UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Climbing the Tree: Structured Hierarchical Representations in Visuomotor Maps

Permalink

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Trach, Juliana E.
McDougle, Samuel David

Publication Date

2022

Peer reviewed

Climbing the Tree: Structured Hierarchical Representations in Visuomotor Maps

Juliana E. Trach (juliana.trach@yale.edu)
Department of Psychology, 2 Hillhouse Ave
New Haven, CT 06511 USA

Samuel D. McDougle (samuel.mcdougle@yale.edu)
Department of Psychology, 2 Hillhouse Ave
New Haven, CT 06511 USA

Abstract

Humans are uniquely adept at extracting structure from the world around them. It is well known that people often form hierarchical task representations during learning, even when a task does not explicitly necessitate a hierarchical representation. Still, how individuals capitalize on this structure to facilitate behavior is an open question. In the present study, we address this question by carefully examining patterns of response time switch costs in a hierarchically structured visuomotor association learning task, to adjudicate between multiple models of behavior. We find evidence that participants do appear to navigate through a hierarchical representation of stimulus-response associations as they prepare responses, rather than maintaining a non-hierarchical, flat model or being primarily affected by changes in stimulus features. These results establish the existence of hierarchical mental representations even for static visuomotor mappings, and imply that such representations are internally navigated in an orderly manner during action selection.

Keywords: hierarchy; structure learning; action selection; cognitive control; cognitive map

Introduction

An impressive aspect of human cognition is the ease with which we provide structure to our mental representations. For example, previous research demonstrates that humans tend to organize mental representations *hierarchically*. This has been observed in several domains of cognition, ranging from spatial reasoning (McNamara, 1986) and cognitive control (Badre et al., 2010), to object categorization (Theves et al., 2021; Collins & Quillian, 1969) and reinforcement learning (Collins & Frank, 2013). Moreover, individuals often spontaneously extract hierarchical representations from tasks even when hierarchy is not essential, or even beneficial, for accomplishing a task (Collins, 2017), further supporting its fundamental role in the organization of human behavior.

Hierarchies are especially useful when people learn about sequences of events that have a predictable temporal structure, as observed in decades of studies involving cued motor sequences (Kahn et al., 2018; Nissen & Bullemer, 1987), task sequences (Schneider & Logan, 2006; Trach et al., 2021), or passive learning of sequences of visual stimuli (Mark et al., 2020). Indeed, the abstract structuring of information with respect to time is thought to be a fundamental aspect of higher-level cognition (Behrens et al., 2018; Dehaene et al., 2015). Inspired by research on

temporally abstracted hierarchies, here, using a *de novo* visuomotor learning task with no sequential structure, we asked if and how abstract hierarchical representations are formed and navigated in "action space." In particular, this study builds on previous work by specifically examining how these hierarchical representations are *navigated*.

How can hierarchical mental representations be inferred from behavior? Previous studies have typically focused on "switch costs" at various levels of abstraction to study the learning and implementation of hierarchical structures (Korb et al., 2017; Mayr & Bryck, 2005; Schneider & Logan, 2006; Trach et al., 2021). The term "switch cost" here refers to elevated response times (RTs) or error rates when an individual must switch tasks or responses on consecutive trials, relative to repeating a task or response (for a review, see Monsell, 2003). Switch costs can occur at different levels of abstraction (i.e., task switch costs versus response switch costs; Korb et al., 2017), and are reduced but not eliminated with practice (Berryhill & Hughes, 2009; Strobach et al., 2012). While there has been much debate on what exactly drives switch costs and what makes them so robust, they are generally assumed to reflect the time needed to "reset" one's cognitive state, offloading a previous task or response and preparing for the next task or response, or adjusting to an outcome that deviated from one's expectation (Kahn et al., 2018; Monsell, 2003; Strobach et al., 2012).

To illustrate, consider that we store slightly different motor programs for Macs versus PCs. Now take the goal of copying some text: once you've specified which type of computer you're using (the upper level of a putative action hierarchy), you then select either 'ctrl-c' or 'command-c' (the lower level). Switching from a PC to a Mac should thus infer a "cost" relative to switching between two PCs; a previously active "PC node" on the putative action hierarchy must be departed for the correct "Mac node."

