

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Structured dynamics of hierarchical action selection

Permalink

<https://escholarship.org/uc/item/00q1v22d>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Trach, Juliana E.
McDougle, Samuel David

Publication Date

2023

Peer reviewed

Structured dynamics of hierarchical action selection

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

Extensive work has examined how individuals structure information in cognitive representations. However, the *dynamics* of how these structured representations are implemented in real time and how cognitive and motor processing interact during action selection have received relatively less attention. Here, we use computational modeling to closely examine the dynamics of hierarchical action selection at the scope of individual decisions. We had participants learn eight stimulus-action associations with latent hierarchical structure. They then engaged in a forced response task where we manipulated the amount of time participants had to prepare each response. We find evidence that hierarchical action selection is a top-down, serialized process but appears to occur without bottlenecks between decision-making and movement. Overall, our results highlight a close coupling between cognitive and motor processing during hierarchical action selection.

Keywords: hierarchy, structure learning, action selection, forced response, dynamics

Introduction

Hierarchical cognitive representations are a hallmark of human cognition. There is evidence across a wide range of domains that people utilize these structured representations to facilitate behavior (Behrens et al., 2018; Collins et al., 2014; Theves et al., 2021). Further, there is evidence that even simple visuomotor routines can be represented hierarchically (Logan, 2018; Trach & McDougle, 2022). While the structure of these representations has been widely studied, there has been limited work examining how they are *implemented* in real time. Here, we use psychophysics and computational modeling to characterize the dynamics of hierarchical decision-making and action selection within the scope of individual decisions.

Previous work addressing hierarchical cognitive representations largely utilizes differences in reaction time (RT; Dykstra et al., 2022) or variation in the magnitude of RT “switch costs” to make inferences about the structure of cognitive representations (Collins, 2017). “Switch cost” in this context refers to slowed RTs on trials where participants must switch tasks or responses relative to trials where they repeat the same task or response. For example, a participant will respond slower to the color of a presented stimulus if they were asked to respond to the shape on the previous trial (e.g., Schneider & Logan, 2006). Such switch costs are thought to reflect cognitive effort involved in “offloading”

previously task-relevant information and preparing to process newly task-relevant information (Schneider & Logan, 2006; Strobach et al., 2012). In the context of hierarchical tasks, switch costs are an especially useful tool as they can occur at multiple levels of a task (e.g., ‘task’ switch costs versus ‘response’ switch costs; Korb et al., 2017; Mayr & Bryck, 2005) and increase with the level of abstraction such that switch costs are higher at superordinate levels of a task relative to subordinate levels (task switch costs > response switch costs; Collins, 2017; Collins et al., 2014). Thus, researchers can use the magnitude of switch costs to infer which task features are superordinate to others from behavior.

Take for example the goal of copying or pasting on Macs or PCs (**Figure 1A**). Each of these operating system-goal pairs necessitates a unique motor plan (e.g., command-c versus command-v or ctrl-c versus ctrl-v). While all four motor plans are highly similar, you might have the intuition that switching between operating systems (ctrl-c → command-c) is more difficult than switching between copying and pasting within one operating system (ctrl-c → ctrl-v). Thus, based on this switch cost, we can infer that we represent this type of task hierarchically, where operating system is superordinate to goal which in turn governs response output. While this approach is useful for studying the *structure* of representations to a certain extent, it does not reveal the dynamics of how these structures are implemented in real time, particularly within the scope of the execution of individual actions.

Such hierarchical decision-making is canonically conceptualized as a top-down, serialized process (e.g., Sigman & Dehaene, 2005). That is, decisions at superordinate task levels must be resolved before proceeding down through the hierarchical cognitive representation to subordinate levels and finally to the motor response level. In the context of our example, this view suggests that you must finish determining the operating system before you begin to consider the goal level and finally the actual movement. Some behavioral and neural work has challenged this ‘bottleneck’ view, providing evidence of parallel processing of information at different task levels (Cellier et al., 2022; Ranti et al., 2015). In one recent example, Cellier et al. (2022) used EEG to reveal persistent and temporally overlapping representations of information across hierarchical levels of a cognitive control task. Parallelization in this context is thought to be possible due to a hypothesized hierarchical gradient of representational abstraction in the prefrontal cortex (Badre &

