nature human behaviour

Article

https://doi.org/10.1038/s41562-025-02217-2

Mental graphs structure the storage and retrieval of visuomotor associations

Received: 6 June 2024

Accepted: 16 April 2025

Published online: 2 June 2025

Check for updates

Juliana E. Trach[1] ✉ & Samuel D. McDougle[1,2]

Much of human memory takes the form of cognitive graphs that allow us to relate and generalize knowledge. The influence of structured memory in the motor system is less clear. Here we examine how structured memory representations influence action selection when responses are retrieved from newly learned, hierarchical visuomotor maps. Human participants ($N=182$) learned visuomotor mappings with (or without) an imposed latent structure that linked visual stimulus features (for example, colour or shape) to intuitive motor distinctions, such as hands and pairs of fingers. In participants who learned structured visuomotor mappings, transitional response times indicated that retrieving the correct response from memory invoked the 'traversal' of a structured mental graph. Forced-response experiments revealed similar computations within individual trials. Moreover, graph-like representations persisted even after multiple days of practice with the visuomotor mappings. Our results point to direct links between internal computations over structured memory representations and the preparation of movements.

Graph-like structures are ubiquitous in human cognition, from organizing spatial knowledge to representing hierarchical plans and abstract sequences[1–4]. Graph-like mental representations have also been implicated in the learning and planning of action sequences[5,6], suggesting that such representations may be fundamental to both memory and cognitive control. What is less clear is the relationship between structured memory and the real-time dynamics of selecting a single movement. Are structured memories strictly in the domain of cognition? Or might traces of internal memory structures also be echoed in the motor system?

Consider piano sight-reading, where clefs (for example, bass versus treble), note locations (for example, the third line on the staff) and accidentals (for example, sharps and flats) are combined to determine key presses (Fig. 1a). How does the sight-reading musician rapidly navigate their memory of symbol-to-finger mappings to produce fast, accurate finger movements? One possibility is that navigating this internal mapping and generating an action are strictly separable computational stages. That is, musicians may implement an algorithm that parses features of the stimulus and queries an internal representation of the relevant symbol-finger mapping. Then, only when the decision about the desired response is complete, they shuttle the result (for example, D flat/right index finger) to their motor system. Alternatively, we hypothesize that people may automatically prepare relevant motor commands while in the process of querying visuomotor memory. For example, determining the clef might potentiate the fingers of one entire hand, and then determining the exact note might initiate the movement of one finger on that hand. This type of coupling between decision-making and movement would be consistent with work in both humans and animal models suggesting that sensory evidence accumulation is echoed in the motor system[7–16]. However, in our proposal, this coupling is not dependent on the accumulation of perceptuomotor evidence (for example, dot motion tasks) but can be revealed even when an abstract memory representation mediates between perception and action. If supported, this idea would blur the lines between structured memory retrieval and movement selection processes.

Here, we aimed to test these competing hypotheses using a variant of an arbitrary visuomotor association learning task. Our primary goal was to understand how these more complex types of decisions—ones that require retrieving information from structured mental graphs—may interact with the motor system on short timescales. To that end, our task motivated participants to use stimulus features (colour, shape and pattern) to determine correct responses, similar to musicians considering notes on the page as they prepare to play a note. Critically, individual features could be associated with different 'levels' of an intuitive motor hierarchy (Fig. 1b). The hierarchically structured

[1]Department of Psychology, Yale University, New Haven, CT, USA. [2]Wu Tsai Institute, Yale University, New Haven, CT, USA. ✉e-mail: juliana.trach@yale.edu

https://doi.org/10.1038/s41562-025-02217-2

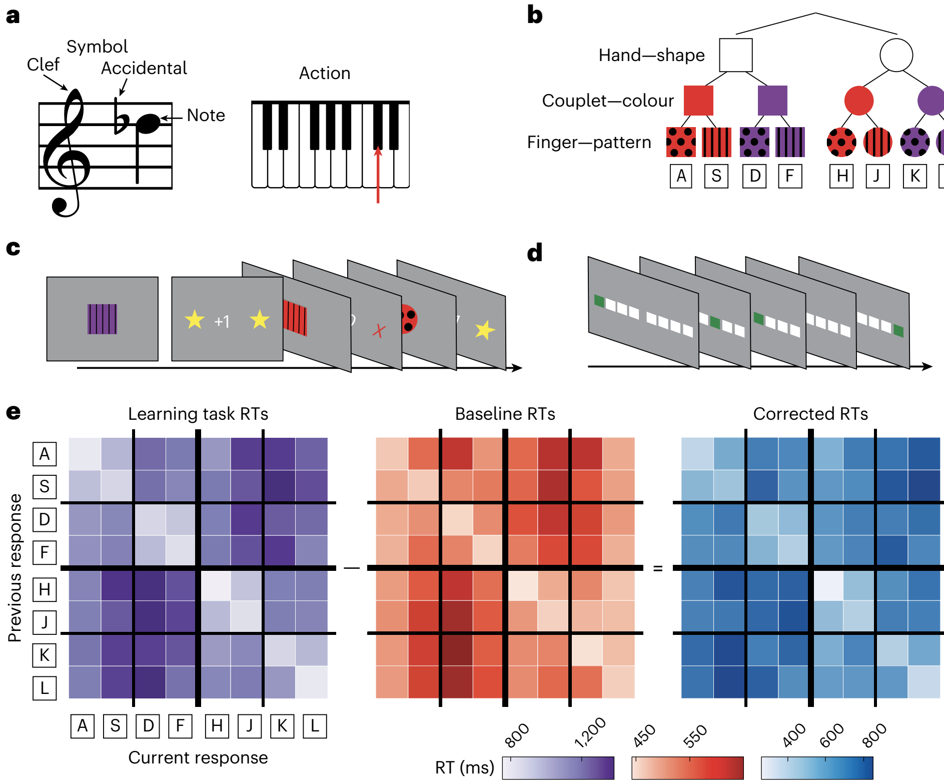

Fig. 1 | Task design and baseline correction. **a**, Example musical notation and the associated action/key. **b**, Example visuomotor mapping with the correct key press for each of the eight stimuli and illustration of one specific feature-to-level assignment. Visuomotor mappings were counterbalanced across participants. **c**, Task schematic for the learning task. **d**, Task schematic for the RT baseline task. **e**, Visualization of RT baseline correction, with average RTs for each pairwise transition between fingers for the learning and RT baseline tasks (Experiment 1) depicted as heat maps. Baseline RTs are subtracted from learning task RTs to yield corrected RTs.

mapping allowed us to make precise predictions about behaviour and ask whether the structure of an internal visuomotor mapping guides rapid movement preparation.

We established that a simple measure—participants' trial-by-trial response times—could reveal the latent structure of a learned visuomotor mapping (Experiment 1). We then accounted for various potential alternative explanations, such as intrinsic response time costs when switching between different fingers, in control experiments (Experiments 2 and 3). Next, we used a variant of a forced-response-time paradigm to characterize within-trial retrieval dynamics[17–19] (Experiments 4 and 5) by interrupting the retrieval of learned visuomotor associations at various time points during deliberation and measuring the resulting errors. We found evidence that people sequentially 'prune' the structured visuomotor mapping from top to bottom during the preparation of single finger movements. These data could be described by a simple computational model in which stimulus features were prioritized during action selection on the basis of the mapping structure to dynamically potentiate difference clusters of potential motor actions. Finally, we show that interactions between action selection and the structure of the learned mapping persist even after extensive practice (eight days). We speculate that retrieving actions from a structured visuomotor memory invokes a navigation-like computation over a cognitive graph or neural state space[2,20,21] and propose that this process can dynamically shape motor planning.

Results

Reaction times reflect the structure of learned mappings

Our goal was to understand how structured visuomotor mental representations, akin to note–key pairings in music sight-reading (Fig. 1a), dynamically interact with the motor system during action selection. To do this, participants ($N = 40$) engaged in a visuomotor learning task (Fig. 1b,c) and a reaction time (RT) baseline task (Fig. 1d). During the task, the participants used trial-by-trial feedback to learn an eight-to-eight deterministic visuomotor mapping (Fig. 1b,c). To embed structure into the mapping, stimuli varied along three features (colour, shape and pattern), and we assigned each dimension to a level of an intuitive motor hierarchy (hand > finger-couplet > finger)[22]. For example, if shape was assigned to the highest level of the hierarchy—hand—then the shape of the stimulus determined which hand contained the correct response (Fig. 1b). We confirmed that these mappings were intuitive to participants in a naive group of participants with a survey (Methods and Supplementary Fig. 1).

We used variation in RT between pairs of trials ('transitional RTs') as an index of whether people were sensitive to the structure of the mappings[23,24]. Transitional RT analyses involve classifying RTs on the basis of features of the current and previous trials or responses, rather than considering each trial independently. We isolated the impact of the mapping structure on transitional RTs and controlled for intrinsic finger-to-finger switch costs by subtracting the mean RT for each of the pairwise transitions between fingers measured during the RT baseline task from the RTs in the learning task for the same finger-to-finger transitions (Fig. 1e). The remaining variation in RT switch costs during the learning task therefore should not be driven by intrinsic finger transition biases but rather by the latent structure of the mapping. We conducted all primary analyses on these corrected RTs.

We compared participant transitional RTs to the predictions of three theoretical models designed to explain transitional RTs: (1) a Hierarchical graph model (our hypothesis), (2) a Feature-Based model and (3) a Flat model (Fig. 2). We entered the corrected RTs (for consecutively correct trials only; Methods) into three linear mixed-effects models that operationalized each theoretical model and assessed the model fit for each.

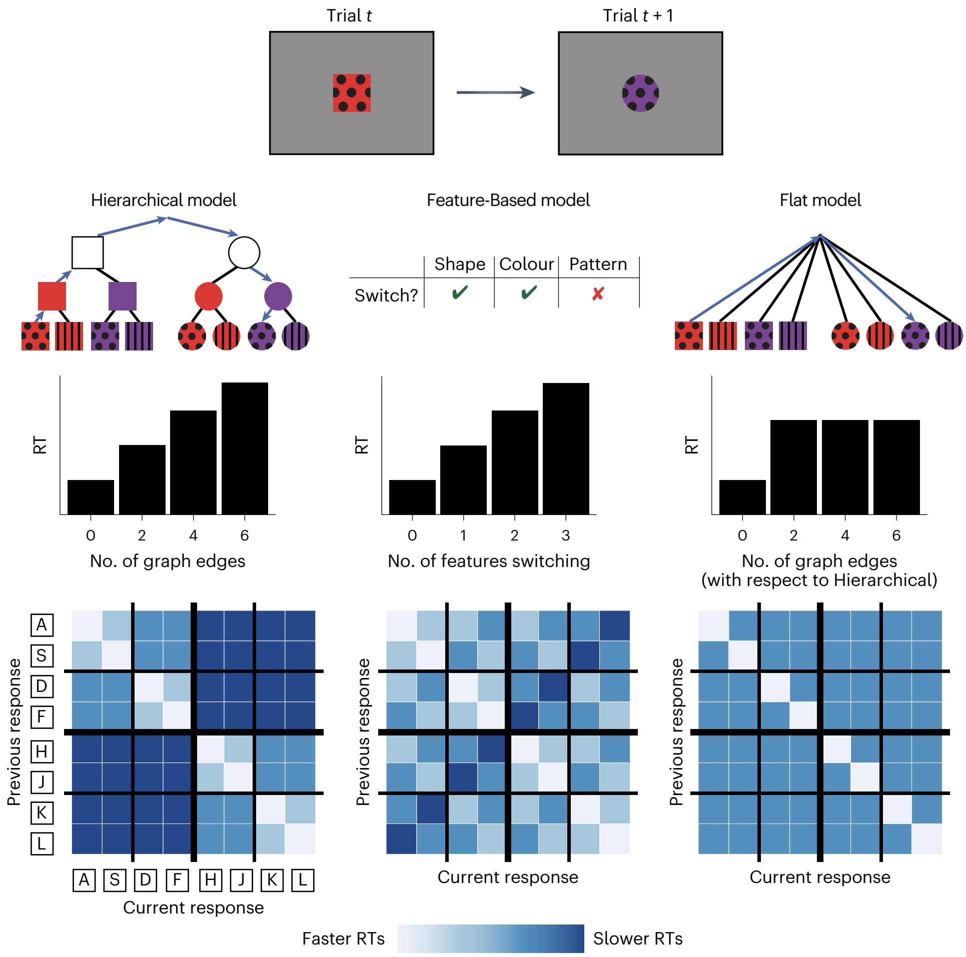

Fig. 2 | Theoretical models. Example trial and illustrations of three theoretical models of behaviour with predicted transitional RTs under each model (upper, summarized by path distance or number of features switching across trials; lower, 8 × 8 transitional structure).