In the current work, we closely analyzed action switch costs to determine if participants form and navigate hierarchical representations in real time as they select actions from a visuomotor map. Participants performed a simple visuomotor learning task where they integrated trial-by-trial feedback to determine the correct button press response to eight distinct visual stimuli (**Figure 1**). Stimuli were presented in a randomized order, thus removing any temporal structure from the task. The stimuli varied along three visual feature dimensions – shape, color, and pattern – with two potential values at each level. Crucially, each feature
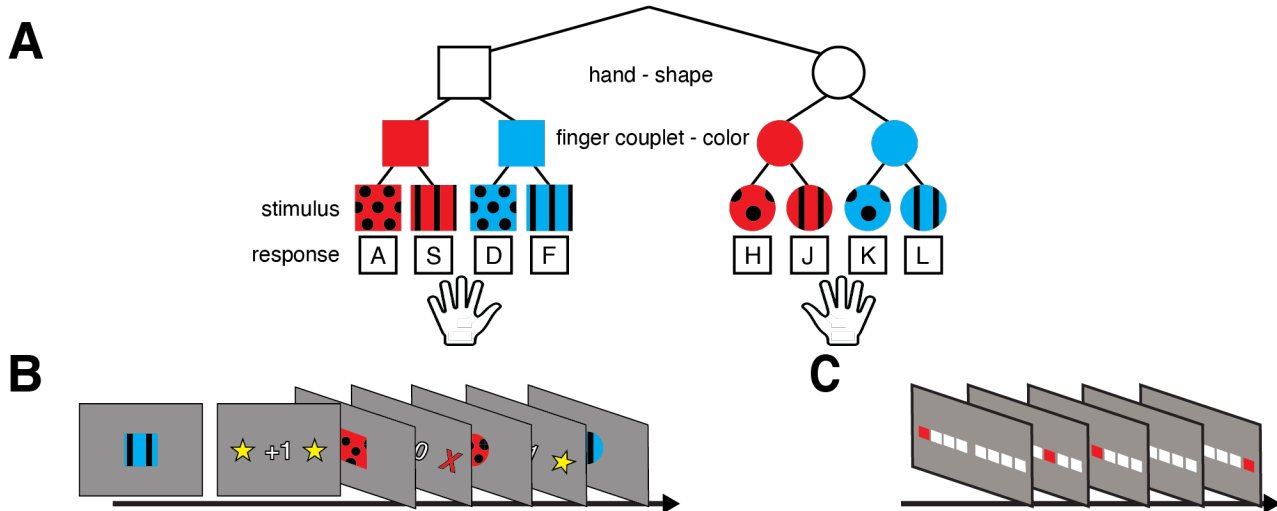
Figure 1. Task figure. (A) Schematic of stimulus-response rule structure, with shape at the top level (hand), color at the middle (finger couplets), and pattern at the bottom (fingers within couplets). Assignment of specific feature dimension to levels of the motor hierarchy was counterbalanced across participants. (B) Sequence of example trials for the learning task. Participants saw shape stimuli and then received trial-by-trial error feedback. Sequences of stimuli were randomized. (C) Sequence of example trials for the baseline motor task.

dimension was assigned to one distinct level of an intuitive, spatially organized motor hierarchy (**Figure 1A**; top: hand; middle: finger "couplets"; bottom: unique finger; Collins & Frank, 2016), such that there was a discoverable hierarchical rule structure embedded in the task. We examined interactions between finger, hand, and couplet switch costs to determine if participants learned to represent the hierarchical structure of the stimulus-response associations, and, if so, how they might be navigating this representation during action selection.

Models Considered

We first specified, *a priori*, how an agent might represent and navigate this visuomotor mapping. We considered three general models that yield dissociable predictions about RT switch costs. We were particularly interested in how hand switch costs interacted with couplet switch costs. We will refer to trial types of interest as "hierarchy-level(switch/repeat)." For example, hand(repeat)/ couplet(switch) refers to a trial where participants responded with the same hand but a different couplet as the trial before.

Flat Model. As our task is fundamentally a simple stimulus-response learning paradigm and participants need not acquire any hierarchical rules to perform it well, it is possible that participants extract a non-hierarchical cognitive representation of the mapping. That is, participants could learn the eight correct stimulus-response associations without representing any relationships between the feature dimensions and levels of the motor hierarchy. In this "flat" model, switches from one stimulus to any other stimulus are equivalent, and, after correcting for baseline motor switch costs (**Figure 2C,** described more below), switch costs should not be modulated by any particular stimulus features

(**Figure 2**). In other words, while we expect responses to repeated stimuli to be facilitated (Bertelson, 1963), all trials where some form of stimulus change occurs should produce comparable RT switch costs. For instance, trials where the correct hand or finger couplet switches from the previous trial –hand(repeat)/couplet(switch), hand(switch)/couplet(repeat), and hand(switch)/couplet (switch) trials – should not show significantly different RTs.