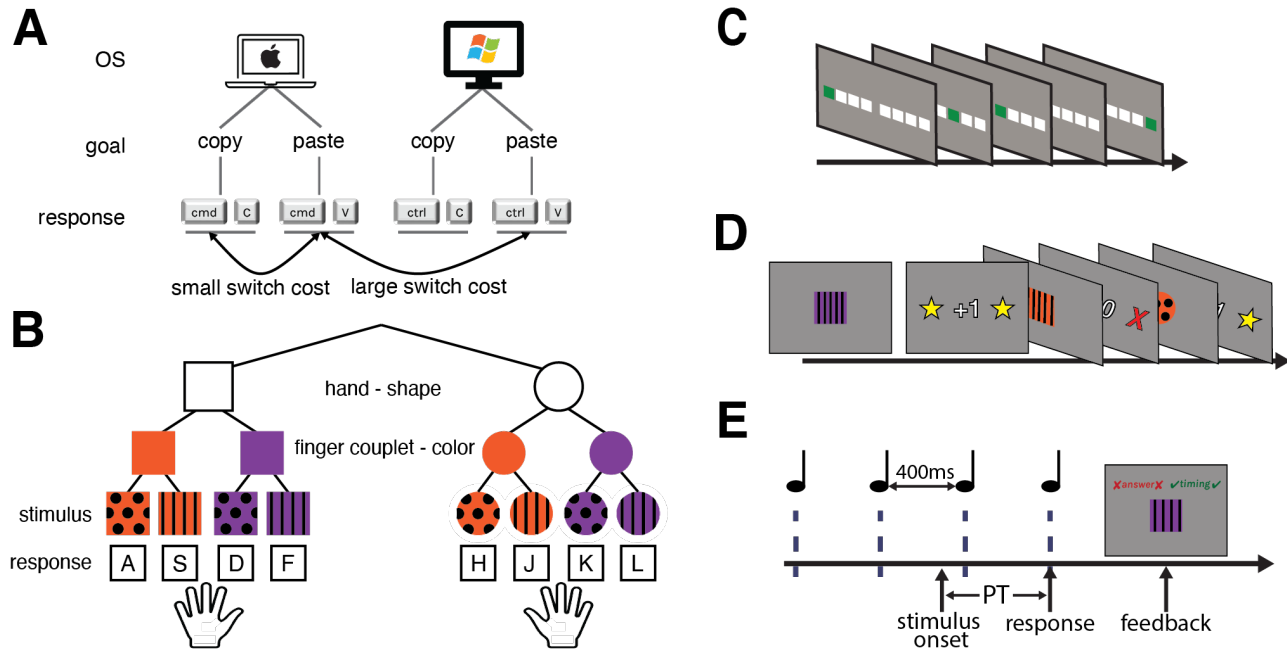


Figure 1. A) Diagram of hierarchical organization of operating system-goals where operating system is superordinate to goal. B) Example structure of stimulus-action associations in the learning task. Hand is associated with shape (top level), then couplet with color (middle level), and finally pattern with finger (bottom level). C) Illustration of example trials from the motor baseline task. D) Illustration of example trials of learning task. E) Schematic of individual forced response trial. Participant aligns response with fourth beep. Stimulus onset can occur at any point during the trial interval. Preparation time (PT) is the time between stimulus onset and response. Timing and accuracy feedback displayed after response.

Desrochers, 2019). Importantly, such work does not consider the role of motor processing in their analyses and it remains an open question whether parallel hierarchical processing occurs in the motor domain at the level of individual action selection (e.g., choosing one effector over another).

We address the limitations of typical hierarchical decision-making tasks and expand on existing literature by using a forced response paradigm (Figure 1E) in conjunction with a stimulus-action learning paradigm (Trach & McDougle, 2022). For this paradigm, participants first learned an eight-to-eight stimulus-action mapping (Figure 1B) from visual stimuli to finger responses, spanning the two hands (no thumbs). The mapping followed an intuitive hierarchy where different stimulus features (color, shape, pattern) were associated with different spatial/motor features of the task (hands, pairs of fingers, left/right positions within those pairs). After training to asymptote, participants entered a forced response phase (Hardwick et al., 2019; McDougle & Taylor, 2019). Here, they had to synchronize their responses to a precisely-timed auditory cue played at the end of a metronomic countdown. We manipulated the amount of time participants had to prepare each response by varying when the visual stimulus was displayed during the countdown. For example, if the stimulus was displayed early in the countdown, participants would have a longer time to prepare their response before having to respond at the end of the trial. On other trials, the visual stimulus might be displayed toward the end of countdown, right before a response had to be made; here, participants essentially had to guess to respond at the

correct time. Using this paradigm, we measured how hierarchical action selection unfolds over time by examining the types of errors people made at different preparation times. In this way, we can characterize decision-making and motor planning dynamics continuously over time, rather than being limited to responses at the end of the decision-making process as in traditional reaction time (RT) measures.