First, the Hierarchical model used distance within the visuomotor mapping structure (that is, the number of graph edges between the current and previous stimuli in the hierarchical structure; [0,2,4,6]) to predict RTs. We use the term 'path distance' to refer to the number of graph edges between consecutive responses. In this model, longer path distances should be associated with slower transitional RTs. Second, the Feature-Based model predicted transitional RTs using the number of stimulus features that switched on a given transition ([0,1,2,3]). This model posits slower transitional RTs when more features of the stimulus change across trials but does not impose any structure or treat switching of specific features differently. Importantly, the design of the stimuli and underlying structure in the mapping ensured that the Hierarchical and Feature-Based models predict different patterns of RT behaviour (Fig. 2, bottom). Third, if participants were simply learning one-to-one associations between stimuli and actions but not representing any latent structure (a flat representation), switching from any response to any other response should incur comparable RT switch costs. The predictor for the Flat model was thus whether the stimulus repeated or switched across trials, with the assumption that repeating the same response facilitates RT. In addition to these theoretically grounded models, we included a Physical Distance model and a Nearest-Neighbour model (Methods) to operationalize a linearly modulated attentional effect where attention could be biased towards the previous response and thus responses far away from the previous response would take longer to prepare.

Our primary analyses provided convergent evidence that participants represented the latent hierarchical structure in the mapping: corrected transitional RTs scaled with the path distance through a graph of the visuomotor mapping (Fig. 3a,c; two-tailed paired t-test, Bonferroni-corrected $\alpha = 0.05/6 = 0.008$; 6-distance versus 4-distance: $t_{39} = 3.24$; $P = 0.002$; Cohen's $d = 0.30$; 95% confidence interval (CI), (17.8, 77.3); 6-distance versus 2-distance: $t_{39} = 8.50$; $P < 0.001$; Cohen's $d = 1.35$; 95% CI, (148.6, 241.3); 6-distance versus 0-distance: $t_{39} = 13.6$; $P < 0.001$; Cohen's $d = 2.73$; 95% CI, (315.2, 425.4); 4-distance versus 2-distance: $t_{39} = 7.12$; $P < 0.001$; Cohen's $d = 1.05$; 95% CI, (105.5, 189.3); 4-distance versus 0-distance: $t_{39} = 13.11$; $P < 0.001$; Cohen's $d = 2.44$; 95% CI, (272.9, 372.5); 2-distance versus 0-distance: $t_{39} = 13.92$; $P < 0.001$; Cohen's $d = 1.59$; 95% CI, (149.8, 200.8)). In other words, RTs monotonically increased with our hypothesized path distance metric (Fig. 3c). This result was further supported by the mixed-effects models: the Hierarchical model reliably produced the best fit to the behaviour, compared with the four competing models (Bayesian information criterion (BIC): Hierarchical, 330,109; Feature-Based, 330,171, Flat, 330,581, Physical Distance, 330,790; Nearest-Neighbour, 330,563; Fig. 3d). We replicated this result in a separate group of participants ($N = 27$) who were trained on a modified hierarchical structure (Supplementary Fig. 2). We also found that participant errors reflected the mapping structure (Supplementary Fig. 3). These findings suggest that participants learned and mentally represented the structure of the visuomotor mapping. Furthermore, the structure affected trial-by-trial RTs and errors.

Unstructured mappings yield different RT profiles

One potential concern about these results is that the RT baseline correction might not sufficiently control for intrinsic motor switch costs.

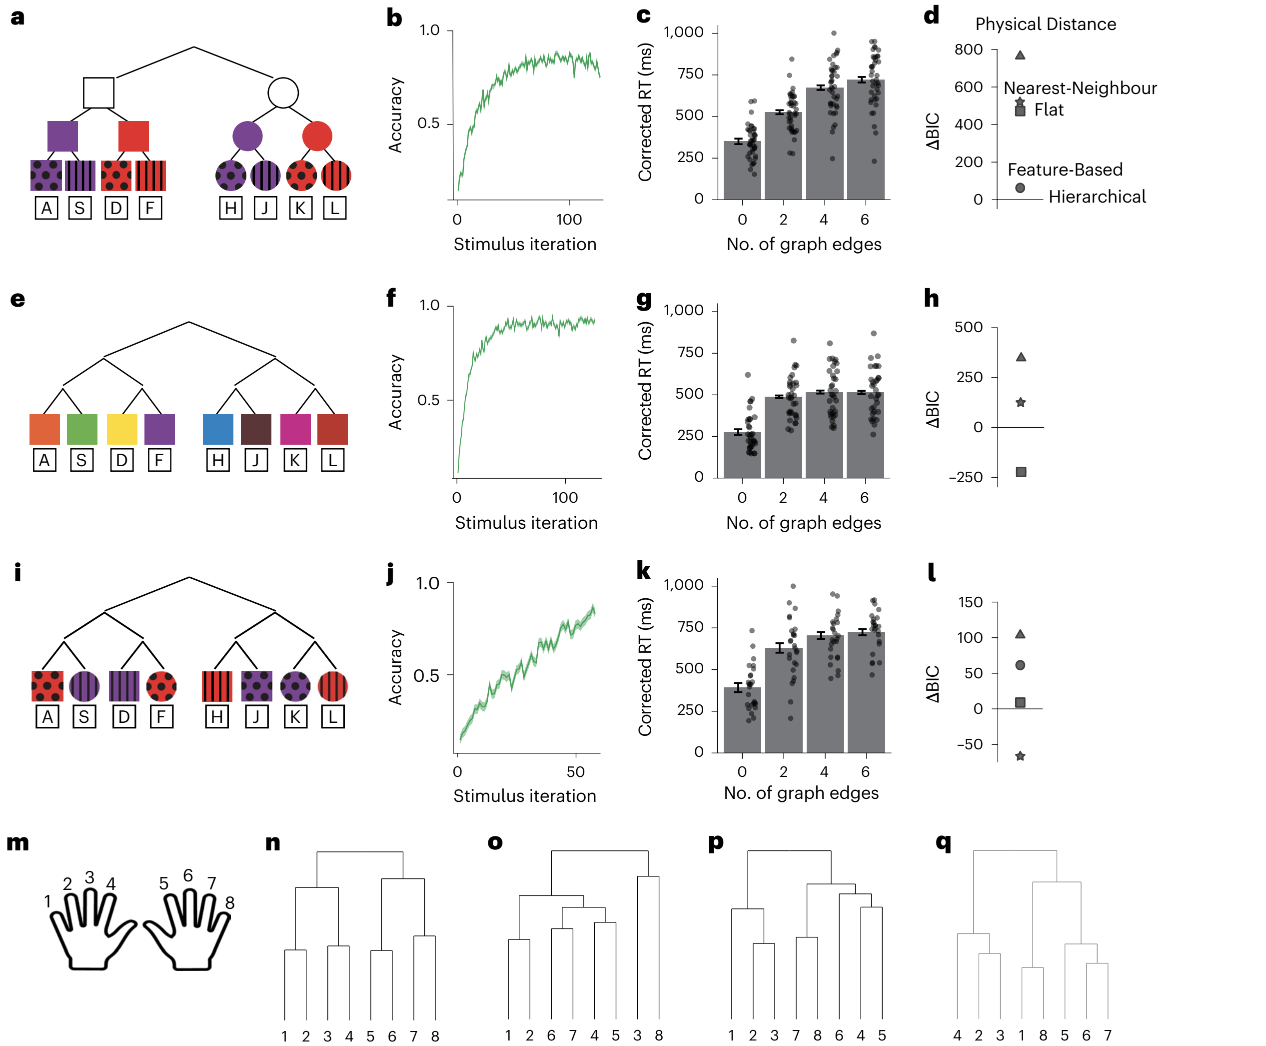

Fig. 3 | Results for Experiments 1–3. a, Example mapping in Experiment 1 ($N = 40$). **b**, Learning curve from Experiment 1. RTs and corrected RTs over the course of the task are plotted in Supplementary Fig. 10. **c**, Corrected RTs plotted by path distance (consecutively correct trials only) from Experiment 1. **d**, Linear mixed-effects modelling results from Experiment 1 plotted as the difference in BIC between the Hierarchical model (line at 0) and the alternatives (Feature-Based, Flat, Physical Distance and Nearest-Neighbour). Positive values indicate that the Hierarchical model was the best fit for participant behaviour. Models below the 0 line outperformed the Hierarchical model. **e**–**h**, Experiment 2 ($N = 33$)—same as Experiment 1. Note that H includes only three points for the Flat, Nearest-Neighbour and Physical Distance alternative models since stimuli vary only along one feature. **i**–**l**, Experiment 3 ($N = 26$)—same as Experiment 1. **m**, Legend for mapping the hierarchical clustering dendrogram results onto fingers. **n**–**p**, Hierarchical clustering results for Experiments 1–3. Hierarchical clustering reproduced the latent structure of the visuomotor mapping in Experiment 1. Unstructured mappings (that is, Experiments 2 and 3) yielded idiosyncratic dendrograms. **q**, Hierarchical clustering results for the RT baseline task. The data in **b**, **c**, **f**, **g**, **j** and **k** are presented as mean values ± 1 s.e.m.

We implemented two additional control studies that addressed this concern in different ways. In both, we removed the latent structure from the task and trained participants on 'unstructured' mappings (Fig. 3e,i). In Experiment 2, we used simpler stimuli than in Experiment 1 to ensure that there was no possibility of extracting any latent structure in the mapping. This experiment was designed to rule out the possibility that the results of Experiment 1 would arise from any eight-to-eight stimulus–action mapping regardless of structure. In Experiment 3, we constructed 'unstructured' mappings using the same three-feature stimuli from Experiment 1 (Supplementary Fig. 4). In these experiments, we quantified path distance using the same logic as the hierarchical mapping (from Experiment 1) and performed the same RT analyses.

In both cases, the data were not well fit by the Hierarchical model. In Experiment 2, we found that the Flat model was a better fit for participant behaviour (that is, lower BIC; Fig. 3h; BIC: Hierarchical, 325,873; Flat, 325,546; Physical Distance, 326,074; Nearest-Neighbour, 325,683). Transitional RTs were facilitated on repeat trials but were effectively equivalent across trials where the stimulus changed (Fig. 3g; two-tailed paired t-test, Bonferroni-corrected $\alpha = 0.05/6 = 0.008$; 0-distance versus 2-distance: $t_{32} = 11.98$; $P < 0.001$; Cohen's $d = 1.75$; 95% CI, (175.9, 248.0); 0-distance versus 4-distance: $t_{32} = 11.31$; $P < 0.001$; Cohen's $d = 1.83$; 95% CI, (197.2, 283.7); 0-distance versus 6-distance: $t_{32} = 10.68$; $P < 0.001$; Cohen's $d = 1.80$; 95% CI, (192.4, 283.2); 2-distance versus 4-distance: $t_{32} = 2.39$; $P = 0.023$; Cohen's $d = 0.21$; 95% CI, (4.3, 52.7); 2-distance versus 6-distance: $t_{32} = 2.23$; $P = 0.033$; Cohen's $d = 0.18$; 95% CI, (2.3, 49.4); 4-distance versus 6-distance: $t_{32} = -0.35$; $P = 0.725$; Cohen's $d = 0.02$; 95% CI, (−17.8, 12.5)). For Experiment 3, the Nearest-Neighbour model was the best-fitting model (Fig. 3l; BIC: Hierarchical, 58,706; Flat, 58,716; Feature-Based, 58,768; Physical Distance, 58,811; Nearest-Neighbour, 58,642), and RTs did not linearly increase with path distance (Fig. 3k; Bonferroni-corrected $\alpha = 0.05/6 = 0.008$; 0-distance versus 2-distance: $t_{25} = 5.5$; $P < 0.001$; Cohen's $d = 1.48$; 95% CI, (148.3, 325.9); 0-distance versus 4-distance: $t_{25} = 9.04$; $P < 0.001$;

Cohen's $d = 2.17$; 95% CI, (241.0, 383.2); 0-distance versus 6-distance: $t_{25} = 10.58$; $P < 0.001$; Cohen's $d = 2.56$; 95% CI, (267.3, 396.5); 2-distance versus 4-distance: $t_{25} = 2.19$; $P = 0.038$; Cohen's $d = 0.45$; 95% CI, (4.6, 145.4); 2-distance versus 6-distance: $t_{25} = 2.78$; $P = 0.010$; Cohen's $d = 0.6$; 95% CI, (24.7, 165.0); 4-distance versus 6-distance: $t_{25} = 0.76$; $P = 0.457$; Cohen's $d = 0.14$; 95% CI, (−34.1, 73.7)). Taken together, these results refute the possibility that our previous results were an artefact of linking our mapping to physical effectors and suggest that latent mapping structure impacts the pattern of RTs.