Feature Switch Model. Another possibility is that switch costs are modulated by the number of feature dimensions ("task sets") that change across successive trials. The fundamental prediction of this model is that RTs will be higher on trials where more feature dimensions of the stimulus change versus those where fewer dimensions change, perhaps reflecting a type of attentional switch cost or so-called "task set reconfiguration" (Monsell, 2003). Due to our design, various combinations of hand and finger couplet switches (or repeats) can be dissociated from the number of features that change across trials. For example, on trials where there is a hand(switch)/couplet(repeat), one or two stimulus features could have changed, whereas hand(switch)/couplet(switch) trials can be accompanied by two or three feature changes, but never by one. The precise switch cost predictions of the Feature Switch model are detailed in **Figure 2.** We note that while somewhat distinct, this model makes the same switch cost predictions as an alternative model that considers the number of switches at the *motor* level (e.g., hand(switch)/couplet(repeat) should have a similar RT as hand(repeat)/couplet(switch)). We necessarily combine these models here.

Tree Model. Alternatively, participants could form an action representation that is isomorphic to the hierarchical rule
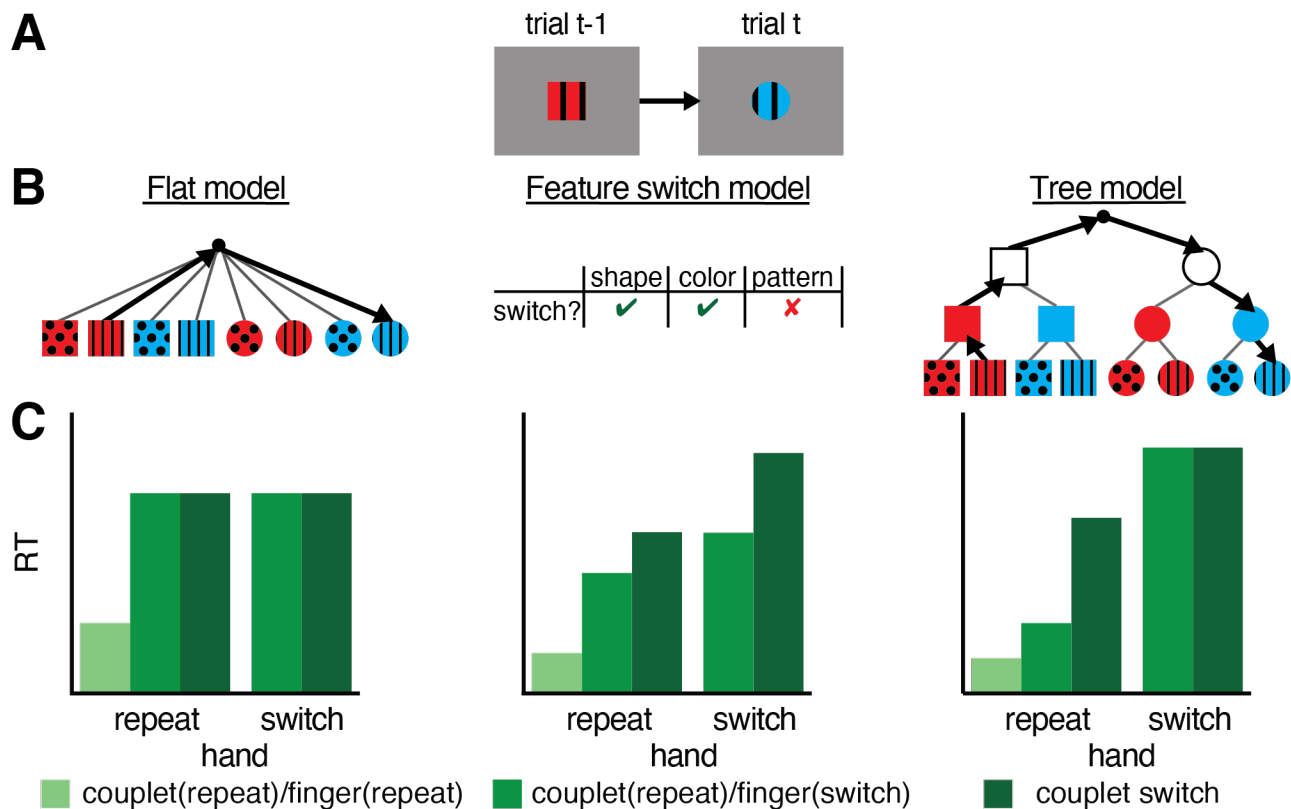
Figure 2. (A) Example sequence of two trials. (B) Schematic of the Flat, Feature switch, and Tree models' predictions for response selection on the trial sequence depicted in panel A. (C) Predictions of the pattern of switch costs for hand, couplet, and finger switches (and repeats) under each model.