We also employed computational modeling to simulate different behavioral strategies in our hierarchical task. Overall, we found evidence that during hierarchical decision-making and action selection, levels of the hierarchy are processed serially, but likely without bottlenecks – that is, multiple levels can be resolved at once as people select an action but are resolved from “top to bottom.” This finding thus unites serial and parallel models of hierarchical processing, arguing for a mixture of both. Further, our results reveal a strikingly tight coupling between cognitive and motor processing: instead of motor output simply being the end result of a completed cognitive or perceptual process, we find that different subsets of motor actions are dynamically potentiated as hierarchically represented sensory input is interpreted. This discovery suggests that selecting actions from a structured sensorimotor mapping involves rapid, orderly hierarchical processing, and involves a continuous flow of information from cognition to action.

Methods

Participants

40 participants (N = 22 female; mean age =19.5; range =18-22) were recruited through the student subject pool at Yale University. They received 1 course credit/hour for their participation. We excluded 4 participants that did not show reliable evidence of learning (i.e., got fewer than 25% of trials correct on at least 4 of the 8 stimulus-action associations). We planned to exclude participants that did not respond to >75% of trials in either the learning or forced response phase and participants that were unable to respond on time during the forced response task at least 25% of the time, however no subjects met those criteria. Our final sample included 36 participants (N = 19 female; mean age =19.6; range =18-22).

Task Design

Experimental sessions were 1 hour in duration and consisted of four phases: 1) a motor baseline task (4 min), 2) forced response task training (3 min), 3) the stimulus-action learning task (27 min), and 4) forced response task (20 min).

Motor Baseline Task. Participants first engaged in a cued-response task to measure intrinsic finger-to-finger switch costs prior to the learning task (**Figure 1C**). During this task, eight squares were displayed on the screen that were spatially aligned with the participant's hands on the keyboard (left hand: A,S,D,F; right hand: H,J,K,L). On each trial, one of the squares would change from white to green and participants were instructed to press the key that was aligned with the green square as quickly and accurately as possible. The trial ended once the participant responded with the correct finger. After the correct response, all the squares turned back to white for 100 ms before the next trial began. Participants executed 489 trials of this task and the trial sequence included 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). Trials where participants responded incorrectly on their first attempt and the first five trials of the task were excluded from analysis.

Next, we used this task to introduce the forced response procedure to participants. On each trial, participants heard four beeps each spaced 400 ms apart. Participants were instructed to respond in synchrony with the last of the four beeps, regardless of when the stimulus they were responding to appeared on the screen. During this practice phase, the green square always appeared with the first beep in the sequence so participants had 1.2 s to prepare their response before responding in time with the last of the four beeps. They received feedback on whether their response was correct and whether they responded at the correct time during the practice. A correct response was signaled if participants responded with +/- 100 ms of the fourth tone.

Learning Task. During the learning task, participants were instructed to use trial-by-trial reward feedback to figure out the correct response to a visual stimulus (**Figure 1D**). Like the motor task, participants used eight fingers to respond (left hand: A,S,D,F; right hand: H,J,K,L). Each button was uniquely and deterministically associated with one visual

stimulus. On each trial, participants would see the stimulus (3.5s), make a response on the keyboard, and then receive binary feedback as to whether their response was correct or not (750ms). During the main task, they saw 56 iterations of each of the eight stimuli (448 total learning trials). The goal was to use the feedback to figure out which action was associated with each of the eight stimuli.

Participants were given instructions and executed a short practice block with three emojis as the visual stimuli to ensure they understood the instructions before beginning the main learning task. They were also shown the eight stimuli for the main learning task in a random order before beginning.