Model-free clustering of RTs reproduces latent structure

We performed a model-free clustering analysis to reconstruct the mappings from transitional RT data[25,26] (R packages cluster and factoextra). We predicted that this analysis would reproduce the hierarchical structure of the mapping for Experiment 1 but yield idiosyncratic structures when there was no consistent structure. We calculated average RTs for each pairwise transition between responses to obtain a transitional RT profile for each target stimulus. We then calculated the Euclidean distances between the RT profiles and performed clustering on the transitional RT profiles. Finally, we visualized the inferred structure of the mapping with dendrograms (Fig. 3n–p).

The clustering algorithm faithfully reproduced the latent structure of the mapping from Experiment 1 (Fig. 3n). In contrast, experiments without latent structure yielded idiosyncratic dendrograms (Fig. 3o–q). This data-driven analysis further demonstrates that individuals learned and used the structure built into the task when it was available and reiterates our finding that the format of memory representations can be inferred from a simple behavioural measure—transitional RTs.

Experiments 4 and 5: within-trial dynamics of action selection

The results from Experiments 1–3 show that the structure of a visuomotor mapping shapes trial-by-trial action selection. We posited that this result could arise from individuals mentally traversing an internal representation of the visuomotor mapping—a cognitive graph—to retrieve correct responses, similar to the traversal of structured mental representations of non-motor content found in other domains[27–29]. We hypothesized that this latent traversal process would automatically communicate with the motor system, potentiating relevant sets of actions in real time as people query nodes in the cognitive graph. In the context of our task, participants might sequentially potentiate (or prune) responses by considering those that share the top-, then mid- and finally low-level features of the stimulus to arrive at the correct response. This sequential dynamic is structured, as the latent structure of the mapping directly shapes the action selection process. In contrast, an unstructured dynamic would describe a process where the latent structure of the mapping is not evident in the action selection process.

We tested this hypothesis using a paradigm designed to elicit responses at different points during action selection on each trial. We trained two groups of participants on either the structured visuomotor mapping used in Experiment 1 (Experiment 4) or an unstructured visuomotor mapping, as in Experiment 2 (Experiment 5). We then compared the probabilities of different types of errors as a function of how long participants had to prepare their responses on a given trial.

Learning task results: replication of Experiments 1 and 2

The learning phase results replicated the results of Experiments 1 and 2, respectively (Supplementary Fig. 5). That is, the Hierarchical model was the best fit for participants trained on the structured mapping (BIC: Hierarchical, 117,458; Feature-Based, 117,656; Flat, 117,930; Physical Distance, 117,981; Nearest-Neighbour, 118,395), and the Flat model was the best fit for participants trained on the unstructured mapping (BIC: Hierarchical, 70,254; Flat, 70,134; Physical Distance, 70,350; Nearest-Neighbour, 70,381). The model-free hierarchical clustering algorithm again reliably reproduced the latent structure in the task for participants trained on a structured mapping (Experiment 4) but not an unstructured mapping (Experiment 5; Supplementary Fig. 5).

Action preparation dynamics accord with mapping structure

After learning, the participants performed a forced-response task (Fig. 4a). During this task, the participants heard four beeps on every trial and were instructed to synchronize their response with the fourth beep, regardless of whether they felt prepared to respond[18,19]. We varied the stimulus onset during each trial to manipulate the amount of preparation time (PT; that is, the time between stimulus onset and the fourth beep) that participants had (100 ms–1.2 s). On some trials, participants had sufficient time to plan and execute their responses, and on other trials they were forced to guess. Our analyses focus on the types of errors participants made as a function of PT. We examined the within-trial dynamics of action selection, rather than considering the influence of previous trials, as the forced-response paradigm necessitates a large number of error trials, which are difficult to interpret in transitional RT analyses (but see Supplementary Fig. 6 for analysis of previous trial effects).

Our main question was whether participants sequentially visit nodes on the cognitive graph of the visuomotor mapping and simultaneously prepare relevant sets of actions (structured action preparation) or not (unstructured action preparation). Consider again the example of the pianist: the structured preparation hypothesis describes a model of behaviour where they parse the musical notation and, at the same time, potentiate movements associated with different levels (clef, note location, accidental and so on) of an implied cognitive graph. The unstructured hypothesis, in contrast, describes a dynamic where they first determine the appropriate action (cognitive stage) before potentiating that response in the motor system (motor stage).

Crucially, these two hypotheses predict different patterns of errors. For structured action preparation, we expected top-level errors (that is, hand errors) to be less frequent than mid- or low-level errors (couplet or finger-level errors), since resolving the top level of the response should occur before resolving the lower two levels. Similarly, mid-level (couplet) errors should be less frequent than low-level errors, and low-level errors should be the most frequent, since it takes the most amount of time to finally resolve the subordinate level. There should thus be an orderly procession of error-type probabilities from top to bottom. In contrast, if participants do not plan any finger movements prior to a terminal decision, we would not expect any orderly progression of error types (unstructured action preparation). We note that while the structure of the mapping is hierarchical, the structured pattern of behaviour does not necessitate a strictly hierarchical dynamic where superordinate features of the mapping must be fully resolved before lower-level features are processed; a structured dynamic could also arise from stimulus features being prioritized from top to bottom but with partially parallel processing of features. We focus here instead on evidence for structured versus unstructured action selection, rather than subtle differences between variants of a structured dynamic.

We classified responses as top-, mid- or low-level errors on the basis of shared features between the target stimulus and the stimulus associated with the response participants made (Fig. 4b). For example, if the target response was key D, then responding with key F would be considered a low-level error (that is, the hand and couplet are correct, but the wrong finger was chosen), key A or key S would be a mid-level error (that is, the correct hand was chosen, but not the correct couplet) and responding with any finger on the right hand would be a top-level error. We then normalized that probability by the number of responses that were classified as each type of error to account for the fact that there were more ways to commit top-level errors (four responses) than mid-level (two responses) or low-level (one response) errors.

We found clear evidence of structured action preparation for participants trained on the structured mapping: the probability of errors at different levels of the task stacked in an orderly fashion, such

https://doi.org/10.1038/s41562-025-02217-2

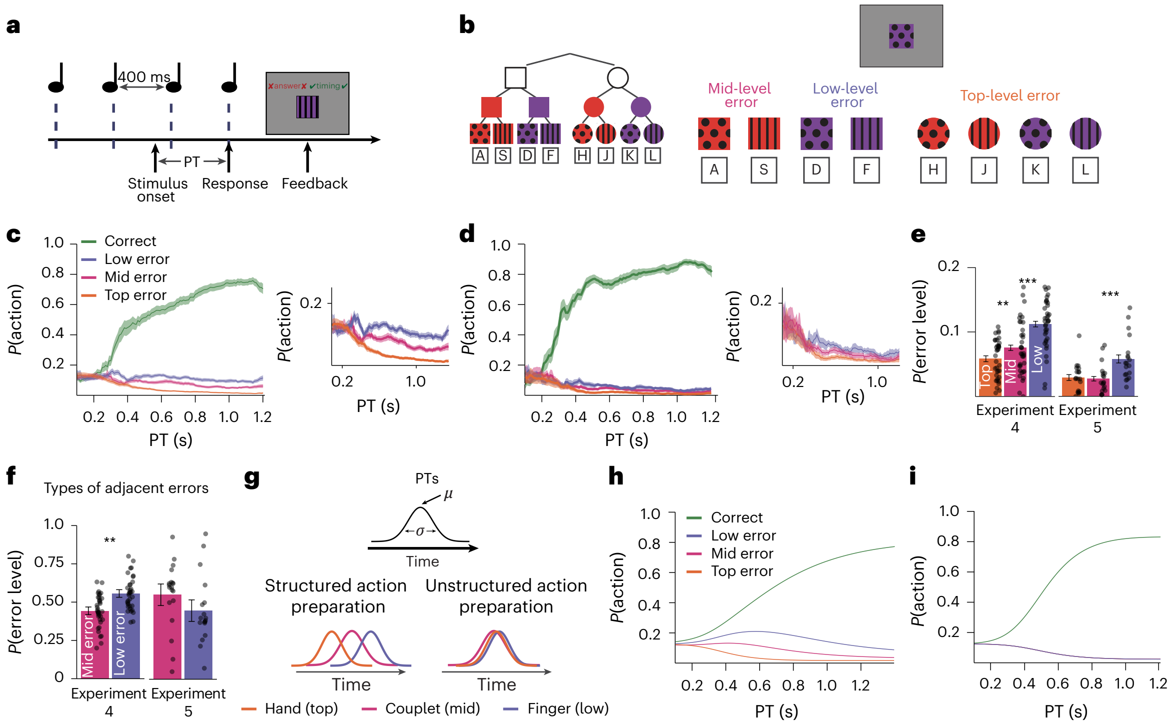

Fig. 4 | Forced-response design and results. a, Diagram of a forced-response trial. The vertical arrows indicate the stimulus onset and response cue. PT, preparation time. **b**, Error coding for an example trial. Errors were coded on the basis of the highest level feature that mismatched between the target stimulus and the stimulus associated with the response the participant made. **c**, Forced-response results for Experiment 4 ($N = 36$). The inset is the same, without the line for correct trials for a clearer view of error probabilities. **d**, Forced-response results for Experiment 5 ($N = 19$). The inset is the same, without the line for correct trials. **e**, Left, probability of errors that share two visual features with the target in Experiment 4 ($N = 36$). Two-tailed, Bonferroni-corrected t-test, top versus mid, $P = 0.002$; mid versus low, $P < 0.001$. Right, same plot for Experiment 5 ($N = 19$) with errors labelled with the same structure as Experiment 4. Two-tailed, Bonferroni-corrected t-test, top versus mid, $P = 0.715$; mid versus low, $P < 0.001$. **f**, Probability of a within-couplet (purple; that is, low-level) error versus an across-couplet (pink; that is, mid-level) error when there are two possible adjacent finger errors. The left panel shows results for Experiment 4 ($N = 36$; two-tailed t-test, mid versus low, $P = 0.003$), and the right panel shows results for Experiment 5 ($N = 19$; two-tailed t-test, mid versus low, $P = 0.310$). **g**, Model logic and predictions of structured and unstructured action preparation models. **h**, Model simulation with averaged fitted parameters for the best-fitting hierarchical three-σ model from Experiment 4. **i**, Model simulation with averaged fitted parameters for the best-fitting flat model from Experiment 5. Note that lines for errors at different mapping levels are on top of each other. The data in **c**–**f** are presented as mean values ± 1 s.e.m. $^{**}P < .01$; $^{***}P < 0.001$.

that top-level errors were the least frequent and resolved quickly, while low-level errors were the most frequent and resolved slowly (Fig. 4c; see also Supplementary Fig. 7a–c). We did not see this pattern for participants trained on the unstructured mapping (Experiment 5; Fig. 4d and Supplementary Fig. 7d–f). An analysis of variance (ANOVA) on the normalized probability of errors at each level across all PTs (Experiment (4 versus 5) × Error Level (top, mid, low)) revealed significant main effects of Experiment ($F_{1,52} = 23.7$, $P < 0.001$, $\eta^2 = 0.31$) and Error Level ($F_{1.65,87.51} = 55.4$, $P < 0.001$, $\eta^2 = 0.51$). Crucially, we found a significant interaction between Experiment and Error Level, reflecting the different selection dynamics between the two experiments (Experiment × Error Level: $F_{1.65,87.51} = 9.34$, $P < 0.001$, $\eta^2 = 0.15$).

Post hoc t-tests showed that the probability of errors stacked in an orderly manner, such that low-level errors were the most common and top-level errors were the least common for participants trained on the structured mapping (Supplementary Fig. 8; low-level versus mid-level errors: $t_{35} = 6.38$; $P < 0.001$; Cohen's $d = 0.98$; 95% CI, (0.02, 0.05); mid-level versus top-level errors: $t_{35} = 7.95$; $P < 0.001$; Cohen's $d = 1.1$; 95% CI, (0.03, 0.04)). In contrast, in Experiment 5, low- and mid-level errors did not differ in their frequency ($t_{18} = 1.51$; $P = 0.149$; Cohen's $d = 0.48$; 95% CI, (−0.006, 0.04)), but top-level errors were less frequent than mid- or low-level errors (top versus mid: $t_{18} = 2.99$; $P = 0.008$; Cohen's $d = 0.58$; 95% CI, (0.02, 0.004); top versus low: $t_{18} = 3.36$; $P = 0.003$; Cohen's $d = 1.0$; 95% CI, (0.05, 0.01)). Taken together, participants in Experiment 4 arrived at the correct action by pruning the visuomotor mapping in real time, while participants in Experiment 5 displayed a different pattern of errors.