structure embedded in the task. In this model, on each trial, the participant must "navigate" from the most recent response "node" through the hierarchical representation to arrive at the response node appropriate for the current trial (**Figure 2**). This model predicts that switch costs are modulated by the "distance" between the previous response and the next response as determined by the hierarchical map. This distance metric can be defined as the number of edges on the graph of a tree that one traverses between successive responses. For example, a hand(repeat)/couplet(repeat) trial can require traversal of either 0 or 2 edges, depending on whether the exact stimulus was repeated or not. Further, hand(repeat)/couplet(switch) trials are always 4 edges apart, and any trial where you switch hands requires traversing 6 edges. Crucially, in contrast to the feature switch model, the tree model predicts no difference in RT between couplet repeats versus couplet switches when accompanied by a hand switch.

Methods

Participants

41 participants (N = 13 female; mean age =32.4; range =21-57) were recruited on the online platform Prolific to participate in the study. Participants were screened for normal or corrected-to-normal vision and fluency in English to ensure that instructions were understood. They were compensated $10/hr for their time and were consented via an online form before proceeding to the task. To limit our analysis to the phenomena of interest (i.e., learned representations), we excluded participants who did not show reliable evidence of learning (i.e., got fewer than 25% of trials correct on at least 4 of the 8 stimulus-response associations), leaving us with a functional sample of 26 participants (N = 6 female; mean age = 32.9; range =22-56).

Task Design

The experimental session lasted approximately 30 min and consisted of a simple baseline motor task (4 min), instructions and practice with the learning task (2 min), and the learning task (24 min).

Motor Baseline Task. The purpose of the motor baseline task was to establish intrinsic finger switch costs (i.e., pre-learning visuomotor switch costs) for each participant. On each trial, participants saw eight empty white squares in a row on the screen (**Figure 1C**). One of the squares was then filled red while the others stayed white. Participants pressed the key that corresponded to the location of the cued red square (left hand: A,S,D,F; right hand: H,J,K,L). Once the participant
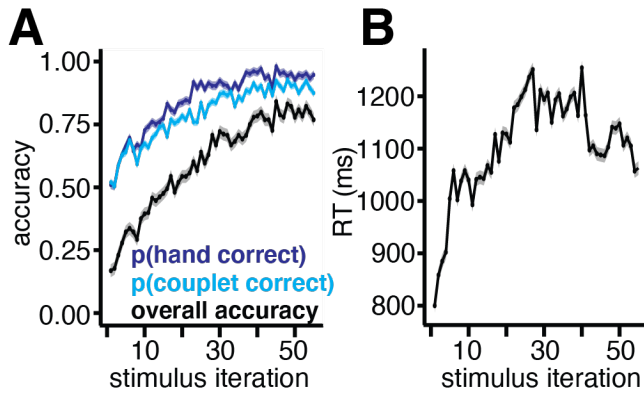
Figure 3. (A) Response accuracy (i.e., probability correct) over iterations of the stimuli, averaged across the 8 stimuli. (B) Average RT over iterations of the stimuli. Error shading is ±1 SEM.

made the correct response, all squares turned white for 100 ms before the next target was cued. Each target remained red until the correct response was made. All possible transitions between fingers (e.g., right hand pinky to left hand index), including finger repeats (e.g., right hand index to right hand index) were experienced a minimum of 4 times in the trial sequence in order to establish a stable baseline RT for each transition. Participants completed 433 trials of the motor baseline task before proceeding to a further instruction phase and practice with the learning task. The first five trials and trials where participants did not make the correct response on their first try were excluded for analysis, however the subsequent trial was included as participants had to make a correct response before proceeding.

Learning Task. After completing the motor task, participants were informed that the next task would involve using feedback to learn the correct keypress responses to eight unique visual stimuli (**Figure 1B**). On each trial,

participants had 2.5s to press a key in response to a visual stimulus. After they made their response, they received binary feedback concerning whether their response was correct (1s). Their objective was to learn the correct response for each stimulus. They were informed that each stimulus was associated with exactly one unique correct response.

Before proceeding to the main task, participants executed a short practice block using three emojis as visual stimuli and the H, J, and K keys. They were then reminded to position their left hand on the A, S, D, and F keys and their right hand on the H, J, K, and L keys and shown a preview of the stimuli that were going to be used in the main task before the task began. During the learning task, they responded to each of the eight stimuli over 56 iterations per stimulus (448 trials total). We chose 56 iterations of each association as we expected this to be sufficient exposure for participant learning, while also keeping task duration short enough to be reliably run online. The trial sequence was randomized.

