Cognitive flexibility versus stability via activation-based and weight-based adaptations

Shengjie Xu*, Tom Verguts, Senne Braem

Department of Experimental Psychology, Ghent University, Belgium

***Corresponding Author:**

Shengjie Xu

Department of Experimental Psychology

Ghent University

Henri Dunantlaan 2, 9000 Ghent

Belgium

Email: Shengjie.Xu@ugent.be

Author Notes

Orcid:

Shengjie Xu: 0000-0003-4604-372X

Tom Verguts: 0000-0002-7783-4754

Senne Braem: 0000-0002-2619-8225

Acknowledgements

This work was supported by an ERC Starting grant awarded to S.B. (European Union's Horizon 2020 research and innovation program, Grant agreement 852570). T.V. and S.B. were supported by FWO project G010319N.

Conflicts of Interest: None

Abstract

Humans are remarkably efficient at adapting to different contextual demands by exerting optimal levels of cognitive flexibility versus stability for switching between different tasks. Here, we developed a recurrent neural network to simulate behavioral indices of cognitive flexibility versus stability and investigated its dynamics. We observed that behavioral adaptation to high- versus low- task switch frequency conditions can be achieved via either fast but transient activation-based adaptations (in activation space), or slow but more enduring weight-based adaptations (shallow versus deep task attractor settings in weight space). Interestingly, like humans, the model further learned to associate contextual features to different weight-based task attractor settings and use this knowledge to shift along a flexibility-stability continuum when encountering these same contexts, suggesting that it is optimal to use such contextual features. In sum, our framework sheds new light on classic measures of cognitive flexibility versus stability when people must switch between multiple tasks, through the lens of activation-based versus weight-based adaptations.

Keywords: cognitive control, activation- and weight-based adaptation, flexibility-stability trade-off, task attractor

Introduction

Much of our everyday behavior requires switching between different task goals. For example, preparing a meal typically involves switching between chopping vegetables, boiling water, seasoning the sauce etc. This highlights the need for cognitive flexibility by efficiently shifting task focus for successful task performance. However, we often also need to maintain focus on a single task (cognitive stability). For instance, it is important to not be distracted by other tasks when chopping vegetables with a sharp knife. Therefore, it is critical to strategically regulate task goals in response to different needs for cognitive flexibility versus stability (Braem & Egner, 2018; Cools, 2019; Dreisbach & Fröber, 2019; Egner, 2023; Garner & Dux, 2023; Goschke, 2013; Hommel, 2015; Musslick & Cohen, 2021).

Cognitive flexibility versus stability is usually studied using task switching paradigms that assess performance differences in speed or accuracy between task switching versus repetition conditions (Kiesel et al., 2010). The modulation of cognitive flexibility versus stability can be seen as a meta-control process, i.e., regulating cognitive control according to the context (Eppinger et al., 2021; Wang, 2021). People can learn to be flexible as indicated through smaller switch costs when faced with higher switching frequencies (Braem & Egner, 2018; Dreisbach & Haider, 2006; Dreisbach & Mendl, 2024; Monsell & Mizon, 2006), and optimally shift between cognitive flexibility versus stability depending on contextual demands (Fröber & Dreisbach, 2017; Xu et al., 2024), demand avoidance (Brosowsky & Egner, 2021), time costs (Mendl & Dreisbach, 2024; Mittelstädt et al., 2018) or reinforcement history (Braem, 2017; Held, Vermeylen, Dignath, et al., 2024; Held, Vermeylen, Krebs, et al., 2024). More broadly, several studies emphasize the importance of this regulation in the development of executive functions, mental disorders, and social interaction (Crone & Steinbeis, 2017; Goschke, 2014; Hommel & Colzato, 2017; Ruel et al., 2021; Stucke & Doebel, 2023). However, it remains unclear how these regulations of cognitive flexibility versus stability are learned and implemented.

Importantly, learning different needs for cognitive flexibility versus stability is thought to occur at different speeds (Egner, 2023). For example, the online regulation of cognitive control in response to currently experienced task switching frequencies occurs fast, but dissipates quickly (Fröber et al., 2022). Humans seem to gradually increase task-relevant activation over successive repetition trials, while increasing switch readiness over switching trials – potentially through a flexible use of different gating mechanisms that manipulate inputs and outputs of working memory (Chatham & Badre, 2015; O'Reilly & Frank, 2006). In contrast to these fast adjustments, the formation of enduring context-control knowledge, i.e., associations between environmental features and different needs for cognitive flexibility (e.g., working memory gating settings), requires more experience and likely reflects slower forms of learning (Braem et al., 2024; Xu et al., 2024).

Although several computational models have proposed different control modules to exert cognitive control (Brown et al., 2007; Gilbert & Shallice, 2002; Musslick & Cohen, 2021; Verguts & Notebaert, 2009), how these models shift between levels of cognitive flexibility versus stability was either preset and not learned by the model, or not studied computationally. Recent neural network studies argue that different task attractor settings underlie this regulation of cognitive control (Jaffe et al., 2023; Musslick & Bizyaeva, 2024; Musslick & Cohen, 2021; Ueltzhöffer et al., 2015). However, also here it remains an open question how one can spontaneously learn different task attractor settings for shifting between cognitive flexibility and stability required in various contexts.