To embed hierarchical structure in the task, we varied the visual stimuli on three features – shape, color, and pattern – and assigned each feature to a level of an intuitive motor hierarchy (hand > finger-couplet > finger, **Figure 1B**; Trach & McDougle, 2022). Each feature had two possible values (e.g., color: purple or orange) in order to create the eight stimuli for the learning task. For example, if color was associated with the top level of the motor hierarchy (i.e., hand), then the correct actions for all of the purple stimuli would be on one hand and the correct actions for all of the orange stimuli would be on the other hand. Participants were not informed about this structure. We counterbalanced the assignment of stimulus feature to motor-hierarchy level across participants to ensure that no effects were driven by differences in feature salience. Importantly, all features were relevant to determining the correct actions (i.e., the correct response could not be determined without attending to *all* of the three features), however participants did not need to represent the task hierarchically to perform it successfully. We only include consecutively correct trials (i.e., trials where the participant made the correct response on that trial *and* the previous trial) in our analyses. These trials generally occur later in the learning task once the mapping is well-learned.

Forced Response Task. After participants learned the eight stimulus-action pairs during the learning task, we tested response retrieval dynamics using a forced response paradigm. Before beginning, participants were reminded that the correct actions for each stimulus that they had just learned would remain the same in the final phase.

On each trial, participants heard four, high-pitched beeps spaced 400 ms apart (**Figure 1E**). They were instructed to time their response to coincide with the final of the four beeps (i.e., 1.2s after the first beep), regardless of when they were prepared to respond. In order to test our primary questions, we manipulated the amount of preparation time (PT; i.e., the amount of time between when the stimulus was displayed and the fourth beep) that participants had on each trial by varying when the stimulus was displayed during the trial interval. For example, if the visual stimulus appeared with the initial beep, the participant would have 1.2s of PT on that trial. On other trials, the stimulus would be shown later during the sequence of beeps, leaving participants with less PT on those trials. We varied PT from 100ms-1.2s by uniformly sampling in this

range across trials. Thus, on some trials participants would have sufficient time to plan and execute correct actions, whereas they might be forced to rush (or guess) on trials where PT was very short. Our primary analyses involved examining patterns of errors at different amounts of PT in order to characterize the action preparation process over the whole response interval.

Participants received binary feedback on whether they had made the correct action and whether they had responded at the right time on every trial (**Figure 1E**). They executed 770 forced response trials with short breaks every approximately 96 trials. Participants were encouraged to primarily attend to the timing of their responses and try to respond to the best of their ability on every trial. Trials with responses before the stimulus appeared or after 1.3s were excluded from analysis.

Results

Learning task results. Participants performed the task well and showed learning across the session ($t(35) = 522.14, p < .001$; **Figure 2A**). Participants were more accurate at responding with the correct hand as compared to the correct couplet overall, providing initial evidence that they were representing the hierarchical structure of the task ($p(\text{Correct hand}) = .88$; $p(\text{Correct couplet}) = .82$; $t(35) = 5.86, p < .001$). However, by the end of the learning phase (i.e., in the last 5 iterations of each stimulus), hand- and couplet-level learning were not statistically different ($p(\text{Correct hand}) = .98$; $p(\text{Correct couplet}) = .96$; $t(35) = 1.98, p = .056$).

Before examining reaction times (RTs) to test whether participants were representing the task hierarchically, we corrected all RTs in the learning task using RT switch costs measured during the baseline motor task. This process controlled for potential baseline biomechanical effects of switching between specific fingers. For example, if response transitions across hands take longer than response transitions within hand due to intrinsic biomechanical constraints, this would inflate the effects we observe in the task. To correct for this, we subtracted the mean RT for each finger-to-finger transition in the baseline task from the RTs of the same finger-to-finger transition in the learning task (Trach & McDougle, 2022). This correction accounts for variation in

RT driven by intrinsic motor constraints and isolates RT variation related to learning the task's structure.

We tested if participants were representing the task hierarchically in two ways: First, we compared switch costs at different levels of the task. Previous work indicates that RT switch costs should be larger at higher versus lower levels of a hierarchical task (Collins et al., 2014; Collins & Frank, 2016). We found that corrected RT switch costs on correct trials did in fact vary in size across levels of the task (repeated measures ANOVA RT switch cost \times level(finger, couplet, hand): $F(2,70) = 25.28, p < .001$). Specifically, switch costs were larger at superordinate levels of the task as compared to subordinate levels (hand vs couplet: $t(35) = 3.23, p = .0027$; hand vs finger: $t(35) = 7.08, p < .001$; couplet vs finger: $t(35) = 3.91, p < .001$), providing further evidence that participants were representing the task hierarchically.

