shape	.01	-4.44	0.000	-.06	-4.06	0.000	-.05	-3.42	0.001
Y	One step navigation accuracy			Diagonal 2 step navigation (corner/L)			Straight 2 step navigation (straight/I)		
	Coef.	z	p> z	Coef.	z	p> z	Coef.	z	p> z
Intercept	.77	25.43	0.000	.64	17.01	0.000	.69	19.20	0.000
Cohort	-.06	-1.58	0.114	-.05	-1.02	0.309	-.05	-1.08	0.282
Condition	.06	2.62	0.009	.13	5.07	0.000	.05	2.20	0.028
Cohort:condition	.03	0.92	0.357	-.03	-.91	0.361	-.01	-.40	0.686
Shape	-.15	-4.90	0.000	-.23	-5.76	0.000	-.09	-2.40	0.018
Y	Proximity: 2 v 1 diag			Proximity:2v1 straight			Proximity: 2v3		
	Coef.	z	p> z	Coef.	z	p> z	Coef.	z	p> z
Intercept	.73	35.02	0.000	.73	34.09	0.000	0.700	31.41	0.000
Cohort	-.03	-1.10	0.271	-.01	-.54	0.590	-.00	-.02	0.984
Condition	.03	1.50	0.133	.08	3.98	0.000	.04	2.09	0.037
Cohort:condition	.04	1.57	0.117	.01	0.37	0.714	.03	1.07	0.286
Shape	-.02	-.93	0.351	-.08	-3.70	0.000	-.10	-4.18	0.000

**Table 2. Main text results split by torus and grid cohorts**

*Y ~ cohort \* condition + (1 | participant)*

Y : overall training accuracy	GRID			TORUS		
	Coef.	z	p> z	Coef.	z	p> z
Intercept	.94	78.50	0.000	0.90	51.38	0.000
Cohort	-.01	-.66	0.512	-0.09	-6.07	0.000
Condition	-.08	-7.81	0.000	-0.02	-0.72	0.473
Cohort:condition	.01	.99	0.321	-0.00	-0.15	0.884
Y : overall two-step accuracy	GRID			TORUS		
	Coef.	z	p> z	Coef.	z	p> z
Intercept	.66	16.95	0.000	0.51	11.53	0.000
Condition	.04	1.97	0.049	0.14	3.99	0.000
Cohort	-.01	-.09	0.925	-0.09	-1.41	0.159
Cohort:condition	.01	0.24	0.813	-0.05	-1.06	0.289
Y : diagonal two-step accuracy	GRID			TORUS		
	Coef.	z	p> z	Coef.	z	p> z
Intercept	.63	15.51	0.000	0.42	8.67	0.000
Condition	.09	3.05	0.002	0.17	4.08	0.000
Cohort	.01	0.18	0.855	-0.10	-1.51	0.132
Cohort:condition	-.03	-0.69	0.489	-0.03	-0.58	0.562
Y: straight two-step accuracy	GRID			TORUS		
	Coef.	z	p> z	Coef.	z	p> z
Intercept	0.687	16.83	0.000	0.60	13.64	0.000
Condition	-0.005	-0.17	0.862	0.10	3.01	0.003
Cohort	-0.021	-0.38	0.702	-0.07	-1.13	0.257
Cohort:condition	0.045	1.18	0.240	-0.07	-1.47	0.141
Y: overall proximity accuracy	GRID			TORUS		
	Coef.	z	p> z	Coef.	z	p> z
Intercept	0.714	31.63	0.000	0.66	30.97	0.000
Condition	0.02	1.12	0.262	0.08	4.17	0.000
Cohort	0.02	0.73	0.463	-0.05	-1.71	0.087
Cohort:condition	0.04	1.62	0.105	0.02	0.57	0.565

**Table 3. Additional analyses reported in the main text**

The analysis conducted to determine whether there was an interaction between type of two-step navigation trial and our spatial manipulation was specified as *accuracy ~ cohort \* condition\*trial type + shape + (1 | participant)* The analysis querying whether participants' overall test accuracy increased from the first to the second test block was specified as *accuracy ~ cohort \* condition\*test block + shape + (1 | participant)* The results from these analyses are reported in the table below.