Here, we start from the observation that the impact of experience in a neural network can manifest through two distinct forms: either through changes in the node activation, or through changes in the connection strength between these nodes, i.e., weights (Hummos et al., 2022; O'Reilly & Munakata, 2000; Rao & Ballard, 1999). We propose that a dynamic regulation of flexibility versus stability in cognitive control, can be achieved through both activation-based and weight-based adaptations. First, task activations are influenced by the previous (activation) state. For example, during task switching, starting from a state closer to the attractor of an old task (Fig 1A, brown trajectory) will typically result in a final state further away from the attractor of a new task, relative to starting from a state closer to the attractor of a new task (Fig 1A, blue trajectory). As a result, task switching will be better in the situation represented by the brown trajectory. Such activation-based adaptation manifests itself mostly as a product of recency, or “inertia”, demonstrating the influence of lingering task activation states in working memory, without necessarily learning control parameters (context-control learning) which are instantiated in neural network weights. However, second, the network can also learn different task attractor settings that favor either flexibility or stability via its weights (weight-based adaptation). For example, in a condition with more frequent task switching, requiring higher cognitive flexibility, shallow attractors are favored (Fig 1B, pink dashed line), resulting in lower overall task activations but more efficient updates during task switching; instead, in a condition where there are less switches, deeper attractors are favored (Fig 1B, green solid line), improving performance in task repetitions at a cost of worse performance during task switching. Both types of control adaptations, activation-based and weight-based, can co-exist and indeed mutually determine one another in neural networks (Russin et al., 2024). Here, we propose they jointly contribute to the regulation of cognitive flexibility versus stability.

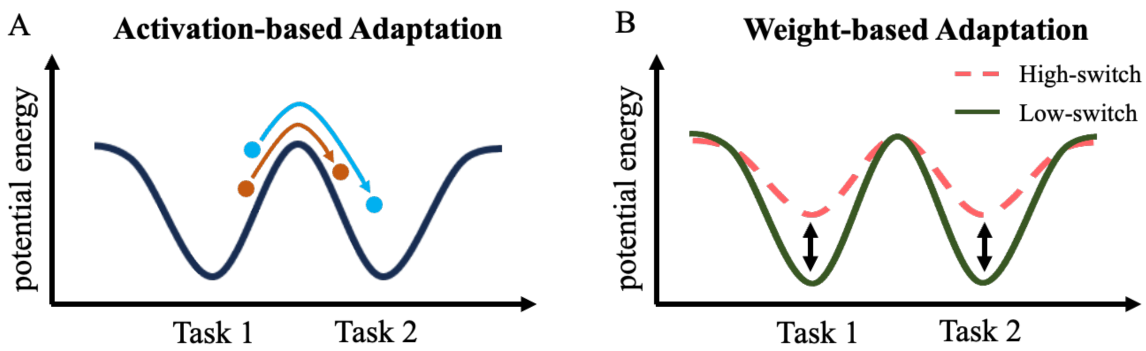

Fig. 1. Illustrations of weight-based and activation-based adaptations. (**A**) An example of activation-based adaptations. When a task goal changes from Task 1 to Task 2, trajectories of blue and brown dots represent corresponding task activation adjustments. Despite the same attractor depth (fixed weights), task switching performance is better for blue dots as its initial state (the left blue dot) starts less deep into the attractor of an old task. (**B**) Weight-based adaptations can be visualized as attractors (Task 1 and Task 2 on the x-axis) with different depths (y-axis) in an energy landscape. The depth of attractors is determined by weights. In a high-switch condition, shallow attractors (pink dotted line) facilitate task switching compared to deep attractors (dark green full line). Humans can adjust the attractor depth to adapt to different task switching conditions, showing context-specific control regulations.

We developed our Learning Control Dynamics (LCD) model to investigate these computational principles of learning and regulating control parameters, based on a recurrent and gating-based neural network (Long Short-Term Memory, LSTM (Hochreiter & Schmidhuber, 1997). It considers both the current task input and the history of previous tasks when generating an output for each step. Importantly, LCD does not include a preset control module, allowing us to study the learning and emergence of cognitive control during task switching.

A classic marker of dynamic control adaptations to varying needs for cognitive flexibility is the observation of a smaller switch cost in a context with more frequent task switching (Braem & Egner, 2018; Dreisbach & Haider, 2006; Dreisbach & Mendl, 2024; Monsell & Mizon, 2006). Therefore, in Study 1, we first demonstrate how this marker of the flexibility-stability trade-off can be driven purely by fast activation-based adaptations in LCD. This indicates that this form of control adaptation can result from recent task exposure without necessarily changing weights. Next, in Study 2, we investigated a slower weight-based form of control adaptation. We found that LCD prepares for task switching by finding the optimal location in the task space (i.e., activation and weight space) that maximizes the activation of the cued task (i.e., an attractor of the cued task), while keeping an optimal level of cognitive flexibility to switch to other tasks via weight-based adaptations. To understand how LCD learns to navigate between tasks with a focus on either cognitive flexibility or stability, we also examined how LCD differentially prepared for upcoming task switches and repetitions in Study 2 after training on high- versus low-switch conditions.

Recent studies in humans further revealed that we can learn to use contextual features to balance the need for cognitive flexibility versus stability, showing (weight-based) context-specific control regulations. For example, we recently demonstrated that humans can learn to associate their control strategies to be more flexible in high-switch versus low-switch conditions to co-occurring *environment* features (Xu et al., 2024). Moreover, humans can also use *task identity* features to increase switch readiness when encountering tasks that appear more often in switching trials (Fröber et al., 2022; Nack & Yu-Chin, 2024; Siqi-Liu & Egner, 2020). Therefore, we next show how LCD can similarly learn to navigate differentially through task space depending on environmental (Study 3) or task identity (Study 4) features for balancing cognitive flexibility versus stability via weight-based adaptations. Based on those simulations, we finally compared our model simulations with an existing human dataset (Siqi-Liu & Egner, 2020) to test novel predictions that emerged from Study 4.

Methods

The LCD model aims to investigate the computational mechanisms that underlie dynamic adaptations of cognitive control and context-control learning in different switching conditions where participants had to switch between multiple different categorization tasks (see also, Xu et al., 2024). We will first explain the general method, after which we will explain the paradigms, the model architecture, and the training regimes in more detail.