One alternative explanation for the effects observed in Experiment 4 is that participants may be reacting to the visual similarity between the target stimulus and the stimulus associated with the low-level error response. The stimulus associated with the low-level error response shares two of three features with the target stimulus by definition; thus, it is possible that the increased probability of this response was driven by visual similarity. Importantly, there are three stimuli that share two features with the target stimulus in the structured visuomotor mapping: one of these three stimuli is the low-level error, one a mid-level error and one a top-level error, even though they all share two features with the target. We thus compared the probability of making each of these types of errors (Fig. 4e, left). Participants were still most likely to make the low-level versus mid- or top-level errors, even when the stimuli were matched for visual similarity with the target (Bonferroni-corrected $\alpha = 0.05/2 = 0.025$; low- versus mid-level error: $t_{35} = 5.48$; $P < 0.001$; Cohen's $d = 0.92$; 95% CI, (0.02, 0.05); low- versus top-level error: $t_{35} = 7.24$; $P < 0.001$; Cohen's $d = 1.57$; 95% CI, (0.04, 0.07); mid- versus top-level error: $t_{35} = 2.4$; $P = 0.02$; Cohen's $d = 0.46$; 95% CI, (0.003, 0.03)). Furthermore, the pattern of errors was different between Experiments 4 and 5 (mixed-factor ANOVA: Experiment × Error Type: $F_{1.22,64.51} = 14.8$, $P < 0.001$, $\eta^2 = 0.22$). Specifically, while participants in Experiment 4 were more likely to make mid-level errors than top-level errors, this was not the case in Experiment 5 ($t_{18} = −0.39$; $P = 0.699$; Cohen's $d = −0.08$; 95% CI, (−0.01, 0.008)). This strongly suggests that our results were driven not by the visual similarity of the stimuli but rather by the latent structure of the learned visuomotor mapping.

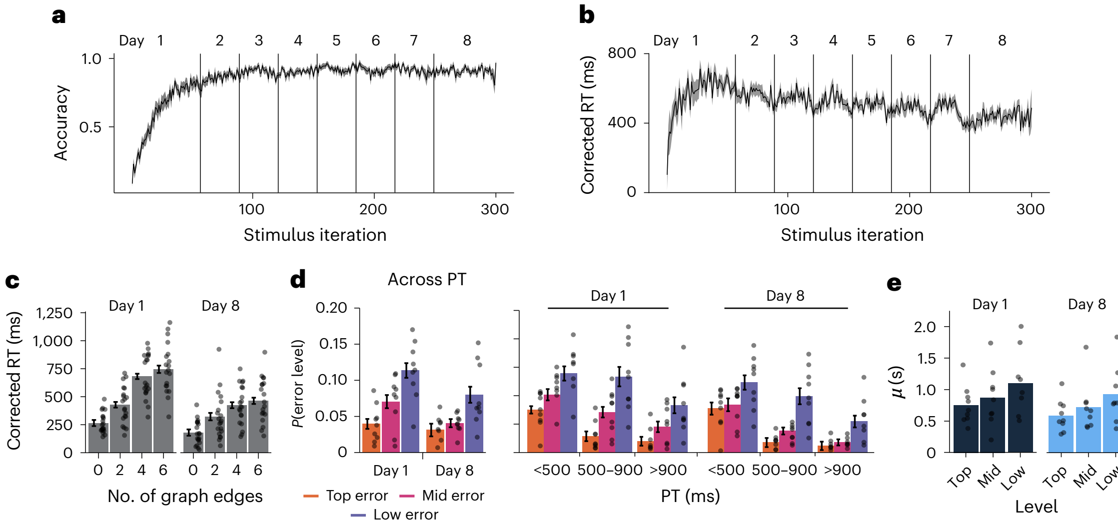

Fig. 5 | Experiment 6 results. a, Learning curve over stimulus iteration. The vertical lines delineate between days of practice. **b**, RT over stimulus iteration. The vertical lines delineate between days of practice. **c**, Corrected RTs plotted by path distance (consecutively correct trials only) for Day 1 and Day 8 ($N = 20$). **d**, Corrected probability of errors at each level of the mapping. The first panel is combined across PTs. The second panel shows the probability of errors at each level of the task in 400-ms bins of PT. $N = 9$. **e**, Values of μ at each level of the mapping on Day 1 and Day 8 ($N = 9$). The data are presented as mean values ± 1 s.e.m.

Another possibility is that spatial proximity to the target response could drive the increased likelihood of low-level errors (that is, action slips with neighbouring fingers). To address this possibility, we compared the probability of making a low-level error to the probability of making a mid-level error that was also adjacent to the target response (on trials where there were responses on either side of the target response). For participants trained on the unstructured mapping (Experiment 5), the two types of adjacent errors were equally probable (Fig. 4f, right; $t_{18} = 1.04$; $P = 0.311$; Cohen's $d = 0.48$; 95% CI, (-0.11, 0.31)), supporting the presence of these spatially driven generic motor errors. In contrast, participants trained on the structured mapping (Experiment 4) were significantly more likely to commit low-level adjacent errors (that is, consistent with the structure of the task) than the alternative adjacent errors (Fig. 4f, left; $t_{35} = 3.20$; $P = 0.003$; Cohen's $d = 1.07$; 95% CI, (−0.2, −0.04)). These control analyses (Fig. 4e,f) provide convincing evidence that the pattern of errors observed in Experiment 4 was a direct result of the structure of the learned mapping.

Computational modelling of distinct action selection dynamics

We used computational models to further clarify the visuomotor processing dynamics and compare the results across Experiments 4 and 5. The basic design of the model posits that preparing an action takes some mean amount of time μ with variance σ and is normally distributed[18] (Methods). These distributions can be transformed into curves that describe the probability of making a specific response as a function of PT by taking their cumulative density. We compared four variants of this model to characterize response dynamics in our structured and unstructured tasks. To adapt this model to our hierarchically structured task, we allowed for three μ parameters to vary freely, one for each level of the mapping (Fig. 4g). We fit one version of this Hierarchical model that included a single σ parameter that was used at each level of the mapping and another version that included separate σ parameters for each level of the mapping. We also fit a Feature-Based model and a Flat model for comparison (Methods). We predicted that one of the Hierarchical models would best capture participant behaviour when they learned a structured mapping (Experiment 4) and that the Flat model would best capture behaviour for participants trained on an unstructured mapping (Experiment 5).

As predicted, the Hierarchical models were the best fit for participant behaviour for participants trained on the structured mapping (Hierarchical-three-σ: summed BIC = 76,517; Hierarchical-one-σ: summed BIC = 76,536; Feature-Based: summed BIC = 77,767; Flat: summed BIC = 78,054). Fitted μ parameters were consistent with a sequential pruning process—μ_{top} (0.615 s) was smaller than μ_{mid} (1.01 s), and μ_{mid} was smaller than μ_{low} (1.26 s; Wilcoxon signed-rank test: top versus mid: $z = 4.32$; $P < 0.001$; effect size, $r = 0.72$; mid versus bottom: $z = 4.07$; $P < 0.001$; effect size, $r = 0.72$). In contrast, participants trained on the unstructured mapping were better fit by the Flat model (Hierarchical-three-σ: summed BIC = 34,863; Hierarchical-one-σ: summed BIC = 34,732; Flat: summed BIC = 30,917). Example model simulations are pictured in Fig. 4h,i. These results provide further evidence of structured action preparation during retrieval of visuomotor associations.

Structured dynamics persist after extensive practice

One open question is whether the traversal of the structured representation is a transient phenomenon that appears only when people are first forming the visuomotor memory, or whether it persists even with extensive practice. To address this question, we had participants ($N = 20$) practise a hierarchically structured mapping (as in Experiment 1) over eight consecutive days (Experiment 6; Methods and Fig. 5a,b). We compared transitional RTs between Day 1 and Day 8 of the study, as in the previous experiments. Comparison of the RT profiles on Day 1 versus Day 8 revealed that RTs again varied as a function of path distance (Fig. 5c; ANOVA: Day × Path Distance: main effect of Path Distance: $F_{3,57} = 64.0$, $P < 0.001$, $\eta^2 = 0.77$), and also that participants were significantly faster during the last session (main effect of Day: $F_{1,19} = 37.54$, $P < 0.001$, $\eta^2 = 0.67$). Additionally, a significant interaction effect suggested that the pattern of RTs was subtly different between the first and last days of practice ($F_{1,19} = 19.73$, $P < 0.001$, $\eta^2 = 0.51$).

To ascertain whether the structured cognitive representation had compressed to a flat representation, we compared the RT profiles on Day 8 to those from the participants in Experiment 5 (where participants were trained on an unstructured mapping). Average RT and accuracy were comparable between the end of the last session of Experiment 6 and the end of free-response phase of Experiment 5 (RT: $t_{36.7} = 0.82$; $P = 0.416$; Cohen's $d = 0.26$; 95% CI, (−61.5, 145.7); accuracy: $t_{30.6} = 0.20$;

$P = 0.839$; Cohen's $d = 0.07$; 95% CI, (−0.04, 0.05)), so this analysis additionally addresses whether the results of Experiment 5 were driven by overall better performance with the simple stimuli rather than the different structure of the mapping. Interestingly, we found that the pattern of RTs on the last day of this study were significantly different from the pattern of RTs in Experiment 5 (mixed-factor ANOVA: Experiment × Path Distance: main effect of Path Distance: $F_{2.0,74.6} = 70.48$, $P < 0.001$, $\eta^2 = 0.66$; Experiment × Path Distance: $F_{2.0,74.6} = 4.53$, $P = 0.014$, $\eta^2 = 0.11$), although there were no overall differences in RT between these samples (main effect of Experiment: $F_{1,37} = 0.83$, $P = 0.368$, $\eta^2 = 0.02$). Furthermore, from the linear mixed-effects modelling, we found that across all eight days, the Hierarchical model was the best fit for participant behaviour (see Supplementary Table 1 for BIC values). Together, these results indicate that participants continued to leverage the latent structure in the mapping to retrieve their responses even after extensive practice.

A subset of these participants ($N = 9$) also completed the forced-response task on the first and last days of the experiment, allowing us to address whether the structured dynamics we revealed in Experiment 4 persist after extensive practice. First, we examined changes in the probability of errors at different levels of the mapping between the first and last days of the experiment (Fig. 5d). We conducted an ANOVA with one factor for Day (1 versus 8) and another for Error Level (top, mid, low) on the normalized probabilities of errors at each level of the task (excluding correct responses). There was a significant effect of Day ($F_{1,8} = 9.75$, $P = 0.014$, $\eta^2 = 0.55$), as participants had overall higher accuracy during the session on the eighth day. In addition, there was a significant effect of Error Level ($F_{2,16} = 19.34$, $P < 0.001$, $\eta^2 = 0.71$), though the interaction was not significant (Day × Error Level: $F_{2,16} = 1.88$, $P = 0.185$, $\eta^2 = 0.19$), indicating that the pattern of errors was similar across the two sessions.

We followed this ANOVA with post hoc t-tests to compare the frequency of errors across levels of the mapping. On Day 1, we found that error probabilities stacked such that low-level errors were more frequent than mid-level errors and mid-level errors were more frequent than top-level errors (Bonferroni-corrected $\alpha = 0.05/2 = 0.025$; mid versus low: $t_8 = 3.55$; $P = 0.008$; Cohen's $d = 1.1$; 95% CI, (0.02,0.07); top versus mid: $t_8 = 3.00$; $P = 0.017$; Cohen's $d = 0.95$; 95% CI, (0.007, 0.05)). On Day 8, low-level errors were again more frequent than mid-level errors (mid versus low: $t_8 = 2.85$; $P = 0.021$; Cohen's $d = 1.1$; 95% CI, (0.008,0.07)); however, the comparison between top- and mid-level errors did not survive correction ($t_8 = 2.36$; $P = 0.046$; Cohen's $d = 0.51$; 95% CI, (0.0002, 0.02)). This pattern of results held when we incorporated PT as well (Supplementary Fig. 9).

Finally, we fit the PT models discussed previously (Fig. 4g; Hierarchical-one-σ, Hierarchical-three-σ, Feature-Based and Flat) to the data from Experiment 6. The Hierarchical models were the best fit for participant data during the first and last sessions (Day 1: Hierarchical-one-σ: summed BIC = 17,680; Hierarchical-three-σ: summed BIC = 17,674; Feature-Based: summed BIC = 18,012; Flat: summed BIC = 18,105; Day 8: Hierarchical-one-σ: summed BIC = 14,260; Hierarchical-three-σ: summed BIC = 14,268; Feature-Based: summed BIC = 14,446; Flat: summed BIC = 14,255), although the additional σ parameters in this case did not produce a substantially better model fit. Furthermore, reductions in the value of the fitted μ parameters reflect the idea that participants are getting faster with practice but not necessarily compressing the representation (Fig. 5e)—the fitted parameters suggest that retrieval is faster on Day 8 but still affected by the latent structure. Taken together, these results demonstrate that the structure of the cognitive representation continued to be echoed in response retrieval dynamics after extensive practice.