Y	Two-step accuracy on the different types of two-step trials			Overall test accuracy (one and two-steps)			
	Coef.	z	p> z	Coef.	z	p> z	
Intercept	.60	16.53	0.000	Intercept	.66	18.85	0.000
Trial type	.12	5.03	0.000	Test block	.02	0.62	0.535
condition	.13	5.42	0.000	condition	.09	2.47	0.014
Trial type : condition	-.08	-2.39	0.017	Test block : condition	.00	0.09	0.932

cohort	-0.04	-1.00	0.317	cohort	-0.07	-1.57	0.117
Trial type : cohort	-0.00	-0.10	0.923	Test block : cohort	.04	0.86	0.389
condition: cohort	-0.03	-0.98	0.329	Condition : cohort	-0.01	-0.22	0.828
Trial type : condition : cohort	.02	.42	0.671	Test block : condition : cohort	-0.02	-0.34	0.733
shape	-0.16	-4.24	0.000	shape	-0.17	-9.46	0.000

**Table 4. Pairwise comparisons (grouped vs disjoint) for the four experimental groups for the main text results**

Task	Shape	Group	t	df	p
Training accuracy overall	grid	blocked	7.30	48	0.000
	grid	semiblocked	7.35	42	0.000
	torus	blocked	6.17	45	0.000
	torus	semiblocked	6.34	43	0.000
Training accuracy initial training	grid	blocked	6.93	48	0.000
	grid	semiblocked	7.45	42	0.000
	torus	blocked	6.18	45	0.000
	torus	semiblocked	5.70	43	0.000
Training accuracy refresher	grid	blocked	6.12	48	0.000
	grid	semiblocked	4.35	42	0.000
	torus	blocked	4.15	45	0.000
	torus	semiblocked	4.93	43	0.000
One-step navigation accuracy	grid	blocked	-2.02	48	0.049
	grid	semiblocked	-0.30	42	0.763
	torus	blocked	-3.44	45	0.001
	torus	semiblocked	-3.73	43	0.001
Overall two-step navigation accuracy	grid	blocked	-2.53	48	0.015
	grid	semiblocked	-1.89	42	0.065
	torus	blocked	-2.37	45	0.022
	torus	semiblocked	-4.38	43	0.000
Diagonal two-step accuracy	grid	blocked	-2.33	48	0.024
	grid	semiblocked	-2.93	42	0.005
	torus	blocked	-3.20	45	0.000
	torus	semiblocked	-4.28	43	0.000
Straight two-step accuracy	grid	blocked	-1.77	48	0.082
	grid	semiblocked	0.15	42	0.879
	torus	blocked	-0.89	45	0.378
	torus	semiblocked	-3.36	43	0.002
Overall proximity judgements	grid	blocked	-3.66	48	0.001
	grid	semiblocked	-1.09	42	0.281
	torus	blocked	-4.92	45	0.000
	torus	semiblocked	-4.32	43	0.000
2v3 diagonal	grid	blocked	-2.11	48	0.040
	grid	semiblocked	-0.88	42	0.386
	torus	blocked	-2.79	45	0.008
	torus	semiblocked	-2.19	43	0.034

2v1 diag	grid	blocked	-1.91	48	0.062
	grid	semiblocked	0.44	42	0.661
	torus	blocked	-4.13	45	0.000
	torus	semiblocked	-2.49	43	0.017
2v1 straight	grid	blocked	-2.79	48	0.007
	grid	semiblocked	-1.81	42	0.077
	torus	blocked	-3.92	45	0.000
	torus	semiblocked	-3.73	43	0.001
2v2	grid	blocked	0.67	48	0.504
	grid	semiblocked	0.08	42	0.935
	torus	blocked	-0.37	45	0.715
	torus	semiblocked	-0.11	43	0.910
Map reconstruction	grid	blocked	-1.50	48	0.141
	grid	semiblocked	-1.00	42	0.322