To simulate task switching, we asked the model to respond to different task goals on different trials, which was to respond to one dimension of a multidimensional stimulus, dependent on the task cue (Fig. 2A). Therefore, LCD needed to learn to attend to the appropriate stimulus feature based on the task cue, as the other stimulus features could lead to inappropriate responses. Task switching occurs when LCD received two different task cues on two successive trials; while task repetition occurs when the task cue was the same on two successive trials (Fig. 2B).

We always trained LCD in two steps. First, in the task rule learning step (Fig. 2C), we trained LCD's non-recurrent layers how to perform each task. By analogy, this can be thought of as the instruction phase in humans where participants are first equipped with the general task rules. Second, we froze all weights learned in task rule learning step and trained an LSTM layer feeding into LCD's task control module, which we refer to as the task control learning step (Fig. 2D). This LSTM layer received all inputs and was trained on different task switching frequencies that could, in some simulations, be predicted by contextual features. Its output was projected to a feedforward layer of the control module to generate a multidimensional output that served as a new task cue input (i.e., task activations) to the output module for behavioral responding. In this way, the model could learn to employ contextual features as a function of recent history, to exert cognitive control.

Behavioral paradigm

On each trial, LCD received a m-dimensional stimulus, where m is the number of tasks. Each stimulus consisted of m random (real) numbers between zero and one (Fig. 2A). The task goal was to evaluate if the number in the cued location is larger or smaller than 0.5. The task cue was one-hot encoded and was used to indicate which location of the stimulus to respond to (Fig. 2A). Each training sequence contained multiple trials (details given in Model Training Regimes). Switching frequency was defined as the proportion of trials where the task cue was different from the preceding trial, within a single training sequence. Stimuli and training sequences were always randomly generated (constrained by the switching frequencies that was aimed for), so no exact same sequence or stimulus was ever repeatedly used. In Study 3, we also used the environment cues to indicate which switching condition LCD is in (Fig. 2A). This environment cue shared the size of stimulus and task cue inputs, consisting of multiple zeros (low-switch) or ones (high-switch). Hence these environment cues served as a predictive feature of switching conditions. In other studies, the environment cues were always set to zero.

In Study 1 and 2, we manipulated the switching frequency, high-switch (90% task switching) versus low-switch condition (10% task switching), and trained LCD on them separately. In Study 3, we introduced two types of environment cues (Fig. 2A, right panel), which were predictive of high-switch or low-switch conditions respectively, to LCD to investigate whether LCD can learn context (*environment*)-control knowledge and therefore use environment cues for future control regulations. In Study 4, similarly, we used task identity as a predictive feature to investigate context (*task identity*)-control learning at the task level. There were four tasks in Study 4 for a balanced design with two tasks being associated with high-switch while another two being associated with low-switch. In first three studies, switching conditions were manipulated by using a switching frequency of 10% or 90% for training in the low-switch or high-switch contexts respectively, In Study 4, overall switching frequency was set to 50%, while two tasks were randomly selected to appear more in switching trials (90%) and the other two appeared more in repetition trials (90%).