The stimuli for the learning task consisted of eight images that varied on three dimensions: shape, color, and pattern. For each dimension (e.g., shape), there were two possible values (e.g., circle, square). The unique combinations of the two values for each of the three dimensions thus yielded 8 distinct stimuli (e.g., blue striped circle, red dotted square, etc.) that were each deterministically associated with one response. To embed structure in the task, we assigned each dimension to a level of an intuitive motor hierarchy (hand > finger-couplet > finger). For example, if shape was assigned to the highest level of the hierarchy – hand – then the shape of the stimulus determined which hand the correct response was on (**Figure 1**). The assignment of which specific dimension corresponded to which level of the motor hierarchy was counterbalanced across participants. Importantly, participants were never explicitly told about the rule structure and, further, did not need to learn the structure to perform the task effectively.
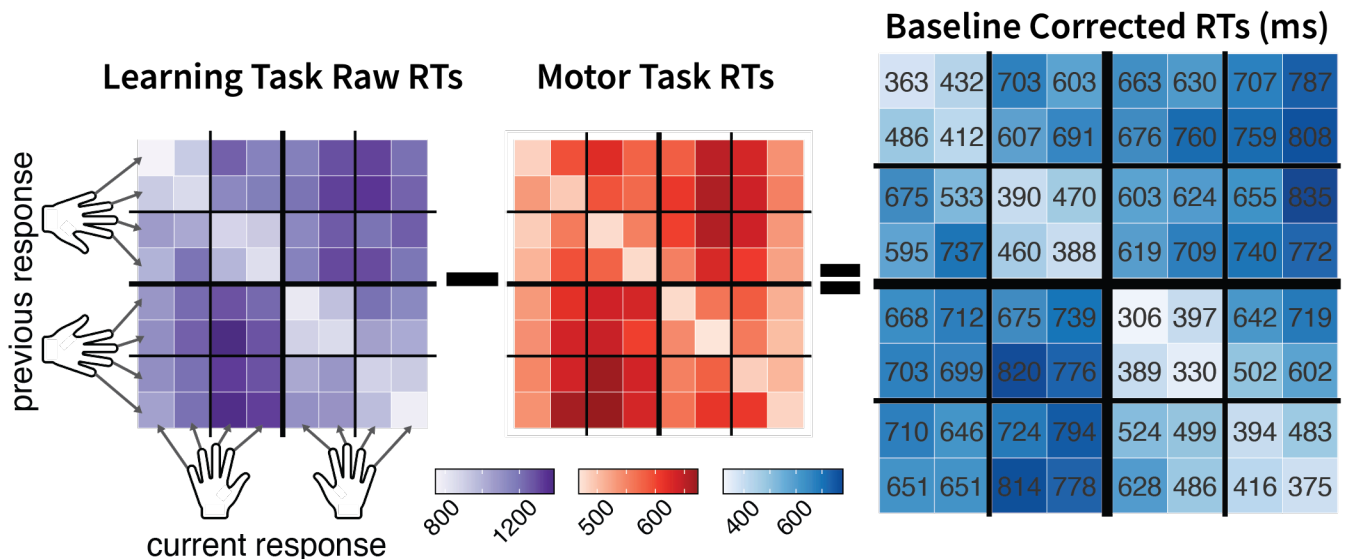

Figure 4. Heatmaps for pairwise finger transitions in learning task (purple), motor baseline task (red), and the baseline corrected RTs (blue) following subtraction. Legends are in ms for each heatmap.

Results

Learning Curves. Participants showed significant learning over the course of the task (**Figure 3A**; chance accuracy = 12.5%; mean accuracy = 60.0%, SD = 15.6%; $t(25)$ = 15.55, $p < .001$). Moreover, participants were more likely to respond with the correct hand (hand accuracy: M = 84.6%, SD = 8.8) as compared to the correct couplet (couplet accuracy: M = 79.5%, SD = 10.2; $t(25)$ = 12.91, $p < .001$). This latter result provides initial evidence that participants may have represented the task hierarchically, in that they better acquired the top level rule versus the lower level rules (Theves et al., 2021).

Departing from conventional RT learning curves, the shape of the RT curve showed that RTs slowed as participants begin to acquire the stimulus-response associations, then plateaued, and then got faster once more practice was attained (**Figure 3B**). We speculate that this "inverted" RT learning function may reflect the additional cognitive effort required for participants to acquire and navigate a hierarchical representation.