**Table 5. All pairwise comparisons**

Shape	Condition	Group	Type of 2-step	t	P (Bonferroni corrected)
Torus	Grouped	Blocked	Diag	5.16	0.00013
Torus	Grouped	Blocked	Straight	10.35	0.00000
Torus	Disjoint	Blocked	Diag	6.52	0.00000
Torus	Disjoint	Blocked	Straight	9.38	0.00000
Torus	Grouped	Semi-blocked	Diag	5.46	0.00005
Torus	Grouped	Semi-blocked	Straight	10.28	0.00000
Torus	Disjoint	Semi-blocked	Diag	9.38	0.00000
Torus	Disjoint	Semi-blocked	Straight	13.76	0.00000
Grid	Grouped	Blocked	Diag	14.22	0.00000
Grid	Grouped	Blocked	Straight	14.76	0.00000
Grid	Disjoint	Blocked	Diag	13.88	0.00000
Grid	Disjoint	Blocked	Straight	15.49	0.00000
Grid	Grouped	Semi-blocked	Diag	11.98	0.00000
Grid	Grouped	Semi-blocked	Straight	20.09	0.00000
Grid	Disjoint	Semi-blocked	Diag	16.20	0.00000
Grid	Disjoint	Semi-blocked	Straight	14.80	0.00000
Shape	Condition	Group	One-step	t	P (Bonferroni corrected)
Torus	Grouped	Blocked	One-step	7.72	0.00000
Torus	Disjoint	Blocked	One-step	10.14	0.00000
Torus	Grouped	Semi-blocked	One-step	8.50	0.00000
Torus	Disjoint	Semi-blocked	One-step	12.43	0.00000
Grid	Grouped	Blocked	One-step	14.12	0.00000
Grid	Disjoint	Blocked	One-step	17.36	0.00000
Grid	Grouped	Semi-blocked	One-step	17.83	0.00000
Grid	Disjoint	Semi-blocked	One-step	17.30	0.00000

Table 5. One-sample t-test statistics comparing the two-step accuracy of different groups and conditions to chance. We compared to chance (Grid: .125 for diagonal two-steps and .0625 for straight two-steps, .25 for one-steps; torus: .125 for diagonal and .125 for straight two-steps; .25 for one-steps) using a one-sample t-test.

Supplementary Figure 2. All stimuli used



## Data availability

Code for the analyses presented in this paper is available at [https://github.com/LeonieGli/Glitz\\_et\\_al\\_2026](https://github.com/LeonieGli/Glitz_et_al_2026) [↗](#). Data is available at <https://osf.io/592dj/overview> [↗](#).

## Acknowledgements

This research was funded by a Wellcome Trust Discovery Award (<https://wellcome.org/grant-funding/schemes/discovery-awards> [↗](#)) (227928/Z/23/Z) to C.S., also supporting Z.L. It was also funded by UKRI (MR/W008939/1 to H.C.B.) and MRC (MR/W01971X/1 supporting L.G.). The Medical Research Council Centre of Research Excellence in Restorative Neural Dynamics is supported by MRC (UKRI 936). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

## Additional information

### Author contributions

LG and CS designed research; LG collected data; LG analysed and visualised data; LG wrote the initial draft; CS, HCB, ZL and LG edited and revised the draft; CS and HCB supervised the project

### Funding

Funder	Grant reference number	Author
Wellcome Trust (WT)	<a href="https://doi.org/10.35802/227928">https://doi.org/10.35802/227928</a>	Christopher Summerfield
UK Research and Innovation (UKRI)	MR/W008939/1	Helen C Barron
UKRI   Medical Research Council (MRC)	MR/W01971X/1	Helen C Barron

## References

- Behrens T. E., Muller T. H., Whittington J. C., Mark S., Baram A. B., Stachenfeld K. L., Kurth-Nelson Z. (2018) What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* **100**:490-509 <https://doi.org/10.1016/j.neuron.2018.10.002> | [PubMed](#)
- Bellmund J. L. S., Deuker L., Navarro Schröder T., Doeller C. F. (2016) Grid-cell representations in mental simulation. *eLife* **5**:e17089 <https://doi.org/10.7554/eLife.17089> | [PubMed](#)
- Bellmund J. L. S., Gärdenfors P., Moser E. I., Doeller C. F. (2018) Navigating cognition: Spatial codes for human thinking. *Science* **362**:eaat6766 <https://doi.org/10.1126/science.aat6766> | [PubMed](#)
- Cai D. J., Aharoni D., Shuman T., Shobe J., Biane J., Song W., ... Silva A. J. (2016) A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* **534**:115-118 <https://doi.org/10.1038/nature17955> | [PubMed](#)
- Constantinescu A. O., O'Reilly J. X., Behrens T. E. J. (2016) Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**:1464-1468 <https://doi.org/10.1126/science.aaf0941> | [PubMed](#)
- Dekker R. B., Otto F., Summerfield C. (2022) Curriculum learning for human compositional generalization. *Proceedings of the National Academy of Sciences* **119**:e2205582119 <https://doi.org/10.1073/pnas.2205582119> | [PubMed](#)
- Flesch T., Balaguer J., Dekker R., Nili H., Summerfield C. (2018) Comparing continual task learning in minds and machines. *Proceedings of the National Academy of Sciences* **115**:E10313-E10322 <https://doi.org/10.1073/pnas.1800755115> | [PubMed](#)