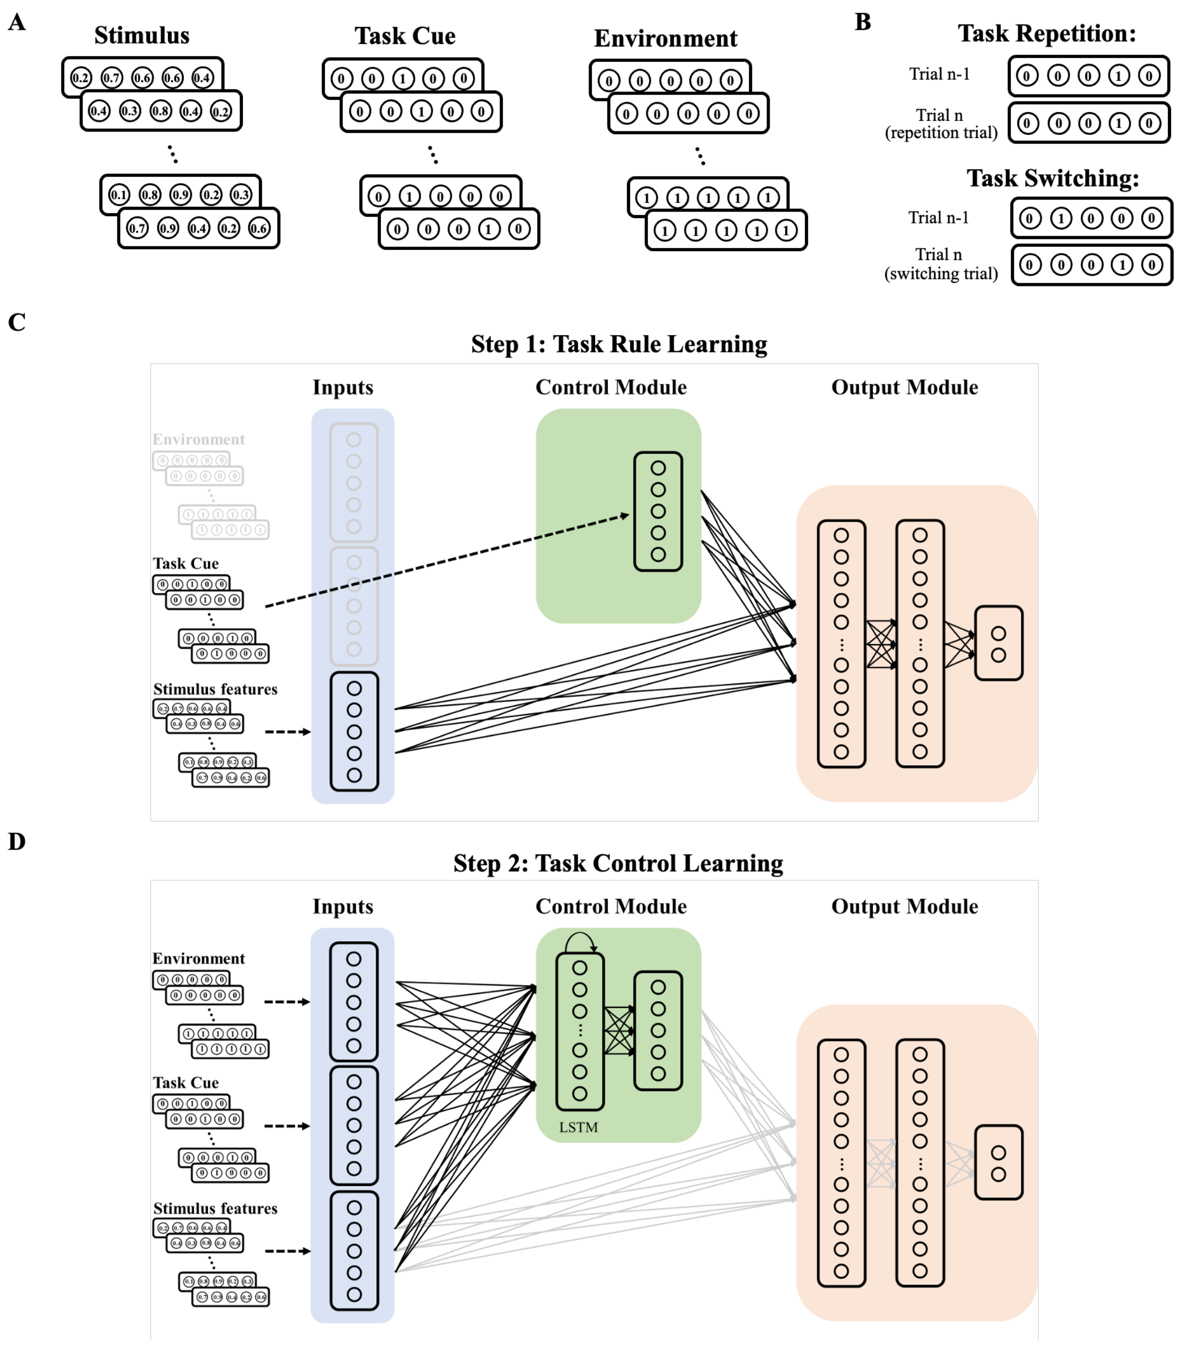

Fig. 2. Training LCD to adaptively switch between different tasks. (**A**) An example of three types of inputs used in model training, stimulus, task cue and environment cue. On each trial, LCD received one row of each input. For each task, the goal is to evaluate if a cued value is below or above 0.5. The stimulus comprises multiple random values in a range of zero to one, excluding the ambiguous 0.5. On each trial, a task cue indicates which feature of the stimulus LCD needed to evaluate, one for the cued stimulus feature and zeros for the others. For instance, if the third value in a task cue input is one (as in the example), LCD should learn to evaluate the third value in a stimulus input on that trial. The environment cue (only varied in Study 3) was provided by multiple separate nodes but consisted of either all zeros or ones indicating the level of task

switching frequency. An environment cue with all zeros implied task switching was less frequent in this environment (i.e., task cues would repeat more often), while the environment with all ones implied more task switching instead. In simulation studies where the environment inputs did not vary, this input was always set as all zeros. (**B**) Illustrations of task repetition and task switching. (**C**) and (**D**) show the LCD's model architecture and our training regime. In the first training step (**C**), the task rule learning step, stimulus and task cue inputs were used to train LCD to learn how to perform tasks one by one. In the second step (**D**), the task control learning step, LCD was trained to optimize its task control parameters and task switching behaviors in response to a list of inputs (i.e., an input sequence) with different switching frequencies. All weights learned from the first step (grey arrows) were fixed in this step. During this phase, an LSTM layer in the control module was trained so that the control module could generate task activations based on all inputs. Different from a binary task cue, the control module output is continuous in a range of zero to one to signify the intensity of the corresponding task activation, with a higher value indicating a greater degree of task activation.

Learning Control Dynamics Model (LCD) Architecture

Control Module

The control module has two layers, a LSTM layer and a feedforward layer. The LSTM layer has 32 units, with a sigmoid recurrent activation function and a ReLU regular activation function. The feedforward layer has *m* units, with a sigmoid activation function. This ensures that the output of the control module on each trial contains *m* values from zero to one, which can then serve as task activations. We used Keras' default settings for all remaining model configurations. In each trial, the control module receives one stimulus, one task cue and one environment cue. These three inputs were concatenated before sending to the first (LSTM) layer of the control module.

Output Module

The output module has three feedforward layers. The first two layers have 16 units with a ReLU activation function. The last layer has two units, with a softmax activation function, to generate the model's choice in each trial. We used Keras' default settings for all remaining model configurations. On each trial, the output module received one stimulus, one task cue (in the *task rule learning step*, see below) or task activations (in the *task control learning step* and simulations, see below), and then generated its choice. These two inputs were concatenated before sending to the first layer of the output module.

Model Training Regimes

Model training and simulations were implemented in Keras, version 2.13.1. There were two training steps in total, with 100 training runs in each study. In the *task rule learning step* (Fig. 2C), we trained LCD's non-recurrent layers (output module) how to perform each task. More specifically, the output module was trained in the *task rule learning step* to generate choices based on task cue inputs and stimulus inputs. As mentioned above, this can be thought of as the instruction phase in humans where participants are first equipped with the general task rules. Second, LCD learned how to regulate its control parameters in response to different needs for cognitive flexibility versus stability at the *task control learning step* (Fig. 2D). This second step allowed learning context-specific cognitive control weights. Similarly, in the *task control learning*, the output module still expected to receive a task cue and a stimulus to generate a choice. In this sense, the output of the control module was supposed to send task relevant information to indicate the subsequent output module how to make its choice. Due to the recurrent nature of the LSTM layer, the output of the control module depends on current inputs and the module's historical states. In this step, we manipulated the switching condition in training sequences. Therefore, the control module should take all inputs into consideration for its output generation, instead of only focusing on a given task cue input in a certain trial.