Second, we entered corrected RTs on correct trials into three linear mixed effects models to further examine whether participants were representing the task hierarchically, using: 1) a Hierarchical model, 2) a Stimulus-based model, and 3) a Flat model. First, the hierarchical model used distance within the hierarchical rule structure (i.e., the number of graph edges between the current and previous stimulus in the hierarchical rule structure; Trach & McDougle, 2022) to predict RTs. Second, the stimulus-based model predicted RTs with the number of stimulus features that switched on a given trial. Third, if participants were simply learning one-to-one associations between stimuli and actions and not representing any latent structure in the task, switching from any response to any other action should incur the same RT switch cost. Therefore, the predictor for the flat model was simply whether the participant responded with the same or a different finger on a given trial. Consistent with previous, we found that the hierarchical model was the best fit for behavior (Hierarchical: BIC = 115,451; Stimulus-based: BIC = 115,631, Flat: BIC = 115,899; **Figure 2B**; Trach & McDougle, 2022). Overall, these findings indicate that participants were representing the latent hierarchical structure of the task and that this structure influenced participant behavior.

Forced response results. Our primary analyses for the forced response phase of the task involved examining patterns of errors at different PTs. To that end, we coded each response participants made based on correctness at each level (e.g., 1/1/0 denotes that the participant responded with the correct hand and couplet but the incorrect finger, whereas 0/1/0 would indicate that the participant had responded with the correct couplet but incorrect hand and finger). For visualization purposes, we only code the finger level as correct if the participant responds with the correct finger on the correct hand, however the couplet level is coded as correct if they respond with the correct couplet on either hand. Thus, each trial can be coded as one of four possible error codes (1/1/0, 1/0/0, 0/1/0, 0/0/0) or as correct (1/1/1). We then used a sliding window of 100 ms across to calculate and visualize the probability of each type of error over the entire range of

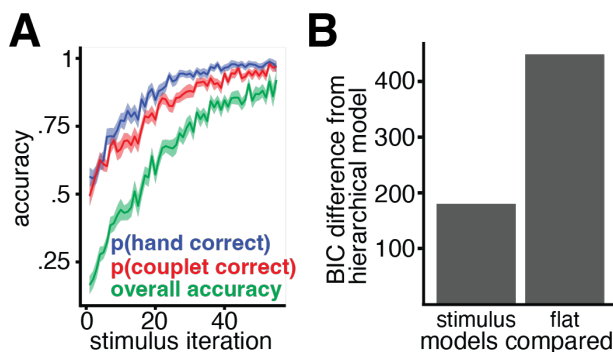


Figure 2. A) Learning curves depicting overall accuracy and the probability of responding with the correct hand or couplet. B) BIC comparison between hierarchical model and stimulus-based or flat model.

presented PTs (**Figure 3A**). At the shortest PTs, participants were at chance for each response type (chance level: .125 for 1/1/1 and 1/1/0, .25 for 1/0/0, 0/1/0, and 0/0/0) and response probabilities evolved with increasing PT. In this way, we can characterize the ongoing influence of the cognitive representation of the task on hierarchical action selection continuously across the trial window.

We designed a computational model to characterize the dynamics of processing across the trial interval and test whether processing at different hierarchical levels of the task happens serially or in parallel (**Figure 3B**). The basic model assumes that action preparation time is normally distributed, and takes some average amount of time, μ , with variance σ . We can calculate the probability of executing a specific response as a function of PT by taking the cumulative density of the RT distribution (**Figure 3B**). We modeled the probability of each response at each PT as a mixture of this action selection probability curve and some degree of random guessing (“lapses”) across all possible actions, with guessing scaled by weighting factor ρ .

To test whether processing across the hierarchical representation happens serially or in parallel, we allowed three separate μ parameters when fitting the model, with one assigned to resolving each level of the task (μ_{hand} , $\mu_{couplet}$, μ_{finger}). Here, each μ describes how long it takes to select the correct action candidates for each level of the task. Thus, if

processing starts with the superordinate task levels before percolating down to subordinate levels, we expect fitted μ_{hand} values to be lower (i.e., faster) than $\mu_{couplet}$ or μ_{finger} . However, if action selection across levels happens fully in parallel, or as the end result of cognitive decision-making processes that do not reach the motor system until completed, we would not expect this pattern of results (i.e., $\mu_{hand} = \mu_{couplet} = \mu_{finger}$).