- Ghandour K.**, Ohkawa N., Fung C. C. A., Asai H., Saitoh Y., Takekawa T., ... Inokuchi K. (2019) Orchestrated ensemble activities constitute a hippocampal memory engram. *Nature communications* **10**:2637 <https://doi.org/10.1038/s41467-019-10683-2> | PubMed
- Goode S.**, Magill R. A. (1986) Contextual interference effects in learning three badminton serves. *Research quarterly for exercise and sport* **57**:308-314 <https://doi.org/10.1080/02701367.1986.10608091>
- Josselyn S. A.**, Tonegawa S. (2020) Memory engrams: Recalling the past and imagining the future. *Science* **367**:eaaw4325 <https://doi.org/10.1126/science.aaw4325> | PubMed
- Krueger K. A.**, Dayan P. (2009) Flexible shaping: How learning in small steps helps. *Cognition* **110**:380-394 <https://doi.org/10.1016/j.cognition.2008.11.014> | PubMed
- Liu X.**, Ramirez S., Pang P. T., Puryear C. B., Govindarajan A., Deisseroth K., Tonegawa S. (2012) Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* **484**:381-385 <https://doi.org/10.1038/nature11028> | PubMed
- Mi Q.**, Summerfield C. (2025) Human curriculum learning of a cue combination task. *PsyArXiv* [https://doi.org/10.31234/osf.io/jr42g\\_v1](https://doi.org/10.31234/osf.io/jr42g_v1)
- Mitchell R.**, Martin C. (1997) Rote learning, creativity and 'understanding' in classroom foreign language teaching. *Language Teaching Research* **1**:1-27 <https://doi.org/10.1177/136216889700100102>
- Momennejad I.** (2020) Learning Structures: Predictive Representations, Replay, and Generalization. *Current Opinion in Behavioral Sciences* **32**:155-166 <https://doi.org/10.1016/j.cobeha.2020.02.017> | PubMed
- Noh S. M.**, Yan V. X., Bjork R. A., Maddox W. T. (2016) Optimal sequencing during category learning: Testing a dual-learning systems perspective. *Cognition* **155**:23-29 <https://doi.org/10.1016/j.cognition.2016.06.007> | PubMed
- Park S. A.**, Miller D. S., Boorman E. D. (2021) Inferences on a multidimensional social hierarchy use a grid-like code. *Nat Neurosci* **24**:1292-1301 <https://doi.org/10.1038/s41593-021-00916-3> | PubMed
- Rohrer D.**, Taylor K. (2007) The shuffling of mathematics problems improves learning. *Instructional Science* **35**:481-498 <https://doi.org/10.1007/s11251-007-9015-8>
- Russin J.**, Pavlick E., Frank M. J. (2025) Parallel trade-offs in human cognition and neural networks: The dynamic interplay between in-context and in-weight learning. *Proceedings of the National Academy of Sciences* **122**:e2510270122 <https://doi.org/10.1073/pnas.2510270122> | PubMed
- Schacter D. L.**, Loftus E. F. (2013) Memory and law: what can cognitive neuroscience contribute?. *Nature Neuroscience* **16**:119-123 <https://doi.org/10.1038/nn.3294> | PubMed
- Simon D. A.**, Bjork R. A. (2001) Metacognition in motor learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **27**:907 <https://doi.org/10.1037/0278-7393.27.4.907> | PubMed
- Stachenfeld K. L.**, Botvinick M. M., Gershman S. J. (2017) The hippocampus as a predictive map. *Nature Neuroscience* **20**:1643-1653 <https://doi.org/10.1038/nn.4650> | PubMed
- Taylor K.**, Rohrer D. (2010) The effects of interleaved practice. *Applied cognitive psychology* **24**:837-848 <https://doi.org/10.1002/acp.1598>
- Tolman E. C.** (1948) Cognitive maps in rats and men. *Psychological review* **55**:189 <https://doi.org/10.1037/h0061626> | PubMed
- Xu X.**, Padilla A. M. (2013) Using meaningful interpretation and chunking to enhance memory: The case of Chinese character learning. *Foreign Language Annals* **46**:402-422 <https://doi.org/10.1111/flan.12039>
- Yokose J.**, Okubo-Suzuki R., Nomoto M., Ohkawa N., Nishizono H., Suzuki A., ... Nagase M. (2017) Overlapping memory trace indispensable for linking, but not recalling, individual memories. *Science* **355**:398-403 <https://doi.org/10.1126/science.aal2690> | PubMed
- Glitz L.**, Liang Z., Barron H.C., Summerfield C. (2026) How does curriculum learning influence cognitive map formation?. OSF. ID 592dj <https://osf.io/592dj/overview>