Task Rule Learning Step

All four studies shared the same *task rule learning step*. In this step, we directly passed task cue inputs together with stimulus inputs to the output module. There were 400 trials in each training sequence. We updated LCD’s weights every four training sequences. Tasks were totally interleaved in each training sequence. We used Keras’ built-in Adam optimizer, with a learning rate of 0.001, to update model weights based on categorical cross-entropy loss. At the end of each epoch, we evaluated model accuracy on each task separately. The training stopped when LCD reached at least 95% accuracy on each single categorization task.

Task Control Learning Step

The *Task Control Learning Step* was subject to switching condition manipulations in each study. First, we froze all weights learned in the *task rule learning step* and trained an LSTM layer feeding into LCD’s task control module. This LSTM layer received environment cues, stimulus and task cues and was trained on different task switching frequencies that could, in some simulations (see below), be predicted by contextual features. Its output was projected to a feedforward layer of the control module to generate a multidimensional output that served as a new task-cue-like input (i.e., task activations) to the output module for responding.

In Study 1 and 2, we had 100 independent training runs in high-switch and low-switch conditions respectively. Similarly, in Study 3 and Study 4, LCD was trained on high-switch and low-switch conditions together with associated contextual features until reaching 85% accuracy. More specifically, in Study 3, the same LCD model was trained on both high and low switching regimes, and the regime was indicated by the environment inputs (Fig. 2A), five zeros meaning a low-switch condition while five ones meaning a high-switch condition. In each epoch, LCD was trained on one switching regime together with a corresponding environment input. To keep the model’s task switching performance induced by each environment input comparable (Xu et al., 2024), model accuracy on each environment input with a switching frequency of 50% was evaluated at the end of each epoch. The environment where LCD had worse performance was trained in the next epoch. For example, if LCD had lower accuracy in a low-switch environment (i.e., five zeros) with a switching frequency of 50% in the evaluation, then in the next epoch, LCD was trained with a low-switch environment input with a switching frequency of 10%. The first training condition was randomly chosen in Study 3. In Study 4, we employed only four tasks, and the switching condition was indicated by the task identity. Specifically, two tasks were randomly selected to occur more frequently on switching trials (90% switching, high-switch tasks), while the other two tasks were more prevalent in repetition trials (90% repetition, low-switch tasks), with an overall switching frequency of 50% in each epoch.

In each training run, initial LCD weights were sampled from a normal distribution with a mean of zero and standard deviation of 0.001. There were 50 trials in each training sequence, and weights were updated every four training sequences. The last state in each batch was not passed to the next batch during training. Again, no training sequence was repeated. In addition, during training, we regularized the connection between the LSTM layer and the feedforward layer in the control module with a dropout rate of 0.5. We updated model weights according to categorical cross-entropy loss, using the Adam optimizer with a learning rate of 0.001. To match human performance (Xu et al., 2024), in this step, we trained LCD until it reached at least 85% accuracy in a switching condition with a switching frequency of 50% in each context type, which was tested at the end of each epoch. We let LCD freely decide how to optimize its behavior in each switching condition to reach the target accuracy. To ensure a similar training procedure, we discarded models if they did not reach 85% accuracy in this step after 300 epochs.

Simulations and Analyses

Model Simulations

In all simulations, learned weights after model training were frozen to avoid any further learning. Therefore, each simulation served as a testing phase to study how LCD adapts to different switching conditions after learning. We ran 100 simulations in each study. There were 1000 trials in each simulation. We recorded model accuracy as well as task activations in each trial for following analyses.

In Study 1 we studied whether pure activation-based dynamics could contribute to task switch costs. Here, low-switch trained LCD was subjected to either high (90%) or low (10%) task switching frequencies in simulations. Because all weights were learned from the same training regime and fixed during test, weight-based adaptation could not explain any switch cost difference in high-switch versus low-switch testing conditions in Study 1.

In Study 2, conversely, we aimed to investigate the weight-based adaptation. Therefore, we tested the low-switch trained (from Study 1) and a new high-switch trained LCD with frozen weights in a test with unbiased switching frequency of 50%. This time, all (stimulus, task) sequences in Study 2 shared the same switching frequency.

In Study 3, LCD was again tested with frozen weights on a switching frequency of 50% together with environment cues that were predictive of switching frequencies during training, to test environment-control learning. Similarly, in Study 4, four tasks were presented equally often in switching and repetition trials in simulations with a switching frequency of 50% to test task-control learning.

Model Behavioral Analysis

Similar to conventional task switching analyses in humans, we excluded the first and post-error trials in model simulations as they cannot be defined as neither task switching nor task repetition. Then we computed model's behavioral accuracy in switching trials and repetition trials in each context.

Task Activation Adjustment Analysis

We computed task activation adjustment velocities to quantify the extent of adjustments on task activations, via measuring Euclidean distance between m-dimensional task activations $[A_{1,n-1}, A_{2,n-1}, A_{3,n-1}, \dots, A_{m,n-1}]$ and $[A_{1,n}, A_{2,n}, A_{3,n}, \dots, A_{m,n}]$ on two successive trials, trial *n*-1 and trial *n*. Here $A_{i,n}$ stands for task activation in trial *n* on task $i \in \{1, 2, \cdots m\}$. The sign of task activation adjustment velocities depends on whether LCD moved in a direction that increases the activation of the cued task from trial *n*-1 to trial *n*. If trial *n* is a switching trial, then the measured velocity is a switching velocity, otherwise a repetition velocity. To decompose the switching velocity, we computed the task activation changes. We defined three groups of tasks, the previously executed task (i.e., the task performed in trial *n*-1), the current task (i.e., the task performed in trial *n*) and other tasks that were not cued on these two trials. The activation change of each task was computed by subtracting the activation of this task in trial *n*-1 from the activation of the same task in trial *n*. To interpret the repetition velocity, we extracted trials from each simulation where LCD experienced the same task on five successive trials to look at how task activation accumulated over repetition trials. In these five trials, the first one is a switching trial, where LCD switched from another task to the current task and repeated performing this task in following four trials. In the variant version (in Study 4), we instead measured switching and repetition velocities from trial *n* to trial *n*+1 to study the task-specific regulation effect after experiencing high-switch and low-switch tasks. The switching velocity decompositions were also implemented accordingly.