Discussion

Complex memory structures, such as cognitive maps and graphs, are typically studied in the domain of higher-level cognition. But what happens at the interface of structured memory and movement preparation? In these studies, we examined interactions between retrieving visuomotor memories and rapid action selection in the context of graph-like memory representations. We hypothesized that the structure of the memory representation would constrain action selection dynamics and prompt a navigation-like computation over this latent structure, as is evident in other domains[5,6,28,30]. Our results support this hypothesis: when participants were trained on hierarchically structured visuomotor mappings, participant transitional response times closely tracked the structure of the learned visuomotor mapping (Experiment 1). This finding provided behavioural evidence that individuals learned and used that structure to retrieve responses on a trial-by-trial basis. Furthermore, this result held after controlling for intrinsic switch costs between fingers and was abolished when structure was removed (Experiments 2 and 3).

We expanded on this finding in Experiments 4 and 5 to characterize how the learned mappings constrained action selection within the scope of individual trials: participants' errors systematically varied as a function of movement PT in a manner consistent with a structured model of action selection; that is, we found evidence that participants resolved hierarchical levels of the cognitive graph from top to bottom and concurrently potentiated relevant sets of actions. This effect was not seen when the learned mapping did not contain a structured relationship between perceptual features and actions (Experiment 5), helping to rule out explanations based on more basic processes of action preparation[31]. Finally, in Experiment 6 we tested the persistence of this structured retrieval process with increased practice and found that structured action selection dynamics were evident even after eight days of practice. Taken together, our findings point to a dynamic, rapid interaction between a cognitive process—an internal navigation-like computation over a structured memory—and the preparation of movements.

Our results thus suggest that graph-like representational formats can be used in the context of action selection and interface directly with motor preparation processes. It is possible that distances in a low-dimensional neural 'state space' may correlate with the path distances we posit in our study, where navigation could reflect internal control processes involved with traversing or reconfiguring these state spaces[21]. This could occur via sequentially prioritizing different visual features (and simultaneously activating different action sets) at different times. Indeed, similar processes have been proffered to explain classic task-switching effects[32]. Moreover, there is evidence that neural activity and participant RTs can scale with path length through a putative cognitive graph[28,33]. At the neural level, such structured representations are typically believed to exist in traditionally 'non-motor' regions such as the hippocampus and orbital frontal cortex[2,34–38], raising the possibility that these regions might also be involved in storing and accessing structured perceptuomotor mappings.

Our work is also related to a large literature examining information flow between sensory evidence accumulation and the motor system during perceptual decision-making. This work has suggested that pre-movement activity in the motor system can reflect ongoing evidence accumulation processes in perception[10,39–44]. A large body of research using mouse- and eye-tracking methods during simple decision-making tasks also suggests that movements act as continuous read-outs of evolving decision processes, rather than being the output of a terminal decision process[8,11–14,16,45–49]. While this perspective generally aligns with our results, we note that our behavioural approach limits our ability to take a strong stance on the putative 'continuity' of these cognitive–motor interactions. The fully continuous model proposed in previous work represents one extreme version of cognitive–motor interactions during decision-making. The other extreme in this case is a discrete model where decision-making processes fully precede any preparatory activity in the motor system. Indeed, there is evidence that the truth lies somewhere in the middle of these two models—the degree of continuity in information flow depends at least to some degree on a variety of task features and the granularity of the decision[50]. Our results could be interpreted as evidence for a softer version of a discrete model, as long as such a model allows for multiple response selection stages within the scope of individual decisions (that is, at each level of the mapping).

In terms of within-decision dynamics (Experiments 4 and 5), some results in the cognitive control domain contrast with the structured processing dynamic that we identified here. Specifically, research with behavioural and neural recording methods has not found evidence of sequential pruning of potential responses during hierarchically structured cognitive control tasks[51,52]. Instead, people may process hierarchical levels of a structured mapping simultaneously in these cases. Parallelization in these tasks is thought to be possible due to a hypothesized hierarchical gradient of representational abstraction in the prefrontal cortex[53,54], where different areas of the prefrontal cortex can process different task rules/levels in parallel. A useful future direction could be to link previous work on task rules and cognitive control to our current study, where the structured representation was not defined by any strict hierarchical or contextual cues but rather was directly linked to motor effectors.

Our focus on newly learned visuomotor mappings naturally raises questions about how action selection dynamics might evolve with experience and, ultimately, expertise. If participants practised the mapping for longer (that is, becoming 'experts' at the mapping), would they show 'flattening', transforming the structured representation into a direct 8–8 stimulus–response mapping? Experiment 6 argued against this. Instead, action preparation was still closely linked to the latent structure of the mapping even after thousands of trials. Revisiting our pianist learning to parse musical notation while planning finger movements, early in learning the pianist is probably explicitly parsing individual symbols in order (clef → note → accidentals), dynamically and automatically potentiating different actions as they determine the correct key to press. What happens after years of training? Instead of overtrained stimulus–response associations becoming crystallized into flattened 'instances'[55,56], it may be that experts who are overtrained on our task, or even tasks like music sight-reading, still use a sequential parsing algorithm even after extensive practice but simply speed the algorithm up.

Our study has several limitations. First, our modelling is somewhat constrained; additional computations, such as an evidence accumulation threshold for proceeding through levels of the visuomotor mapping, could be added. Relatedly, our model does not delineate between the strictest possible version of a hierarchical processing model, where moving to the next level in the mental mapping can happen only after resolving the previous level, and looser variants of hierarchical dynamics, where different levels may be processed simultaneously but perhaps with some being prioritized over others. In any case, it is likely that there are both parallel and sequential processing dynamics at play during our task—an idea that has been debated for decades[57–59]. Finally, we focused here on navigation 'down' from a superordinate control node to the appropriate response in Experiments 4 and 5 and did not test whether there is evidence for 'climbing back up' the putative tree. Our design is not optimized for these analyses, although our supplemental analyses provide some initial evidence for this effect (Supplementary Fig. 6).

Taken together, our results suggest that cognitive memory structures can directly shape the dynamics of action selection. This work goes beyond previous findings in lower-level perceptual decision-making[10,39–44] by linking the potentiation of actions to higher-level structured memory representations. Our study thus makes new connections between research on cognitive maps and graphs in organizing behaviour and knowledge[2,60,61] and the study of sensorimotor learning and control, perhaps offering a new avenue for understanding the format of mental representations in complex, naturalistic visuomotor skills. Overall, our work raises questions about nominal distinctions between high-level cognitive processing and motor processing and provides evidence in support of a highly interactive, dynamic blending of cognition and action.

Methods

All protocols were approved by Yale University's Institutional Review Board, protocol number 2000027351.

Participants

We recruited participants for Experiments 1–5 from the Yale University undergraduate community (Experiment 1, $N = 44$; modified structure control, $N = 29$; Experiment 2, $N = 34$; Experiment 3, $N = 28$; Experiment 4, $N = 40$; Experiment 5, $N = 20$), and all experiments were conducted in line with a protocol approved by the university's Institutional Review Board. All participants provided informed consent prior to the initiation of any study protocols. The participants received course credit for their participation. All participants reported that they were not colour-blind and had normal or corrected-to-normal vision. We planned a priori to exclude participants that did not show reliable evidence of learning by excluding participants that did not show over 25% accuracy for at least four of the eight visuomotor associations (total exclusions: Experiment 1, $N = 3$; modified structure control, $N = 1$; Experiment 2, $N = 0$; Experiment 3, $N = 1$; Experiment 4, $N = 4$; Experiment 5, $N = 1$). There were also a small number of exclusions due to technical issues (Experiment 1, $N = 1$; Experiment 2, $N = 1$; Experiment 3, $N = 1$). Additionally, we planned to exclude participants who were not attentive to the task by excluding participants who did not respond on at least 75% of the trials in the learning task for Experiments 1–3 and in the learning or forced-response task in Experiments 4 and 5. No participants met this exclusion criterion. After exclusions, we had 40 participants in Experiment 1 (N female, 17; mean age, 20.1 years), 28 participants in the modified structure control (N female, 20; mean age, 19.1 years), 33 participants in Experiment 2 (N female, 18; mean age, 19.7 years), 26 participants in Experiment 3 (N female, 16; mean age, 19.3 years), 36 participants in Experiment 4 (N female, 19; mean age, 19.6 years) and 19 participants in Experiment 5 (N female, 15; mean age, 20.2 years).

Experiment 1

Task design. Experimental sessions for Experiment 1 were approximately one hour long and consisted of an RT baseline task, a task training phase and the learning task. The task was coded in jsPsych (version 6.1.0)[62]. All data reported in this manuscript were collected on a Lenovo IdeaPad 5 (Ubuntu 22.04).

RT baseline task. The participants first completed an RT baseline task to measure intrinsic finger-to-finger switch costs without the influence of the learning task. They performed the baseline task again following the learning task. Both instances of this task were identical and lasted approximately 5 min. During this task, the participants used their left hand (on keys A, S, D and F) and right hand (on keys H, J, K and L) to respond to the position of a target green square on the screen. On each trial, they would see eight squares on the screen that were spatially aligned with their fingers on the keyboard. Seven of the squares were white, while one square was green. Their goal was to press the key aligned with the green square. Once they had made the correct response, all squares turned white for 100 ms before another square turned green to initiate the next trial. The next trial did not begin until the participant made the correct response. We used this constraint to avoid having participants rapidly responding with incorrect responses to expedite the task. Only trials where participants made the correct response on their first attempt were included in analysis. Additionally, we excluded the first five trials in the task from analysis to account for slowed RTs at the beginning of the block. Trial sequences included all pairwise transitions between fingers (including repeating the same finger) a minimum of four times to ensure a stable switch cost estimate for each pairwise transition between fingers.

Learning task. After the motor task, the participants were familiarized with the structure of the learning task with 15 trials of a simplified version of the task. During this practice phase, the participants were instructed to use trial-by-trial feedback to learn the correct key to press (H, J or K) in response to three highly distinguishable emoji stimuli. The participants would see an emoji on every trial and then guess a response before receiving binary feedback as to whether their response was correct or not. Trial duration was unrestricted in this phase.

After familiarization with the basics of the learning task, the participants returned their hands to the keys that they had used during the RT baseline task (left hand: A, S, D, F; right hand: H, J, K, L). Before the learning task began, the participants saw an instruction screen with eight stimuli that would be used during the task arranged in a random order on the screen. Once the task began, the participants saw one stimulus per trial and used trial-by-trial feedback to learn the correct button to press in response to each stimulus. Correct stimulus–response associations were deterministic. The sequence of stimuli was random, such that every trial was independent. On every trial, the participants would see a single stimulus, make a response and then get feedback as to whether their response was correct or not (feedback duration, 750 ms; Fig. 1c). The next trial would proceed after the feedback from the previous trial disappeared. If the participant did not make a response within 2.5 s of viewing the stimulus, the trial would time out, and the participant would receive feedback that they needed to respond more quickly. The participants saw each of the eight stimuli at least 125 times (that is, 125 iterations of each stimulus) during the learning task.

Stimulus and visuomotor mapping design. Each stimulus varied along three features: colour (red, orange, blue or purple), shape (square, circle, triangle or diamond) and pattern (vertical stripes, diagonal stripes, dots or checkerboard). We randomly selected two possible values for each feature (for example, red and blue, square and circle, vertical stripes and dots) for each participant. Thus, all of the combinations of specific features yielded eight unique stimuli per participant.

To embed structure into the task, we assigned each feature to a level of an intuitive motor hierarchy, such that one feature indicated what hand to respond with (top-level), another feature dictated a pair of fingers within each hand or 'couplet' (mid-level) and the remaining feature could be used to determine the correct response within a couplet (low-level; Fig. 1b). For example, if colour was associated with the top level, then all stimuli of one colour would be associated with responses in the right hand and all stimuli of the other colour would be associated with responses in the left hand. Mid- and low-level features were assigned from left to right in extrinsic space (Fig. 1b). There were six possible assignments of features to level (colour > shape > pattern; colour > pattern > shape; shape > colour > pattern; shape > pattern > colour; pattern > shape > colour; pattern > colour > shape), and we counterbalanced the assignment of features to level across participants. The participants were never instructed about the structure.