## Peer reviews

### Reviewer #1 (Public review):

This paper investigates how different learning curricula influence the way that humans piece together directly experienced transitions into a broader cognitive map. When adjacent learning trials were grouped within rows or columns of the map, subsequent navigation through the map was weaker than when adjacent learning trials came from disjoint spaces in the map. The authors speculate that the grouped curriculum resulted in mental fragmentation that made navigation across space more difficult later on.

This is an interesting paradigm that contributes useful new findings in the domain of map learning to the growing literature on curriculum learning. The evidence for a difference between conditions is highly compelling, but, as the authors are very transparent in acknowledging in the Discussion, the evidence for their proposed mechanism - mental fragmentation under grouped learning - is somewhat weak. The study thus presents an intriguing empirical result but not an ironclad mechanistic account.

An alternative - by their account, "less interesting" - explanation is that grouped learning was easier because trials in close succession had overlapping elements, and so participants were not trying as hard or as engaged. There is a literature on spaced (as opposed to massed) learning being better for subsequent memory because it increases retrieval effort. It seems very plausible that this could be going on here, and the control experiment reported in the supplement would not help to rule this out. This literature deserves some discussion.

The Introduction focuses entirely on literature showing advantages in grouped over intermixed learning, setting that up as the most well-motivated expectation from the literature. Upon finding the opposite, the Discussion then mentions that interleaving has been found to be useful in "applied domains", but then returns to how surprising this is in light of recent findings in the category learning literature. But there is a substantial earlier literature on interleaved vs blocked curricula in category learning, very often finding advantages for interleaving. See, e.g., Carvalho & Goldstone, 2015, for a review. There is also a paper showing interleaving advantages in associative inference, Zhou et al., 2023, JEP:G, which is very relevant to several of the discussion section paragraphs. Thus, the treatment of the prior curriculum learning literature is currently sparse.

<https://doi.org/10.7554/eLife.111035.1.sa2>

### Reviewer #2 (Public review):

I think this paper is an excellent and timely contribution. It clearly shows that learning overlapping relationships in a disjoint training schedule (where the overlaps are not encountered close together in time) appears to aid the formation of an integrated associative memory structure (a cognitive map) and supports generalisation. I believe the methods are sound and the results are clear. I only have a couple of methodological questions that may not warrant any changes to the paper (or only very minor changes/additions):

(1) The mixed effects models did not include random slopes for the within-subject factors ("spatial manipulation" and "block"), and so the corresponding fixed effect inferences may be unsafe. Having said that, it is likely that including these slopes may not be warranted given their contribution to the model's fit. I recommend that the authors check this.

(2) The mixed effects models for accuracy appear to model average performance across trials rather than using a generalised linear model with a (e.g.) logit link function and the binomial distribution to characterise performance. I think this is a little sub-optimal, as the latter is

often more sensitive. Nonetheless, it is not in any way wrong; the results are clear enough as is, and there may be a good reason to avoid a non-linear link function, which can alter the interpretation of effects close to the ceiling and floor.