Task Representation Separation

To measure how separately LCD represented each task in each context, we first located a task representation centroid as a reference point in each simulation data. A task representation centroid is defined by averaging task activations over simulation trials, resulting in an m-dimensional vector $[\overline{A_1}, \overline{A_2}, \overline{A_3}, \dots, \overline{A_m}]$, where $\overline{A_i}$ stands for an average task activation of task i, $i \in \{1, 2, \cdots m\}$, and m is the number of tasks. Then we computed a Euclidean distance between the task representation centroid and an m-dimensional task activation in each trial with a correct response. Afterwards, the representation separation of each task was measured by averaging each task's own computed Euclidean distance values.

Model Statistics

We conducted a 2 (context: high and low switching regimes) × 2 (trial type: switching and repetition) analysis of variance (ANOVA) on behavioral accuracy data in all studies, where we set context as a between-subject factor in the study 2 but a within-subject factor in the other studies. We conducted two-sided t tests to compare accuracy between contexts on task switching and repetition trials respectively, task representation separation, switching / repetition velocities, task activation changes and task activation accumulation.

Human Dataset Reanalysis

The human dataset used in Study 4 is from Experiment 3a and 3b of Siqi-Liu & Egner, 2020. Two experiments share a similar design, where participants switched between three tasks with different switching frequencies. We used their data from a high-switch condition with a switching frequency of 70%. There were two biased tasks and one unbiased task in each block. Biased tasks appeared more in switching trials, but unbiased tasks appeared equally in switching and repetition trials, resulting in an overall switching frequency of 70%. We applied the same data preprocessing and exclusion strategy (see original paper for task procedure and analysis details). We relabeled their biased tasks as high-switch tasks and their unbiased task as a low-switch task. To gain higher statistical power, we combined data from two experiments, and therefore reached a sample size of 102. We ran a repeated measure two-way ANOVA of trial type (switching and repetition) and task type (high-switch tasks and low-switch tasks), and then another repeated measure two-way ANOVA of trial type (switching and repetition) and previous task type (after high-switch tasks and after low-switch tasks) to examine whether the switch cost would be smaller in trials with high-switch tasks or after receiving high-switch tasks.

Data and Code Availability

Human behavior data can be found via the link provided in the original paper (Siqi-Liu & Egner, 2020). All code used to train and analyze the model, trained models and simulation results in this manuscript are available: *https://osf.io/cqhvx/*.

Results

Activation-based adaptations in response to different needs for cognitive flexibility

Study 1 evaluated whether the model would show a modulation in switch costs when LCD was trained on the low-switch condition, but tested under low- versus high-switch conditions with its weights frozen during the test. Results showed that the classic observation of smaller switch costs in high versus low switch conditions can indeed be explained by dynamic variations in (accumulated) task activation caused by the experienced high- versus low-switch frequency conditions. Specifically, in a low-switch testing condition, LCD experienced more task repetitions, and therefore had more opportunities to get closer to a given task attractor (e.g., Fig. 3A, blue dots). This helped LCD accumulate task activation over subsequent task repetitions, improving performance on task repetitions but impairing performance on subsequent switch trials. To quantify activation-based adaptation, we computed, similar to humans, the switch cost on LCD's accuracy data. As predicted, we observed a significant switch cost, which was significantly smaller when switching frequency was higher (Fig. 3B). Specifically, LCD had better accuracy in switching trials, but worse in repetition trials when faced with a higher switching frequency (Fig. 3B). Previous work argued that a deeper task attractor is favored when the need for task switching is low, which leads to more separated task representations and vice versa (Jaffe et al., 2023; Musslick &

Cohen, 2021). Therefore, we measured the distance between LCD's activation state of each task and its average activation state (see Methods) to indicate task representation separation in Study 1. Notably, even though the weights were fixed, we observed less separation in task representations (Fig. 3C) when LCD was tested in a high-switch condition. This indicates that LCD dynamically adapted to emerging needs for control, even after its weights are frozen, via activation-based adaptation.

An important consequence of these findings is that they emphasize the importance of using equal testing conditions when studying control regulations. Otherwise, under unequal testing conditions, the observed modulation of switch costs may be driven either by (activation-based) variations in task activations induced by the locally experienced switching frequencies, or instead by (weight-based) enduring changes in control parameters (Braem et al., 2019, 2024; Frings et al., 2023; Wang et al., 2018) or both. In fact, we believe the activation-based variation may be a more likely explanation in most previous studies, given that learning weight-based adaptation of control strategies often requires considerable amounts of task and context-specific training beyond the length of a single experimental session (Braem et al., 2024; Xu et al., 2024).