Motor Baseline Correction. Our primary analyses focused on RT switch costs on correct trials. However, to get reliable measures of learning-driven switch cost effects, we first removed intrinsic costs associated with switching between different fingers/hands in response to visual cues. To accomplish this, we used the RTs from the motor baseline task. As expected, participants exhibited finger switch costs (finger repeat versus finger switch: $t(25)$ = 8.43, $p < .001$) and hand switch costs (hand repeat versus hand switch: $t(25)$ = 12.41, $p < .001$) in the motor baseline task (Popp et al., 2020). To correct the RTs in the learning task, we calculated the average RT for each of the 64 pairwise transition between fingers in the motor baseline task (i.e., 8^2 unique transitions), and then subtracted the average RT for each transition from RTs in the learning task that corresponded to that same transition (**Figure 4**). This critical baseline correction step helps ensure that differences in switch costs are driven by the learned rule structure and not confounded by intrinsic motor biases.

RT Switch Costs. We used corrected RTs from the learning task as our dependent measure in our primary analyses. We excluded trials where participants did not make the correct response and categorized trials as hand and couplet switches versus repeats based on the correct response the participant made on that trial, and the response they made on the previous trial. The number of trials of each type included in analyses did not differ significantly across participants (Fs(1,25) < 2.46, ps > .13).

We entered the corrected RTs from the learning task (for correct trials only) into a repeated measures ANOVA, with couplet trial type (repeat versus switch) and hand trial type (repeat versus switch) as factors. We found that there was a significant main effect of couplet trial type ($F_{1,25}$ = 51.6, $p <$ .001, η_p^2 = .67), hand trial type ($F_{1,25}$ = 77.24, $p < .001$, η_p^2 = .76), and a significant interaction ($F_{1,25}$ = 76.87, $p < .001$, η_p^2
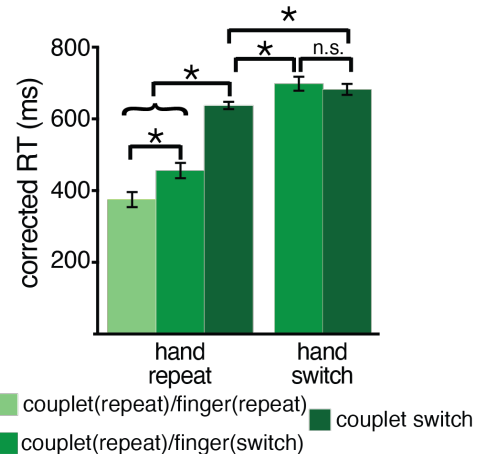

Figure 5. Average corrected RT by finger, couplet, and hand switches and repeats. Error bars are ±1 SEM.

= .76). Moreover, a switch cost was also observed at the lowest level of the hierarchy (i.e., couplet(repeat)/finger(switch) trials versus couplet(repeat)/finger(repeat) trials; $t(25)$ = 3.25, p = .003). In contrast, we did not find differences in accuracy across trial types (hand trial type: $F_{1,25}$ = 1,71, p = .20; couplet trial type: $F_{1,25}$ = 0.07, p = .79; interaction: $F_{1,25}$ = 3.91, p = .06), ruling out the confound of a speed-accuracy tradeoff.

To better relate these results to the various candidate models (**Figure 2**), we followed this analysis with four planned, post-hoc paired t-tests (with Bonferroni correction, 0.05/4: α = .0125). The results are depicted in **Figure 5**. We found evidence of a couplet switch cost within hand ($t(25)$ = 10.89, $p < .001$). In addition, we found that hand(switch)/couplet(repeat) trials were significantly slower than hand(repeat)/couplet(switch) trials ($t(25)$ = 3.10, p = .005) and that hand(repeat)/couplet(switch) trials were slower than hand(switch)/couplet(switch) trials ($t(25)$ = 3.32, p = .003). Finally, as predicted by the tree model, RT switch costs corresponding to hand(switch)/couplet(repeat) versus hand(switch)/couplet(switch) trials were not significantly different ($t(25)$ = 0.85, p = .405). Taken together, the observed switch costs neatly align with the predictions of the tree model.

We note that we also ran this same analysis on correct trials only on the second half of the learning task (i.e., 28 iterations of each stimulus), when participants have substantial practice with the task. The key results aligned with those presented above (Fs > 39, ps < .001), though the comparison between hand repeat-couplet switch and hand switch-couplet switch trials was marginally significant in this case after Bonferroni correction ($t(25)$ = 2.12, p = .044). Overall, the orderly pattern of results strongly supports the tree model, where RT switch costs are modulated by the number of edges separating successive responses on a hierarchical graph (**Figure 2**).