I think the introduction and/or discussion would benefit from contrasting their results with Berens & Bird (2022, PLOS Comp Bio). In this paper, it is shown that blocking the training of discriminations in a linear hierarchy (what we call progressive training) substantially benefited transitive inference performance. This seems at odds with the author's finding that "participants struggle to integrate information across rows and columns, i.e. across groups of transitions that were trained separately in time".

I would really like to know what the authors think about this discrepancy (or, indeed, whether they think there is one at all). Is it possibly because "progressive" learning is some combination of "grouping", "blocking" and "chaining" (where there is a structured overlap between adjacently trained relationships)? Or is it something else, e.g., that there is a fundamental difference between learning associations and discriminations (personally, I lean on this explanation)?

Relevant to this, the authors note that their "findings do contradict recent reports from the category learning literature, where blocking seems to help learning and generalisation (Dekker et al., 2022; Flesch et al., 2018; Noh et al., 2016). It may be that where the goal is not to learn a complex knowledge structure - like a map - but simply to compress exemplars by mapping them onto a smaller number of labels - the benefits of blocking emerge." However, the benefit of progressive (blocked) training in my own work was observed in a task that required learning a complex/relational structure in the form of a transitive hierarchy, which theoretical accounts suggest depends on learning map-like representations (Whittington et al., 2020).

<https://doi.org/10.7554/eLife.111035.1.sa1>

### Reviewer #3 (Public review):

Summary:

This study examines how training regimes influence the formation of cognitive maps. Participants learned two relational maps over three days through pairwise transitions: one map was trained with grouped sequences that followed rows or columns, while the other was trained with disjoint transitions sampled randomly across the map. In addition, the study manipulated the temporal spacing of training blocks (blocked vs. semi-blocked) and tested whether the results generalized across two map geometries (a 5×5 grid and a 4×4 torus).

Furthermore, they run a follow-up experiment (or condition) testing rows and columns shuffled in the grouped condition.

While grouped training produced better performance during learning, the authors report that disjoint training led to superior performance at test on tasks probing the global map knowledge.

Summarising experimental design:

- (1) Map geometry (between-subjects): 5×5 grid vs 4×4 torus
- (2) Training block schedule (between-subjects): Blocked vs Semi-blocked
- (3) Training regime/transition sampling (within-subject): Grouped or Disjoint (Day 1 and Day 2)

### Strengths:

The study addresses a clear and timely theoretical question about how the training regime affects the formation of cognitive maps. A further strength is the well-controlled experimental design, allowing the authors to test their hypotheses in a systematic and informative way.

### Weaknesses:

(1) If I understood correctly, participants learned one map on the first day and the other on the second day, with the training regime (grouped vs. disjoint) counterbalanced across maps. This raises the possibility that experience with one training regime on day one could influence performance on the second day. For example, it would be interesting to examine whether participants who experienced the disjoint regime first showed any differences when learning the grouped regime on the following day. While it may be difficult to fully disentangle such transfer effects from the main training regime effects, it would be informative to test whether performance on the second day depends on the regime experienced on the first day (e.g., whether prior exposure to the disjoint regime predicts performance on the subsequent grouped training, but not vice versa).

(2) The author mentions a control experiment. Did the participants in the control experiment complete only the training phase or also the testing tasks used in the main experiment? If testing was included, it would be informative to report whether performance at test was comparable to that observed in the main experiment. Given that this condition appears to involve blocked transitions while moving across both rows and columns, I would expect performance to fall somewhere between the grouped and disjoint conditions.

(3) Participants' performance did not differ between conditions in the map reconstruction task, suggesting that participants in both the grouped and disjoint regimes were ultimately able to form a cognitive map. Was this task always administered last during the testing session? I wonder whether the explicit request of the reconstruction task could have influenced participants' awareness of the map structure.

(4) The manuscript describes the study as consisting of four experiments (two groups per map shape, differing in the blocked versus semi-blocked schedule). However, based on the design described in the Methods, this appears more accurately characterized as a single experiment with two between factors: map geometry (grid vs. torus) and blocking schedule (blocked vs. semi-blocked) manipulated between participants, and training regime (grouped vs. disjoint) manipulated within participants.

(5) It is not entirely clear to me from the Results section whether performance at test differed between the two map geometries (grid and torus), or whether the reported effects of training regime were consistent across them.

<https://doi.org/10.7554/eLife.111035.1.sa0>