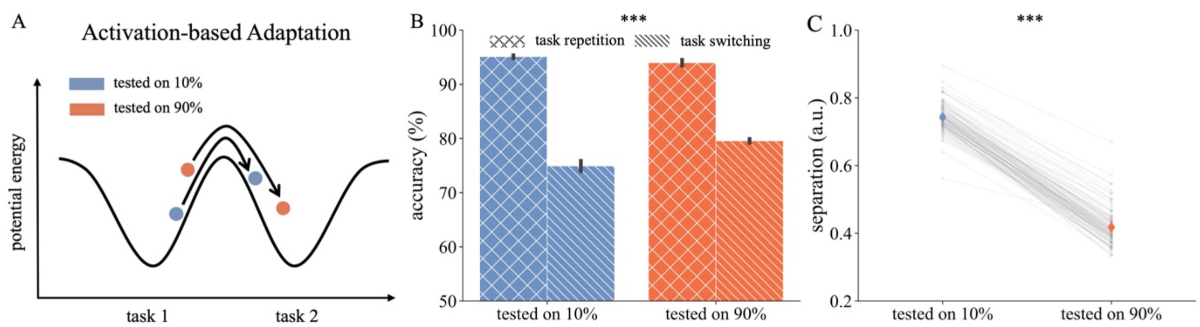

Fig. 3. Activation-based adaptation in Study 1. (**A**) When tested on 10% task switching, LCD has more chances to accumulate task activation over repetition trials. Consequently, LCD is closer to one task attractor, which benefits task repetition, but impairs its performance on task switching, i.e., farther away from the cued task attractor on a switching trial. (**B**) LCD behavioral results of LCD with a low-switch training regime in the testing phase with 10% (blue) and 90% (orange) task switching frequencies respectively. (**C**) LCD task

representation separation results with a low-switch training regime in the testing phase with 10% (blue) and 90% (orange) task switching frequencies respectively. Error bars stand for the 95% confidence intervals, *** represents $p < .001$.

Weight-based adaptations through different task activation velocities

In Study 2, we wanted to investigate whether LCD can also learn such weight-based adaptation in response to different needs for flexibility. Our reasoning was opposite to Study 1, as the two groups of simulation data were now from LCD that was trained on different needs for cognitive flexibility (i.e., weight-based adaptation), but experienced the same proportion of task switching and repetition in the testing phase, hence reflecting effects from weight-based control learning instead of task activation variations induced by different switching frequencies.

Indeed, we found that LCD still showed smaller switch costs in the testing phase when trained on a high-switch condition (Fig. 4B). Specifically, LCD had better accuracy in switching trials, but worse in repetition trials after a high-switch training regime compared to a low-switch training regime. As for task representation separation, LCD with a high-switch training tends to represent tasks closer together (Fig. 4C), suggesting shallower task attractors. Thus, LCD acquired different strategies for the regulation of cognitive control through weight-based learning during low- versus high-switch frequency conditions.

We analyzed the outputs from LCD's control module to unveil how LCD adjusted its task activation to meet the varying needs for cognitive flexibility. Because we implemented multiple tasks, LCD needed to navigate a high-dimensional task space. To quantify the dynamics of task activations from trial to trial, we developed a novel metric that we call task activation adjustment velocity. This metric quantifies how efficiently LCD moved away from its old location in the task space in a correct direction toward the currently cued task, during task switching and task repetition separately. We hypothesized that LCD would strategically modulate its task activation adjustment velocity in response to different needs for control. Specifically, we expected it would show a higher task activation adjustment velocity (switching velocity, Fig. 4D) in switching trials when moving toward the cued task attractor away from an old task attractor, after a high-switch training regime as opposed to after a low-switch training regime. Conversely, LCD had a higher task activation adjustment velocity (repetition velocity, Fig. 4G) while moving toward the same task attractor in a repetition trial, after a low-switch training regime as opposed to after a high-switch training regime. Due to this higher repetition velocity, a low-switch training should drive LCD toward stronger task activation over successive repetition trials.

Using the same simulation results as presented above, we analyzed LCD's task activations after high-switch and low-switch training regimes in a testing phase with an equal switching frequency (50% task switching) in Study 2. In line with our expectation, LCD had a higher switching velocity in switching trials after a high-switch training regime (Fig. 4E). To unpack, LCD increased its activation of the cued task more (Fig. 4F) and reduced the activation of the previously executed task more after a high-switch training. Moreover, LCD also showed a smaller reduction in activation of other tasks after a high-switch training regime. These findings suggest that LCD learned the strategy to move faster toward the cued task attractor in the task space during task switching.

Analogous to the switching velocity metric used in task switching, we also calculated velocity during task repetition. For a comparable measurement, we calculated the repetition velocity on a task repetition on the first repetition trials (i.e., Fig. 4I, from trial 1 to trial 2). Indeed, the repetition velocity was higher after a low-switch training regime (Fig. 4H). To examine whether LCD converged to a stronger task activation after a low-switch training regime, we assessed task activation dynamics during task repetition over multiple successive trials, i.e., repeating a task once (trial 2), twice (trial 3), etc. We expected that LCD would accumulate activation of a cued task over successive repetition trials, but with a different pattern for low-switch and high-switch training regimes. For this purpose, we tracked the activation of the same cued task across five successive repetition trials (Fig. 4I). In the first trial, LCD's activation of the cued task was stronger after a high-switch training regime. It is because this trial is a switching trial (see Methods), where LCD moved from another task to this cued task. However, activation strength of the cued task after a low-switch training regime surpassed those from LCD after a high-switch training regime from the second trial onwards, finally converging to a higher task activation value.