Modified structure control experiment

Task design. For the modified structure control experiment, we changed the task in Experiment 1 to measure whether other hierarchical structures were learnable to participants or whether the results of Experiment 1 were driven by the specific latent structure that we trained the participants on. The task design was identical to that of Experiment 1 except for three key details. First, the participants executed the baseline task only once at the beginning of the experiment, rather than before and after the learning task, to reduce the length of the experiment. Second, the task duration was shorter (approximately 30 min), and the participants saw only 55 presentations of each of the eight stimuli (a point at which participants had generally reached asymptotic performance in Experiment 1). Finally and most importantly, we modified the hierarchical structure of the visuomotor mapping: instead of feature values being assigned spatially in extrinsic space from left to right, we aligned the structure with the mirror symmetry of the motor system (Supplementary Fig. 2a). For example, while the stimuli associated with the left pinkie and right index finger shared mid- and low-level features in Experiment 1, the left and right index fingers shared the mid- and low-level features in this experiment. As in Experiment 1, we counterbalanced the assignment of features to task levels. The task was coded in jsPsych[62] (version 6.1.0).

Experiment 2

Task design. Task design for Experiment 2 was identical to that of Experiment 1, except for the stimuli and structure of the visuomotor mapping. The participants were assigned one of two possible stimulus sets in this experiment. In this case, stimuli varied only along one feature (rather than three), either colour or shape. Some participants saw eight squares of different colours during the task (Fig. 3e; orange, green, yellow, red, blue, pink, brown and purple), and others saw eight different shapes that were all black (square, circle, plus, diamond, pentagon, triangle, crescent moon and star). This change meant that there was no learnable visuomotor structure embedded into the task. All other details were the same as in Experiment 1 (RT baseline task before and after learning, 125 iterations of each stimulus during the learning task). The task was coded in jsPsych[62] (version 6.1.0).

Experiment 3

Task design. The overall task design for Experiment 3 was the same as in the previous three experiments. In this experiment, the participants saw 55 iterations of each stimulus and performed the RT baseline task only once in the beginning of the session. We shortened the duration of the experiment to 55 iterations as participants in longer experiments had generally reached asymptotic performance at this point. We again used the three-feature stimuli described for Experiment 1; however, in this experiment, we created pseudorandomized mappings that minimized the amount of learnable intuitive motor structure in the mapping (Fig. 3i). Our aim was to assess whether the behaviour patterns from participants trained on the structured mappings arose from the use of the three-feature stimuli, rather than from the latent structure that we had embedded into the task. We opted for pseudorandomized mappings rather than fully randomizing the stimulus–response associations because randomizing the limited number of features and stimuli often created somewhat structured mappings (for example shuffled mappings, see Supplementary Fig. 4). The task was coded in jsPsych[62] (version 6.1.0).

Experiment 4

Task design. In Experiment 4, the participants started with the RT baseline task before moving into the learning task. The participants saw approximately 55 iterations of each stimulus in this task. We chose this number of iterations because the previous studies suggested that 55 iterations was sufficient exposure for learning, and this duration allowed us sufficient time for the forced-response task (see below) before participants were too fatigued. Visuomotor mappings followed the same structure as explained in Experiment 1 (Supplementary Fig. 5a). After the participants had learned the mapping during the learning task, they performed a forced-response task with the learned associations. Experiments 4 and 5 were coded in Octave[63] (version 6.4.0) using PsychToolbox[64] (version 3.0.18).

Forced-response task. During the forced-response task, the participants heard four ascending beeps (400 ms apart) on each trial (Fig. 4a). The participants were instructed to time their response with the fourth beep, regardless of whether they felt prepared to respond. We varied the time point at which the stimulus appeared on the screen during the beeps to manipulate the amount of preparation time (PT) participants had to make their responses on a trial-by-trial basis. PT is defined as the interval between when the stimulus appeared and the last beep of the trial. PTs were randomly selected from a uniform distribution from 100 ms to 1.2 s. Thus, on some trials, participants would have sufficient time to prepare, while on others they would have to prepare very rapidly (or guess). The participants were encouraged to respond at the appropriate time on each trial even if they felt that they were guessing. After the participants made a response, they received feedback (750 ms) on whether their response was correct or not and whether they had responded in time with the fourth beep. The participants had a ±50-ms cushion from the exact instructed timing within which they would receive positive timing feedback.

The participants were familiarized with the forced-response task using the same emoji stimuli that were used to familiarize them with the learning task. This forced-response familiarization period occurred after the learning task practice and before the main learning task. After the learning task was completed, the participants executed approximately 760 forced-response trials (95 iterations per stimulus), and the task took approximately 25 min. The participants had the option to take self-timed breaks following each 100-trial block.

Experiment 5

Task design. Experiment 5 followed the same protocol as Experiment 4 (including the shortened learning task duration to accommodate the forced-response task), with one change. The difference from Experiment 4 was that participants were trained (Supplementary Fig. 5f) on unstructured mappings using the stimuli from Experiment 2 (eight squares of different colours or eight black shapes, counterbalanced). Again, the participants first executed the RT baseline task, followed by familiarization with the learning and forced-response paradigms. After familiarization, the participants learned the visuomotor mapping during the learning task. The session ended with the forced-response paradigm described above.

Experiment 6

Participants. We recruited a separate group of participants to participate in the longitudinal version of the task. All participants offered informed consent prior to starting the study. In total, we recruited 27 participants to take part in this study. Seven participants were unable to complete the eight sessions due to illness, withdrawal from the study or non-responsiveness, and we excluded them from analysis. Of the remaining 20 participants ($N = 17$ female; mean age, 19.8 years; range, 18–22 years), two participants missed one session. We opted to still include these participants in our analyses, given the difficulty of collecting the longitudinal data. Additionally, two participants had to reschedule their last sessions for two days after the intended eighth session due to illness. We also opted to include these participants in our analyses. The participants were compensated US\$15 h^{-1} for their time (US\$10 for in-lab sessions and US\$15 for six online sessions) and received a US\$10 bonus for completing the whole study at the end of the eight days.

Task design. The task closely resembled the design of Experiments 1 and 4. The goal of this study was to have the participants practise the associations over eight days to assess whether the structured dynamics that we observed in Experiment 4 would be attenuated by extensive practice. In other words, do participants use the latent structure of the mapping early in learning and then compress to more direct stimulus–response associations once they have additional practice?

To answer this question, we had the participants practise a hierarchically structured mapping (as in Experiments 1 and 4) for eight consecutive days. On the first and last days of the experiment, the participants came in-person to the lab to participate in a 30-min session. This session consisted of the same training procedure as the previous experiments, a motor baseline phase and the learning task (55 iterations of each stimulus as in Experiments 3, 4 and 5). During Days 2–7, the participants were emailed a link each morning for 10 min of online practice with the mapping. The practice phase was identical to the learning task other than being shorter. If participants did not complete the task, they were reminded in the afternoon and again in the evening to complete their practice session.

We also ran a subset of the included participants ($N = 9$) on the forced-response task during the first and last sessions of their participation. For these participants, the first and last sessions were identical to Experiment 4.

Survey data collection

In addition to our six experiments, we collected survey data from 37 participants (N female, 29; mean age, 18.8 years) to measure the participants' intuitions about the learnability of structured versus unstructured mappings in our task. We note that we did not directly measure the intuitiveness of the structured mappings but rather used learnability as a proxy for intuitiveness. The participants were recruited from the Yale undergraduate community to participate in a different study that used the same three-feature stimuli as the tasks presented here. The task that they completed had the same general structure as Experiment 3—participants used trial-by-trial feedback to learn the correct response to eight three-feature stimuli with no underlying structure to the mappings. There were some differences in this protocol, such as a surprise memory test at the end of the session and a subset of stimuli that were presented less frequently than the other stimuli. The survey was presented at the end of the experimental session and took about 5 min to complete. Thus, the participants were familiar with the general task when they responded to the survey but had not been exposed to a structured mapping. In this way, we measured participants' intuitions about how latent structure affects the learnability of these stimulus–key mappings.

The survey presented the participants with two ways to indicate whether they thought structured mappings were easier to learn. First, the participants were asked to create an eight-to-eight stimulus–key mapping that they thought would be "easiest to learn for a new participant". On this question, the participants saw eight three-feature stimuli in a random order and selected one of the eight possible key responses for each stimulus. They were also asked to write a few sentences explaining why they chose the assignment that they had created. Second, we asked the participants to rank eight mappings in order from easiest to hardest to learn. For this question, we displayed eight stimuli above pictures of the keys that participants used during the task (left hand: A, S, D, F; right hand: H, J, K, L). We included each of the six possible counterbalanced assignments of feature to hierarchy level (for example, shape → hand, colour → finger-couplet, pattern → finger) and two unstructured mappings in the options. The mappings were all arranged in external space (as in Experiments 1 and 4). We also asked the participants to explain why they thought that the mappings that they put at first and last would be the easiest and hardest to learn. We focus on this second question in this paper because it speaks more directly to the intuitiveness of the embedded structures.

Analysis

The data and analysis code are available at https://github.com/jetrach/StructuredActionPrepVMDM.

Motor correction. We operationalized structure in the learned visuomotor mappings by linking visual features to intuitive groupings of actions. Because our main interest was how the previous trial's action affected the current trial, we had one major confound to contend with: intrinsic switch costs between the different fingers of each hand. We thus wanted to ensure that any RT effects that we observed as a result of learning the visuomotor mapping were due to the structure of the mapping, rather than generic spatial or biomechanical influences on transitional RTs between fingers. To do this, we calculated the mean RT for each of the 64 pairwise transitions between fingers during the RT baseline task. We then subtracted these RTs from the RTs of trials of the same finger-to-finger transitions during the learning task (Fig. 1e). We thus removed variance in the RTs that was present during the baseline task to isolate the impact of the learning task structure on RT. We performed our further analyses on these baseline-corrected RTs. We excluded the first five trials of the task and RTs that were especially slow, indicating that the participant was not attending to the task (3 s.d. above the mean RT).

Learning task. Our primary analyses in Experiments 1–3 were based on a straightforward logic: that after learning, RTs for correct responses would be influenced by the previous trial in a manner dictated by the learned visuomotor mapping, even though the sequence of stimuli across trials was randomized (that is, every trial was independent). That is, we reasoned that transitional RTs would spontaneously reflect the structure of the visuomotor mappings that people learned[23]. Analyses were conducted in R[65] (version 4.2.1, 2022) or MATLAB[66] (version 9.13.0, 2022a, update 4).

We used repeated-measures ANOVAs and linear mixed-effects models to analyse our data. Additionally, we used two-tailed one- or two-sample *t*-tests or paired *t*-tests where appropriate and corrected for multiple comparisons using a Bonferroni correction. We used Welch *t*-tests when assumptions of normality were violated. We used Cohen's *d* to quantify effect sizes for *t*-tests[67] (effsize package in R, version 0.8.1) and η^2 for ANOVAs. We report 95% CIs on the mean difference for *t*-tests and on partial eta squared (η^2) for ANOVAs. For experiments with structured mappings (Experiments 1 and 4 and the modified structure experiment), we combined across all feature-level mappings as we had no a priori hypotheses for how these assignments might affect the RT results. We statistically justified this choice by conducting ANOVAs within experiments to compare overall RT and accuracy across feature-level assignments (Experiment 1: accuracy: $F_{5,34} = 0.67$, $P = 0.647$; RT: $F_{5,34} = 1.3$, $P = 0.287$; Experiment 4: accuracy: $F_{5,30} = 0.47$, $P = 0.799$; RT: $F_{5,30} = 1.84$, $P = 0.135$). Transitional RT analyses were performed on trials where participants responded correctly to the current and preceding trials (consecutively correct trials). We did this to ensure we were examining response dynamics after the participants had sufficiently learned the mapping, and to avoid confounds from post-error slowing that can occur in reinforcement learning settings[68]. We note, however, that the main results do not qualitatively change if we include all correct trials without conditioning on the previous trial being correct. We removed outlier RTs by excluding RTs under 200 ms where participants would not have had sufficient time to respond. In addition, we excluded the first three trials for each participant to account for task initiation costs.

We designed three linear mixed-effects models to operationalize our three main theoretical models of behaviour. The Hierarchical model used the number of graph edges or path distance between responses (0, 2, 4 or 6) on a given pair of successive trials to predict corrected RTs. If the top-level feature (for example, shape) changed across trials, there were six graph edges between consecutive responses. On trials where the top-level feature repeated but the mid-level feature switched, there were four graph edges between consecutive responses. When the low-level feature switched, there were two edges between responses, and on trials where the exact stimulus repeated there were zero edges between responses. We used the structure we embedded into the mapping to calculate these distances in Experiment 1 and the modified structure control experiment. In Experiments 2 and 3, where there was no latent structure in the mapping, we used the extrinsic space hierarchical structure to calculate these path distances. This approach allowed us to rule out alternative accounts of hierarchical effects in Experiment 1.