We fit our action selection model to participants actual choices and PTs using maximum likelihood estimation. We fit 50 iterations to each subject to avoid local minima in the likelihood surface and constrained μ between .2 and 2s and σ between .01 and .6s. Importantly, while we simplify correctness at the lowest hierarchical level for visualization purposes (as described previously), the model is fit with full information about correctness at each level of the task. The simulated action probabilities based on the across-subject mean parameters are depicted in **Figure 3C**.

Fitted μ parameters provided striking evidence that participants were, in real time, dynamically and serially modulating action probabilities based on the hierarchical task levels (**Figure 3D-E**; $\mu_{hand} = .86$; $\mu_{couplet} = 1.13$; $\mu_{finger} = 1.28$; Wilcoxon signed rank test: μ_{hand} versus $\mu_{couplet}$: $z = -4.48$, $p < .001$; μ_{hand} versus μ_{finger} : $z = -5.20$, $p < .001$; $\mu_{couplet}$ versus μ_{finger} : $z = -4.15$, $p < .001$). That is, participants were fastest at determining the correct hand to respond with, followed by the correct couplet, before ultimately selecting the correct

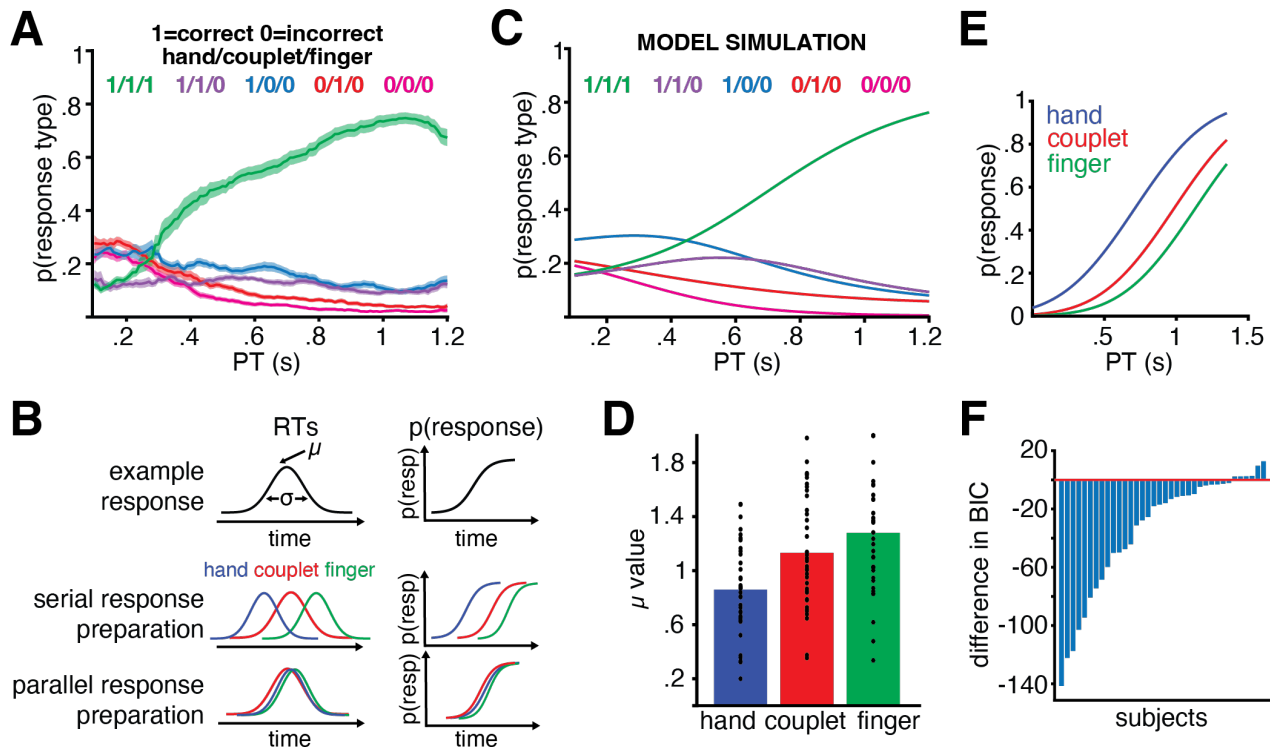


Figure 3. A) Probability of different actions/errors at different PTs over the trial interval. B) Illustration of model and competing hypotheses. Serial action preparation predicts distinct μ values for each level, whereas parallel preparation does not. C) Model simulation of probability of different responses at different PT over the trial interval. D) Mean μ values at each level of the hierarchy from model fitting. E) Response preparation functions based on model fitting results. F) Comparison of BIC for one- μ versus three- μ model by subject. Negative values indicate that three- μ model was a better fit for behavior.

finger. Importantly, we did not *constrain* the model to order these values sequentially – this pattern of results is driven by fitting the model directly to participant behavior.

For completeness, we also fit a model that only allowed for only a single μ parameter (i.e., assuming parallel processing across task levels) to compare with our main model. We found that the three- μ model was a better fit for participant behavior than the single- μ model in 30 out of the 36 participants, even after accounting for its greater number of free parameters (**Figure 3F**; summed BIC three- $\mu = 76,536$; summed BIC one- $\mu = 77,767$). Overall, our results indicate that participants sequentially resolve representational levels from high (superordinate) to low (subordinate) as they select a single action from a structured visuomotor mapping. Further, our results strongly suggest that hierarchical action selection evolves continuously over the course of a single decision, rather than emerging from a bottleneck after the completion of an earlier cognitive decision-making stage.

Discussion