Model Comparisons. Finally, for completeness, we compared linear mixed effects models for each of the three

models we considered. For these regressions, we fit baseline-corrected RTs (on correct trials) with random intercepts for subjects. For the flat model, the predictor simply reflected whether the response switched or was repeated on any given trial. For the feature switch model, the number of feature switches between responses was the predictor. Finally, for the tree model, we set the number of graph edges that needed to be traversed between the previous and current response nodes as the predictor. BIC metrics strongly indicated that the tree model (BIC = 97954) was a better fit for behavior versus both the feature (BIC = 100751; ΔBIC = -2798) and flat (BIC = 98294; ΔBIC = -340) models, providing further empirical support for the tree model.

Discussion

It is well known that humans use hierarchical cognitive representations to organize behavior. Here, we asked how participants navigate these hierarchical representations in real time to select actions from an acquired visuomotor mapping. We provide evidence that participants can represent learned visuomotor associations hierarchically, and that their response times (RTs) between successive actions are modulated by the distance between those actions on a hierarchical tree representation. Speculatively, this latter result may indicate that participants navigate internal hierarchical representations of action maps in a manner that echoes the navigation of actual physical space (Behrens et al., 2018).

While we provide initial evidence supporting this model of cognitive map navigation during action preparation, we note that our results do not perfectly align with a tree model. For example, the overall RT switch cost of hand switches was slightly smaller than a perfect tree model predicts (see **Figures 2** and **5**). While previous work shows that behavioral switch costs are generally larger at superordinate levels of a task hierarchy (Collins & Frank, 2013), a phenomenon that we also observed here, our tree model appeared to overestimate the size of this effect. Discrepancies between a "perfect" tree model (i.e., equal edge lengths) and our results could be, in part, attributable to other variables like feature change costs, interactions between switching and repeating at different levels of a hierarchy (Mayr & Bryck, 2005; Trach et al., 2021), or other known determinants of RTs, such as the urgency to decide (Cisek et al., 2009).

Another possibility is that aspects of our task design might have influenced the alignment of our results with the tree model. For one, ideal analyses would include only RTs on two consecutively correct trials. We were not able to conduct this analysis due to the short duration of our task. Because of this and the high exclusion rate of online studies, we plan to conduct an extended version of this task to replicate these findings. Additionally, it is possible that our correction for motor biases did not perfectly remove effects driven by motor constraints. We also intend to test a version of the task where a hierarchical rule structure cannot be extracted, and compare switch costs in the hierarchical versus non-hierarchical tasks as a more rigorous control for motor bias.

Finally, in future work, we plan to use more controlled RT paradigms to elicit responses at various imposed latencies, allowing us to read out the navigation of a hierarchical representation in "real time." In this case, we can also use error rates to address our central question, potentially providing more evidence for navigation through a hierarchical representation.

Our analysis also rests on strong assumptions about lingering activation of previous response "nodes" as the starting point during putative navigation of the hierarchy (Schneider & Logan, 2007). While the existence of switch costs necessarily implies that previous responses affect current behavior, it is also the case that increasing the interval between two responses can diminish switch costs (Altmann, 2004). In future work we can manipulate these intervals to measure the decay of these effects.

Our results are related to (and partly inspired by) previous work on hierarchical rule structure learning in the reinforcement learning (RL) literature (Collins, 2017; Collins & Frank, 2013, 2016; Eckstein et al., 2019; Eckstein & Collins, 2021). Like our paradigm here, in this previous work participants learned contextual action rules that could be structured hierarchically (although limited to two-level hierarchies). Subjects' behavior in these studies was successfully modeled as a form of hierarchical reinforcement learning. While we have not used RL models to examine the learning process in the current study, we plan to explore this option in future studies, preferably with more learning trials. Our work builds on these findings, showing how a multi-level hierarchy can be abstracted over actions even when future planning is not required. Moreover, our study provides insight into how hierarchical representations are navigated by individuals as they prepare actions; by having more than two levels in our hierarchy, we were able to make graded predictions about RT switch costs.

Here we demonstrated how hierarchical representations can be abstracted over simple actions, similar to how they are abstracted over time (Kahn et al., 2018). We also demonstrate how these representations may be navigated in real time during action selection. Our study also highlights how simple behavioral effects, such as RT switch costs, can help elucidate the complex structure of mental representations. Looking ahead, our findings can inform future work on mental representations of naturalistic visuomotor mappings involved in motor skills that require the selection of actions among multiple effectors, such as typing, playing musical instruments, and dance.