For the Feature-Based model, we used the number of visual features that changed between the previous and current trials (0, 1, 2 or 3) to predict baseline-corrected RTs. Importantly, the stimulus and mapping design ensure that the Hierarchical and Feature-Based models do not predict the same RT behaviour. For instance, trials that are classified as six-edge paths in the Hierarchical model could have one, two or three features that change across that stimulus transition. Similarly, four-edge trials can have either one or two stimulus features changing across the transition. The two models thus make dissociable predictions about participant behaviour. For the Flat model, we modelled whether the stimulus repeated or switched (0 or 1) to predict baseline-corrected RTs. In addition to these three theoretically motivated models, we fit two models that considered the physical distance between responses in difference ways. We fit a Physical Distance model that counted the number of fingers between responses to predict corrected RT (0–7). This model operationalizes a spatially modulated attentional effect where participants are faster to make responses that are adjacent to their previous response. We also fit a Nearest-Neighbour model that predicts facilitated performance on transitions to neighbouring responses, but does not assume a linear effect that extends across both hands (as the Physical Distance model does; 0 for repeats, 1 for adjacent responses and 2 for all other responses). We included random intercepts and slopes for each participant and compared model fits using the BIC.

Forced-response task (Experiments 4 and 5). Our analyses for the forced-response task examined the probability of different types of errors that participants made as a function of PT. To do this, we first calculated the actual PT that participants had on each trial by adding their RT to the planned PT that was hard-coded into the trial (that is, the interval between the visual stimulus appearance and the fourth tone in the countdown). For example, if the stimulus was displayed 700 ms before the response cue on a given trial, and the participant made their response 50 ms after the response cue, the actual PT on that trial would be 750 ms. Similarly, if the participant responded 50 ms before the cue, then the actual PT on that trial would be 650 ms. We did this to quantify PT more accurately on a trial-by-trial basis.

In Experiment 4, we classified errors on the basis of the highest feature level where there was a mismatch between the target stimulus and the stimulus associated with the response that the participant made (Fig. 4b). Thus, there was one correct response, one low-level error, two mid-level errors and four top-level errors possible on each trial. We normalized chance probabilities across error levels by dividing the raw probability of errors at each level by the number of responses associated with that level. In addition to this approach, we conducted primary analyses with errors coded for shared features between the target stimulus and the stimulus mapped to the response that the participant made at each level of the task (that is, without combining probabilities within error levels). Visualizations of this approach are depicted in Supplementary Fig. 7. We excluded trials where participants did not respond within 100 ms of the response cue. Because Experiment 5 operates as a control experiment to rule out generic motor/spatial explanations of our forced-response-time results, we applied the error coding scheme that we used in Experiment 4 to the responses in Experiment 5, as if those stimuli were structured in the same way. To visualize the PT results, we calculated the average probability of making each type of error in a 100-ms sliding window that was moved across the full range of PTs.

Response preparation models. We modified a previous model of response selection[18] to formalize three theoretical models of action selection in our forced-response task. The basic model assumes that the time (T) it takes to prepare an action a is described by a normal distribution with a mean μ and standard deviation σ:

$$T_a = \mathrm{N}(\mu_a, \sigma_a) \tag{1}$$

By taking the cumulative density of this distribution, we get a sigmoidal function p_a that describes the probability of having prepared action a at any given PT. The probability that a given response is prepared is thus dependent on the amount of PT and the μ and σ parameters that describe the response preparation distribution.

We fit four extensions of this straightforward model to the participant data. The first model—the Flat model—assumes that the preparation of the correct action a simply involves an increase in the probability of selecting that action over time (that is, the cumulative density function of equation (1), p_a). Critically, the model treats the selection of any of the seven other actions as equally probable at each time point. This model captures the idea that even if participants mentally represent the structure of the task, only a single action is potentiated at any time during action selection; this would be consistent with a strict separation of deciding on the correct stimulus–response association and preparing actual motor commands. Moreover, this model would be the best candidate for unstructured versions of the visuomotor mapping that have no embedded structure (Experiment 5).

In our second model variant—the Hierarchical model—we extended equation (1) to the preparation of 'groupings' of actions associated with each level, j, of the structured visuomotor mapping (Experiment 4):

$$T_j = \mathrm{N}(\mu_j, \sigma_j) \tag{2}$$

with each level having its own μ and σ free parameters, and where each μ describes how long, on average, it takes for a participant to 'resolve' level j of the learned mapping and prepare the relevant set of actions. Specifically, preparing the top level involves preparing all four actions on the correct hand, preparing the middle level involves preparing the correct couplet on each hand and preparing the lowest level involves preparing the correct left-versus-right finger position across all couplets. According to this model, if sequential resolving from top to bottom of each level of the visuomotor mapping potentiates the associated motor commands in real time, the fitted μ parameters should take the lowest values for the top/hand level (that is, the top is resolved first), middling values for the middle/couplet level and the highest values for the low/finger level. In other words, according to this structured action preparation model, the participant arrives at the correct action by sequentially 'pruning' the visuomotor mapping in real time. We fit an additional variant of this model that included only one σ free parameter that was used at all levels of the structure to see if variance in PTs was comparable across levels.

Lastly, we fit an additional variant of the model to formalize a Feature-Based model of action selection. Here we only allowed for one μ and one σ parameter but maintained the structured preparation of action groupings; this model thus assumes that the learned feature–action associations shaped action selection but with no temporal prioritization of any particular features/levels.

Finally, we note that people often have to guess in the forced-response task given the strict temporal criteria. The goal-oriented action preparation processes described above are thus mixed at each time point with a guessing or 'lapse' process that assumes a uniform probability of any of the eight possible actions being selected, with a mixture parameter ρ that determines the weighting of guessing versus goal-oriented action preparation. This mixture model thus determines the final speed–accuracy probability function P of selecting action a:

$$P_a = (1 - \rho) \times p_a + \rho \times 1/8 \tag{3}$$

Response probabilities P generated by each model at each PT (rounded to the nearest ms) were fit directly to actual participant responses and PTs measured in the forced-response tasks (Experiments 4 and 5). Parameter fits were optimized using the fmincon function in MATLAB. We ran 50 iterations of each fitting procedure for each participant to avoid local minima in the optimized model fits. To simulate the results, we computed the model's response probability functions using the best fit parameters for each participant and then averaged the resulting curves over all participants.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data that support the findings of this study are available via GitHub at https://github.com/jetrach/StructuredActionPrepVMDM.

Code availability

The code used is available via GitHub at https://github.com/jetrach/StructuredActionPrepVMDM. The task code is available upon request. Please refer to the Methods for details on the software used in this project.

References

1. Trach, J. E., McKim, T. H. & Desrochers, T. M. Abstract sequential task control is facilitated by practice and embedded motor sequences. *J. Exp. Psychol. Learn. Mem. Cogn.* **47**, 1638–1659 (2021).
2. Peer, M., Brunec, I. K., Newcombe, N. S. & Epstein, R. A. Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn. Sci.* **25**, 37–54 (2021).
3. Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M. M. Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).
4. Schuck, N. W., Cai, M. B., Wilson, R. C. & Niv, Y. Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* **91**, 1402–1412 (2016).
5. Kahn, A. E., Karuza, E. A., Vettel, J. M. & Bassett, D. S. Network constraints on learnability of probabilistic motor sequences. *Nat. Hum. Behav.* **2**, 936–947 (2018).
6. Rosenbaum, D. A., Kenny, S. B. & Derr, M. A. Hierarchical control of rapid movement sequences. *J. Exp. Psychol. Hum. Percept. Perform.* **9**, 86–102 (1983).
7. Cisek, P. & Pastor-Bernier, A. On the challenges and mechanisms of embodied decisions. *Phil. Trans. R. Soc. B* **369**, 20130479 (2014).
8. Gallivan, J. P., Chapman, C. S., Wolpert, D. M. & Flanagan, J. R. Decision-making in sensorimotor control. *Nat. Rev. Neurosci.* **19**, 519–534 (2018).
9. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
10. Selen, L. P. J., Shadlen, M. N. & Wolpert, D. M. Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J. Neurosci.* **32**, 2276–2286 (2012).
11. Shadlen, M. N. & Kiani, R. Decision making as a window on cognition. *Neuron* **80**, 791–806 (2013).
12. Song, J.-H. & Nakayama, K. Target selection in visual search as revealed by movement trajectories. *Vis. Res.* **48**, 853–861 (2008).
13. Song, J.-H. & Nakayama, K. Numeric comparison in a visually-guided manual reaching task. *Cognition* **106**, 994–1003 (2008).
14. Song, J.-H. & Nakayama, K. Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn. Sci.* **13**, 360–366 (2009).
15. Spivey, M. J. & Dale, R. Continuous dynamics in real-time cognition. *Curr. Dir. Psychol. Sci.* **15**, 207–211 (2006).
16. Spivey, M. J., Grosjean, M. & Knoblich, G. Continuous attraction toward phonological competitors. *Proc. Natl Acad. Sci. USA* **102**, 10393–10398 (2005).

17. Ghez, C. et al. Discrete and continuous planning of hand movements and isometric force trajectories. *Exp. Brain Res.* **115**, 217–233 (1997).
18. Hardwick, R. M., Forrence, A. D., Krakauer, J. W. & Haith, A. M. Time-dependent competition between goal-directed and habitual response preparation. *Nat. Hum. Behav.* **3**, 1252–1262 (2019).
19. Mcdougle, S. D. & Taylor, J. A. Dissociable cognitive strategies for sensorimotor learning. *Nat. Commun.* **10**, 40 (2019).
20. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1467 (2016).
21. Musslick, S. & Bizyaeva, A. Examining cognitive flexibility and stability through the lens of dynamical systems. *Curr. Opin. Behav. Sci.* **57**, 101375 (2024).
22. Collins, A. G. E. & Frank, M. J. Motor demands constrain cognitive rule structures. *PLoS Comput. Biol.* https://doi.org/10.1371/journal.pcbi.1004785 (2016).
23. Dykstra, T., Smith, D. M., Schumacher, E. H. & Hazeltine, E. Measuring task structure with transitional response times: task representations are more than task sets. *Psychon. Bull. Rev.* https://doi.org/10.3758/s13423-021-02035-3 (2022).
24. Reitman, J. S. & Rueter, H. H. Organization revealed by recall orders and confirmed by pauses. *Cogn. Psychol.* **12**, 554–581 (1980).
25. Maechler, M. R., Rousseeuw, P., Struyf, A., Hubert, M. & Hornik, K. cluster: Cluster analysis basics and extensions. R package version 2.1.3 (2022).
26. Kassambara, A. & Mundt, F. factoextra: Extract and visualize the results of multivariate data analyses. R package version 1.0.7 (2020).
27. Brown, T. I. et al. Prospective representation of navigational goals in the human hippocampus. *Science* **352**, 1323–1326 (2016).
28. Rmus, M., Ritz, H., Hunter, L. E., Bornstein, A. M. & Shenhav, A. Humans can navigate complex graph structures acquired during latent learning. *Cognition* **225**, 105103 (2022).
29. Tavares, R. M. et al. A map for social navigation in the human brain. *Neuron* **87**, 231–243 (2015).
30. Karuza, E. A., Kahn, A. E., Thompson-Schill, S. L. & Bassett, D. S. Process reveals structure: how a network is traversed mediates expectations about its architecture. *Sci. Rep.* **7**, 12733 (2017).
31. Rosenbaum, D. A. Human movement initiation: specification of arm, direction, and extent. *J. Exp. Psychol. Gen.* **109**, 444–474 (1980).
32. Jaffe, P. I., Poldrack, R. A., Schafer, R. J. & Bissett, P. G. Modelling human behaviour in cognitive tasks with latent dynamical systems. *Nat. Hum. Behav.* https://doi.org/10.1038/s41562-022-01510-8 (2023).
33. Balaguer, J., Spiers, H., Hassabis, D. & Summerfield, C. Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron* **90**, 893–903 (2016).
34. Ekstrom, A. D. & Ranganath, C. Space, time, and episodic memory: the hippocampus is all over the cognitive map. *Hippocampus* **28**, 680–687 (2018).
35. Theves, S., Neville, D. A., Fernández, G. & Doeller, C. F. Learning and representation of hierarchical concepts in hippocampus and prefrontal cortex. *J. Neurosci.* **41**, 7675–7686 (2021).
36. Whittington, J. C. R. et al. The Tolman–Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249–1263.e23 (2020).
37. Wu, C. M., Schulz, E., Garvert, M. M., Meder, B. & Schuck, N. W. Similarities and differences in spatial and non-spatial cognitive maps. *PLoS Comput. Biol.* **16**, e1008149 (2020).
38. Zhou, J. et al. Complementary task structure representations in hippocampus and orbitofrontal cortex during an odor sequence task. *Curr. Biol.* **29**, 3402–3409.e3 (2019).
39. Burk, D., Ingram, J. N., Franklin, D. W., Shadlen, M. N. & Wolpert, D. M. Motor effort alters changes of mind in sensorimotor decision making. *PLoS ONE* **9**, e92681 (2014).
40. Cisek, P. & Kalaska, J. F. Simultaneous encoding of multiple potential reach directions in dorsal premotor cortex. *J. Neurophysiol.* **87**, 1149–1154 (2002).
41. Hagura, N., Haggard, P. & Diedrichsen, J. Perceptual decisions are biased by the cost to act. *eLife* **6**, e18422 (2017).
42. Moher, J. & Song, J.-H. Perceptual decision processes flexibly adapt to avoid change-of-mind motor costs. *J. Vis.* **14**, 1 (2014).
43. Selen, L. P. J., Corneil, B. D. & Medendorp, W. P. Single-trial dynamics of competing reach plans in the human motor periphery. *J. Neurosci.* **43**, 2782–2793 (2023).
44. Thura, D. & Cisek, P. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* **81**, 1401–1416 (2014).
45. Wong, A. L. & Haith, A. M. Motor planning flexibly optimizes performance under uncertainty about task goals. *Nat. Commun.* **8**, 14624 (2017).
46. Freeman, J. B., Nakayama, K. & Ambady, N. Finger in flight reveals parallel categorization across multiple social dimensions. *Soc. Cogn.* **31**, 792–805 (2013).
47. Magnuson, J. S. Moving hand reveals dynamics of thought. *Proc. Natl Acad. Sci. USA* **102**, 9995–9996 (2005).
48. Alhussein, L. & Smith, M. A. Motor planning under uncertainty. *eLife* **10**, e67019 (2021).
49. Haith, A. M., Huberdeau, D. M. & Krakauer, J. W. Hedging your bets: intermediate movements as optimal behavior in the context of an incomplete decision. *PLoS Comput. Biol.* **11**, e1004171 (2015).
50. Miller, J. Discrete versus continuous stage models of human information processing: in search of partial output. *J. Exp. Psychol. Hum. Percept. Perform.* **8**, 273–296 (1982).
51. Cellier, D., Petersen, I. T. & Hwang, K. Dynamics of hierarchical task representations. *J. Neurosci.* https://doi.org/10.1523/JNEUROSCI.0233-22.2022 (2022).
52. Ranti, C., Chatham, C. H. & Badre, D. Parallel temporal dynamics in hierarchical cognitive control. *Cognition* **142**, 205–229 (2015).
53. Badre, D. & D'Esposito, M. Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat. Rev. Neurosci.* **10**, 659–669 (2009).
54. Badre, D. & Nee, D. E. Frontal cortex and the hierarchical control of behavior. *Trends Cogn. Sci.* **22**, 170–188 (2018).
55. Logan, G. D. Toward an instance theory of automatization. *Psychol. Rev.* **95**, 492–527 (1988).
56. Logan, G. D. Automatic control: how experts act without thinking. *Psychol. Rev.* **125**, 453–485 (2018).
57. Broadbent, D. E. & Broadbent, M. H. P. From detection to identification: response to multiple targets in rapid serial visual presentation. *Percept. Psychophys.* **42**, 105–113 (1987).
58. McLean, J. P., Broadbent, D. E. & Broadbent, M. H. P. Combining attributes in rapid serial visual presentation tasks. *Q. J. Exp. Psychol.* https://doi.org/10.1080/14640748308402123 (1983).
59. Treisman, A. Focused attention in the perception and retrieval of multidimensional stimuli. *Percept. Psychophys.* **22**, 1–11 (1977).
60. Behrens, T. E. J. et al. What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* **100**, 490–509 (2018).
61. Chrastil, E. R. & Warren, W. H. From cognitive maps to cognitive graphs. *PLoS ONE* **9**, e112544 (2014).
62. de Leeuw, J. R., Gilbert, R. A. & Luchterhandt, B. jsPsych: Enabling an open-source collaborative ecosystem of behavioral experiments. *J. Open Source Softw.* **8**, 5351 (2023).
63. Eaton, J. W. GNU Octave and reproducible research. *J. Process Control* **22**, 1433–1438 (2012).
64. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).