Here, we utilized a novel computational model to characterize the real-time dynamics of action selection within the scope of single decisions. The latent hierarchical structure of our stimulus-action mapping allowed us to test whether actions are a result of serialized, hierarchical processing or if processing across task levels can occur simultaneously. Our results strongly suggest that action selection is hierarchical, serialized, and top-down, in contrast to some evidence for purely parallel processing in the domain of cognitive control (Cellier et al., 2022; Ranti et al., 2015). Moreover, our pattern of results illustrates a tight coupling of cognitive, perceptual, and motor processing and shows that movement selection evolves within individual trials along with hierarchical decision-making processes, rather than emerging from a bottleneck (Selen et al., 2012). Beyond novel scientific insights into the action selection process, this work builds on a productive behavioral method to approach questions about the dynamics of cognitive processing (Hardwick et al., 2019; McDougale & Taylor, 2019) and highlights the utility of computational modeling in quantifying these dynamics.

While movement is sometimes considered a final output of cognitive processing (Pashler, 1984), our findings are consistent with a tight coupling between decision-making and motor processes during hierarchical action selection – each stage of perceptual processing (i.e., shape \rightarrow color \rightarrow pattern) biased the probability of different actions being selected in real time. Thus, our results are consistent with work (in non-hierarchical settings) that highlights dynamic

flow of information from ongoing decision-making processes to the motor system (Selen et al., 2012; Spivey et al., 2005). A possible alternative to this result would be that the processing of stimulus features occurs in a top-down, serialized manner but action selection happens as a single stage after the completion of decision-making (**Figure 4**, right panel). If this were the case, we would expect a uniform distribution of error types across the trial interval and equivalent μ values at different task levels from the computational modeling. Future work using neural recording methods with high temporal and spatial resolution (such as MEG) are essential in clarifying interactions between cognitive systems in hierarchical decision-making settings.

Such neural investigation could also inform a more holistic description of the precise temporal dynamics of hierarchical action selection. For example, our results stand in contrast to recent work in the domain of cognitive control that highlight surprisingly parallel processing across representational task levels (Cellier et al., 2022; Ranti et al., 2015). It is possible that there is some degree of parallel processing that we are unable to detect with our behavioral task. The addition of neural recording could allow us to characterize the dynamics of perceptual processing beyond behavioral outputs of responses and examine neural representation across task levels. On the other hand, our model also is not constrained to the strictest serial processing strategy; it does not *require* that superordinate levels are completely resolved before subordinate levels. A strict serial model would force the agent to prune large portions of the action options at each hierarchical level of processing before moving on to the next one. In future work, we can formalize this alternative hypothesis with computational modeling to compare to our current model.