Acknowledgments

References

Altmann, E. M. (2004). Advance preparation in task switching: What work is being done? *Psychological*

Science, *15*(9), 616–622. https://doi.org/10.1111/j.0956-7976.2004.00729.x

Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, *66*(2), 315–326. https://doi.org/10.1016/j.neuron.2010.03.025

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490–509. https://doi.org/10.1016/j.neuron.2018.10.002

Berryhill, M. E., & Hughes, H. C. (2009). On the minimization of task switch costs following long-term training. *Attention, Perception, & Psychophysics*, *71*(3), 503–514. https://doi.org/10.3758/APP.71.3.503

Bertelson, P. (1963). S-R relationships and reaction times to new versus repeated signals in a serial task. *Journal of Experimental Psychology*, *65*(5), 478–484. https://doi.org/10.1037/h0047742

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, *29*(37), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

Collins, A. G. E. (2017). The cost of structure learning. *Journal of Cognitive Neuroscience*, *29*(10), 1646–1655. https://doi.org/10.1162/jocn_a_01128

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229. https://doi.org/10.1037/a0030852.supp

Collins, A. G. E., & Frank, M. J. (2016). Motor demands constrain cognitive rule structures. *PLOS Computational Biology*, *12*(3), e1004785. https://doi.org/10.1371/journal.pcbi.1004785

Collins, A. M., & Quillian, M. R. (1969). *Retrival time from semantic memory* (pp. 191–201).

Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., & Pallier, C. (2015). The neural representation of sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron*, *88*(1), 2–19. https://doi.org/10.1016/j.neuron.2015.09.019

Eckstein, M. K., & Collins, A. G. E. (2021). How the mind creates structure: Hierarchical learning of action sequences. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *43*(43), 8.

Eckstein, M. K., Starr, A., & Bunge, S. A. (2019). How the inference of hierarchical rules unfolds over time. *Cognition*, *185*, 151–162. https://doi.org/10.1016/j.cognition.2019.01.009

Kahn, A. E., Karuza, E. A., Vettel, J. M., & Bassett, D. S. (2018). Network constraints on learnability of probabilistic motor sequences. *Nature Human Behaviour*, *2*(12), 936–947. https://doi.org/10.1038/s41562-018-0463-8

Korb, F. M., Jiang, J., King, J. A., & Egner, T. (2017). Hierarchically organized medial frontal cortex-basal ganglia loops selectively control task- and response-selection. *The Journal of Neuroscience*, *37*(33), 7893–7905. https://doi.org/10.1523/JNEUROSCI.3289-16.2017

Mark, S., Moran, R., Parr, T., Kennerley, S. W., & Behrens, T. E. J. (2020). Transferring structural knowledge across cognitive maps in humans and models. *Nature Communications*, *11*(1), 4783. https://doi.org/10.1038/s41467-020-18254-6

Mayr, U., & Bryck, R. L. (2005). Sticky rules: Integration between abstract rules and specific actions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(2), 337–350. https://doi.org/10.1037/0278-7393.31.2.337

McNamara, T. P. (1986). Mental representations of spatial relations. *Cognitive Psychology*, *18*(1), 87–121. https://doi.org/10.1016/0010-0285(86)90016-2

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. https://doi.org/10.1016/S1364-6613(03)00028-7

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*(1), 1–32. https://doi.org/10.1016/0010-0285(87)90002-8

Popp, N. J., Yokoi, A., Gribble, P. L., & Diedrichsen, J. (2020). The effect of instruction on motor skill learning. *Journal of Neurophysiology*, *124*(5), 1449–1457. https://doi.org/10.1152/jn.00271.2020

Schneider, D. W., & Logan, G. D. (2006). Hierarchical control of cognitive processes: Switching tasks in sequences. *Journal of Experimental Psychology: General*, *135*(4), 623–640. https://doi.org/10.1037/0096-3445.135.4.623

Schneider, D. W., & Logan, G. D. (2007). Retrieving information from a hierarchical plan. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(6), 1076. https://doi.org/10.1037/0278-7393.33.6.1076

Strobach, T., Liepelt, R., Schubert, T., & Kiesel, A. (2012). Task switching: Effects of practice on switch and mixing costs. *Psychological Research*, *76*(1), 74–83. https://doi.org/10.1007/s00426-011-0323-x

Theves, S., Neville, D. A., Fernández, G., & Doeller, C. F. (2021). Learning and representation of hierarchical concepts in hippocampus and prefrontal cortex. *Journal of Neuroscience*, *41*(36), 7675–7686. https://doi.org/10.1523/JNEUROSCI.0657-21.2021