65. R Core Team. R: A language and environment for statistical computing (R Foundation for Statistical Computing, 2022); https://www.R-project.org/
66. MATLAB version 9.13.0 (R2022a) (The MathWorks Inc., 2022); https://www.mathworks.com
67. Torchiano, M. effsize: Efficient effect size computation. R package version 0.8.1 (2020).
68. McDougle, S. D. Post-error slowing during instrumental learning is shaped by working memory-based choice strategies. *Neuroscience* **486**, 37–45 (2022).

Acknowledgements

We thank A. Forrence and T. Adanri in particular for their help with these projects. In addition, we acknowledge M. Ingram, J. Burde, K. Chou, S. Hu, L. Pandolpho and O. Pilkinton for essential support in collecting data. We also thank T. Desrochers and J. Taylor for feedback on early drafts of the manuscript, and the members of the ACT lab and Turk-Browne Lab at Yale University for productive conversations about this work throughout. J.E.T. is supported by the NSF GRFP. S.D.M. is supported by NIH grant no. R01 NS132926.

Author contributions

J.E.T. and S.D.M. designed the paradigm. J.E.T. collected the data. J.E.T. and S.D.M. analysed the data. J.E.T. and S.D.M. prepared the figures and drafted, edited and revised the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41562-025-02217-2.

Correspondence and requests for materials should be addressed to Juliana E. Trach.

Peer review information *Nature Human Behaviour* thanks Aymar de Rugy, Elliot Hazeltine and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

nature portfolio

Corresponding author(s): Juliana E. Trach

Last updated by author(s): Jan 30, 2025

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☐	☒	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☐	☒	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	We used custom task code to collect data. Experiments 1-4 were coded in jsPsych, version 6.1.0. Experiments 5 and 6 were coded in PsychToolBox, version 3.0.18, in Octave, version 6.4.0.
Data analysis	Analyses were conducted in R (R version: 4.2.1, 2022) or MATLAB (Version: 9.12.0, 2022a, Update 4).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

Data is available at https://github.com/jetrach/StructuredActionPrepVMDM.

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	We did not conduct any sex- or gender-based analyses.
Reporting on race, ethnicity, or other socially relevant groupings	We did not report on socially-relevant groupings.
Population characteristics	See above.
Recruitment	For Experiments 1-5, participants were recruited from the subject pool for Introductory Psychology at Yale University. Researchers posted time slots and subject pool students signed up for slots that fit within their schedule. The Introductory Psychology course represents a relatively diverse cross-section of the Yale undergraduate population. We do not anticipate any influence of self selection bias, as the study descriptions that potential participants were presented with were identical across study versions. For Experiment 6, participants were recruited from the broader New Haven community. Once they contacted the lab expressing interest in participating, they were presented with a study description and either opted to participate and scheduled their first session or opted out. For this experiment, sessions were scheduled directly with the researcher via phone or email.
Ethics oversight	The Yale Institutional Review Board approved all protocol. All participants provided informed consent prior to the initiation of any study activities via a IRB-approved consent form.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Our studies used quantitative behavioral data collected during a learning/psychophysics experiment (e.g., RT, response accuracy). Demographic data for participant samples was collected but not used for analyses beyond describing the sample.
Research sample	Our sample includes Yale undergraduate students. Of the 217 participants that were run in the task, 202 (N female = 122; Experiment 1: N = 44, N female = 19, age range = 18-22; Experiment 2: N = 34, N female = 18, age range = 18-23; Experiment 3: N = 28, N female = 18, age range = 18-22; Experiment 4: N = 40, N female = 22, age range = 18-22; Experiment 5: N = 20, N female =16 , age range = 18-22; Experiment 6: N = 27, N female = 20, age range = 18-21; Structure control: N = 29, N female = 20, age range = 18-21; Survey: N = 37, N female = 29, age range = 18-21) contributed usable data. Participants were between 18 and 22 years of age. While this likely not a representative sample of the US population, it is quite typical of similar types of learning and psychophysics studies. We opted for this convenience sample as we did not have specific hypotheses about how other demographic variables might impact study results. The majority of participants were recruited from the Introductory Psychology participant pool and compensated with course credit (Experiments 1-5). This course's enrollment represents a relatively diverse cross section of the general Yale undergraduate population, including students across grade-levels and majors. Participants for experiment 6 were recruited from the New Haven community and compensated $15/hr plus a $10 completion bonus.
Sampling strategy	Our sample was a convenience sample of students in the subject pool at Yale University. We did not perform sample size calculations but instead chose sample sizes that were similar to others reported in related literature (Collins and Frank, 2012; Collins and Frank, 2016; Daniel, Radulescu, and Niv, 2020). For studies 1 and 2 (that were run prior to the modified structure experiment and study 3), we ran larger samples to ensure that we had sufficient data after exclusions (online pilot versions of the task led us to expect higher exclusion). We ran fewer participants in study 3, as exclusions were lower for these in-person samples. For experiments 4 and 5, we opted to run more of our allotted participant hours on Experiment 4, as those results were our main focus, and fewer on Experiment 5, as a control study. For experiment 6, we aimed to recruit 20 participants due to the difficulty of the multi-day design so we brought 27 people in to start the study.
Data collection	Data were primarily collected by the first author or a number of undergraduate research assistants (refered to in the acknowledgments). Data was collected at the Department of Psychology at Yale University at 2 Hillhouse Ave, New Haven, CT in small testing rooms. Experimenters used a common script (with minor changes based on task version) and remained in the room with the participant throughout the duration of the task. Participants were debriefed to the purpose of the study after completion as participation served as an educational requirement for Intro Psych. Experimenters were not blind to the purpose of the study, however they did not know what mapping each participant was assigned, making it impossible for them to influence the performance of the participant in important ways.

Timing	Each sample was collected within one semester. Experiment 1: 3/16/22-5/11/22 (Spring '22) Modified structure: 9/28/22-12/2/22 (Fall '22) Experiment 2: 3/28/22-5/4/22 (Spring '22) Experiment 3: 9/28/22-11/8/22 (Fall '22) Experiment 4: 10/7/22-12/6/22 (Fall '22) Experiment 5: 2/16/23-3/9/23 (Spring '23) Experiment 6: 10/26/22-3/31/23 (Fall '22-Spring 23')
Data exclusions	Taken from manuscript text: We planned a priori to exclude participants that did not show reliable evidence of learning by excluding participants that did not show over 25% accuracy for at least 4 of the 8 visuomotor associations (Total exclusions: Experiment 1: N = 3; Modified structure: N = 1; Experiment 2: N = 0; Experiment 3: N = 1; Experiment 4: N = 4; Experiment 5: N = 1). There were also a small number of exclusions due to technical issues (Experiment 1: N = 1; Experiment 2: N = 1; Experiment 3: N = 1). Additionally, we planned to exclude participants who were not attentive to the task by excluding participants who did not respond on at least 75% of the trials in the learning task for Experiments 1-3 and for the learning or forced response task in Experiments 5-6. No participants met this exclusion criterion. After exclusions, we had 40 participants in Experiment 1 (N female = 17, mean age = 20.1), 28 participants in Modified structure experiment (N female = 20, mean age = 19.1), 33 participants in Experiment 2 (N female = 18, mean age = 19.7), 26 participants in Experiment 3 (N female = 16, mean age = 19.3), 36 participants in Experiment 4 (N female = 19, mean age = 19.6), and 19 participants in Experiment 5 (N female = 15, mean age = 20.2). For experiment 6, we excluded participants that were unable to complete the study for any reason. Our final sample was 20 people (N female = 17, ages 18-22).
Non-participation	No participants revoked consent or declined to participate after signing up for an experimental session in the single-day studies. There were a small number of participants who no-showed for their sessions, however we did not keep systematic record of no-shows and no-shows were not included in any counts reported in the manuscript. For Experiment 6, a small number of participants withdrew consent by either indicating via email that they were unable to proceed or by non-responsiveness to contact. Out of 27 participants, seven withdrew consent either via email for reasons of illness or business in work/classes or stopped responding to emails to complete practice sessions.
Randomization	Counterbalancing of mapping (Experiments 1-5) or stimulus set (Experiments 2 and 5) were randomly determined by assigning numbers to each stimulus-response mapping and sequentially assigning participants to each mapping number. In this way, participants were randomly assigned to their specific task parameters. If two studies were being run during the same semester, participants could choose which experiment to sign up for (although it was not apparent to participants that the two studies were connected), however participation in one study during that semester precluded them from participation in the other study during that semester. This constraint was controlled by settings in the SONA platform to ensure that no human error would allow participants to do both studies. Participants recruited from the broader community (Experiment 6) were verified against lab records before beginning the multiday study to ensure they had not previously participated in any version of the task. There were no experimental groups in Experiment 6 -- all participants completed the same task with the same design.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Plants

Seed stocks	none
Novel plant genotypes	none
Authentication	none