In future analyses, we also plan to extend our model to integrate contextual information about each trial to precisely capture participant behavior. Specifically, we can develop models that integrate information about the previous trial to ask if different types of action transitions influence speed-accuracy functions. One limitation of our current model is that the hand response function shows slightly above chance-level performance even at the shortest PTs. This effect in the model could emerge from differences in response function for hand-repeat versus hand-switch trials, consistent with extensive evidence that trial history impacts hierarchical processing and behavior (Korb et al., 2017; Mayr & Bryck, 2005; Strobach et al., 2012).

This study represents an important step forward in understanding the dynamics of hierarchical processing in human cognition and motor control. We expand beyond previous studies by examining the implementation of hierarchical representations, and characterize the interplay of cognitive and motor processing during hierarchical decision-making. We provide evidence for serialized, top-down hierarchical processing in the context of action selection and a tight coupling between hierarchical decision-making and action selection, furthering our theoretical understanding of the interplay between cognitive systems in decision-making.




Figure 4. Example models of cognitive-motor (C-M) processing.

References

- Badre, D., & Desrochers, T. M. (2019). Hierarchical cognitive control and the frontal lobes. In *Handbook of Clinical Neurology* (Vol. 163, pp. 165–177). Elsevier. <https://doi.org/10.1016/B978-0-12-804281-6.00009-4>
- 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>
- Cellier, D., Petersen, I. T., & Hwang, K. (2022). Dynamics of Hierarchical Task Representations. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0233-22.2022>
- 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., Cavanagh, J. F., & Frank, M. J. (2014). Human EEG Uncovers Latent Generalizable Rule Structure during Learning. *Journal of Neuroscience*, *34*(13), 4677–4685. <https://doi.org/10.1523/JNEUROSCI.3900-13.2014>
- Collins, A. G. E., & Frank, M. J. (2016). *Motor Demands Constrain Cognitive Rule Structures*. <https://doi.org/10.1371/journal.pcbi.1004785>
- Dykstra, T., Smith, D. M., Schumacher, E. H., & Hazeltine, E. (2022). Measuring task structure with transitional response times: Task representations are more than task sets. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-021-02035-3>
- Hardwick, R. M., Forrence, A. D., Krakauer, J. W., & Haith, A. M. (2019). Time-dependent competition between goal-directed and habitual response preparation. *Nature Human Behaviour*, *3*(12), 1252–1262. <https://doi.org/10.1038/s41562-019-0725-0>
- 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. *Journal of Neuroscience*, *37*(33), 7893–7905. <https://doi.org/10.1523/JNEUROSCI.3289-16.2017>
- Logan, G. D. (2018). Automatic Control: How Experts Act Without Thinking. *Psychological Review*, *125*(4), 453–485. <https://doi.org/10.1037/rev0000100>
- 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>
- McDougle, S. D., & Taylor, J. A. (2019). Dissociable cognitive strategies for sensorimotor learning. *Nature Communications*, *10*(40), 1–13. <https://doi.org/10.1038/s41467-018-07941-0>
- Pashler, H. (1984). Processing Stages in Overlapping Tasks: Evidence for a Central Bottleneck. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(3), 358–377.
- Ranti, C., Chatham, C. H., & Badre, D. (2015). Parallel temporal dynamics in hierarchical cognitive control. *Cognition*, *142*, 205–229. <https://doi.org/10.1016/j.cognition.2015.05.003>
- 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>
- Selen, L. P. J., Shadlen, M. N., & Wolpert, D. M. (2012). Deliberation in the Motor System: Reflex Gains Track Evolving Evidence Leading to a Decision. *Journal of Neuroscience*, *32*(7), 2276–2286. <https://doi.org/10.1523/JNEUROSCI.5273-11.2012>
- Sigman, M., & Dehaene, S. (2005). Parsing a Cognitive Task: A Characterization of the Mind’s Bottleneck. *PLOS Biology*, *3*(2), e37. <https://doi.org/10.1371/journal.pbio.0030037>
- Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*(29), 10393–10398. <https://doi.org/10.1073/pnas.0503903102>
- 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>
- Trach, J. E., & McDougle, S. D. (2022). Climbing the Tree: Structured Hierarchical Representations in Visuomotor Maps. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *44*(44). <https://escholarship.org/uc/item/3kq252jz>