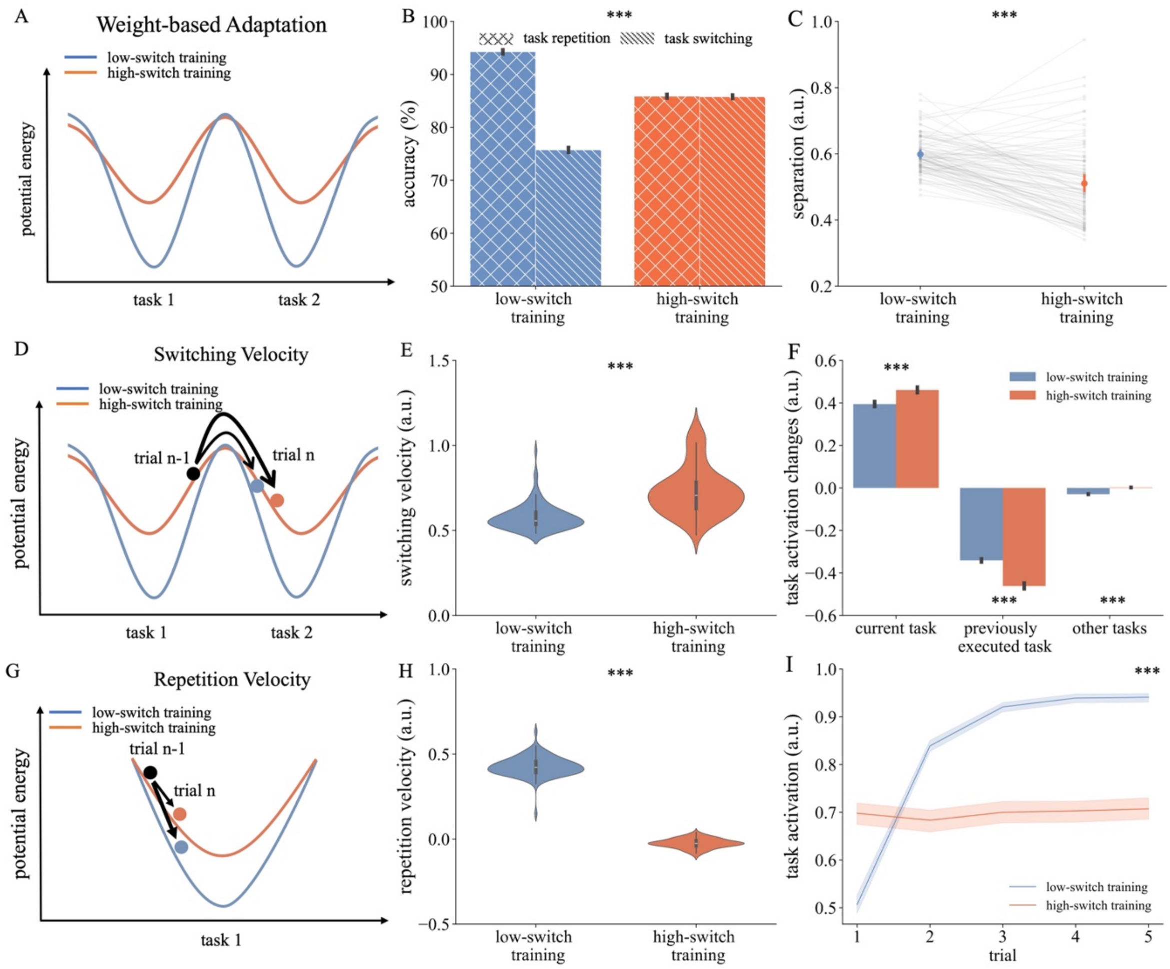

Fig. 4. Weight-based adaptation in Study 2. (**A**) A visual illustration of weight-based adaptation. LCD learned to reduce attractor depth to deal with task-switching after a high-switch training regime (orange line). In a low-switch training regime, LCD increased attractor depth (blue line). Therefore, a high-switch training regime enables LCD to move its task activation state faster in a switching trial, resulting in a task activation state closer

to the cued task from trial n - 1to trial n. (**B**-**C**) Behavioral accuracy and task representation separation results of simulation data from low-switch and high-switch trainings in the testing phase with 50% task switching respectively. (**D**) A visual illustration of switching velocity. A shallower attractor facilitates the movement from one task on trial n – 1 to another task on trial n, resulting in a final task activation state closer to the new task attractor. (**E**) LCD switching velocity results in the testing phase after low-switch and high-switch training. (**F**) Task activation trial-by-trial changes on current task, previously executed task and other tasks during task switching. Here, current task means the task in the current switching trial and previously executed task means the task performed in the last trial, while other tasks refer to three remaining tasks. (**G**) A visual illustration of repetition velocity. In a repetition trial, LCD benefitted from a deeper attractor with a low-switch training, enabling it to move toward the task attractor more efficiently. (**H**) LCD repetition velocity results in the testing phase after low-switch and high-switch training. (**I**) LCD task activation accumulation over task repetitions in the testing phase after low-switch and high-switch trainings, with numbers on x-axis representing repetition trial numbers. Error bars and shade area stand for the 95% confidence intervals. In violin plots, white dots in boxplots indicate medians and upper and lower hinges of the boxplot correspond to the first and third quartiles. *** represents $p < .001$.

Learning environment-specific cues signaling varying needs for cognitive flexibility

To examine whether LCD emulates this type of environment-control learning at the behavioral level, we conducted Study 3 by presenting different switching frequencies (high versus low) associated with different environment cues, and trained LCD to see if it was able to learn these environment-control contingencies. We expected LCD to learn these environment-specific control parameters and therefore established an environment-specific control regulation strategy (Fig. 5A). Indeed, we observed smaller switch costs for the environment cues associated to a high-switch condition (Fig. 5B), suggesting that LCD learned to apply environment-specific task switching strategies. LCD also tended to have less separated task representations while cued by a high-switch environment cue (Fig. 5C), allowing for a higher degree of cognitive flexibility.

Similar to results in Study 2, LCD learned to use different task activation adjustment velocities depending on the environment it was in. On the one hand, when presented with a high-switch environment cue, LCD learned to use a faster switching velocity during task switching (Fig. 5D), as opposed to when presented with a low-switch environment cue (Fig. 5E). Specifically, when LCD was presented with a high-switch environment cue, the cued task activations increased more, with more reduction in activation of the previously executed task but less reduction in activations of other tasks (Fig. 5F). On the other hand, LCD had a larger repetition velocity (Fig. 5H), and converged to a larger value (Fig. 5I, trial 5), when performing task repetitions when presented with a low-switch environment cue. These results support that LCD can learn environment-control associations, showing environment-specific regulations of control to achieve optimal levels of cognitive flexibility versus stability.

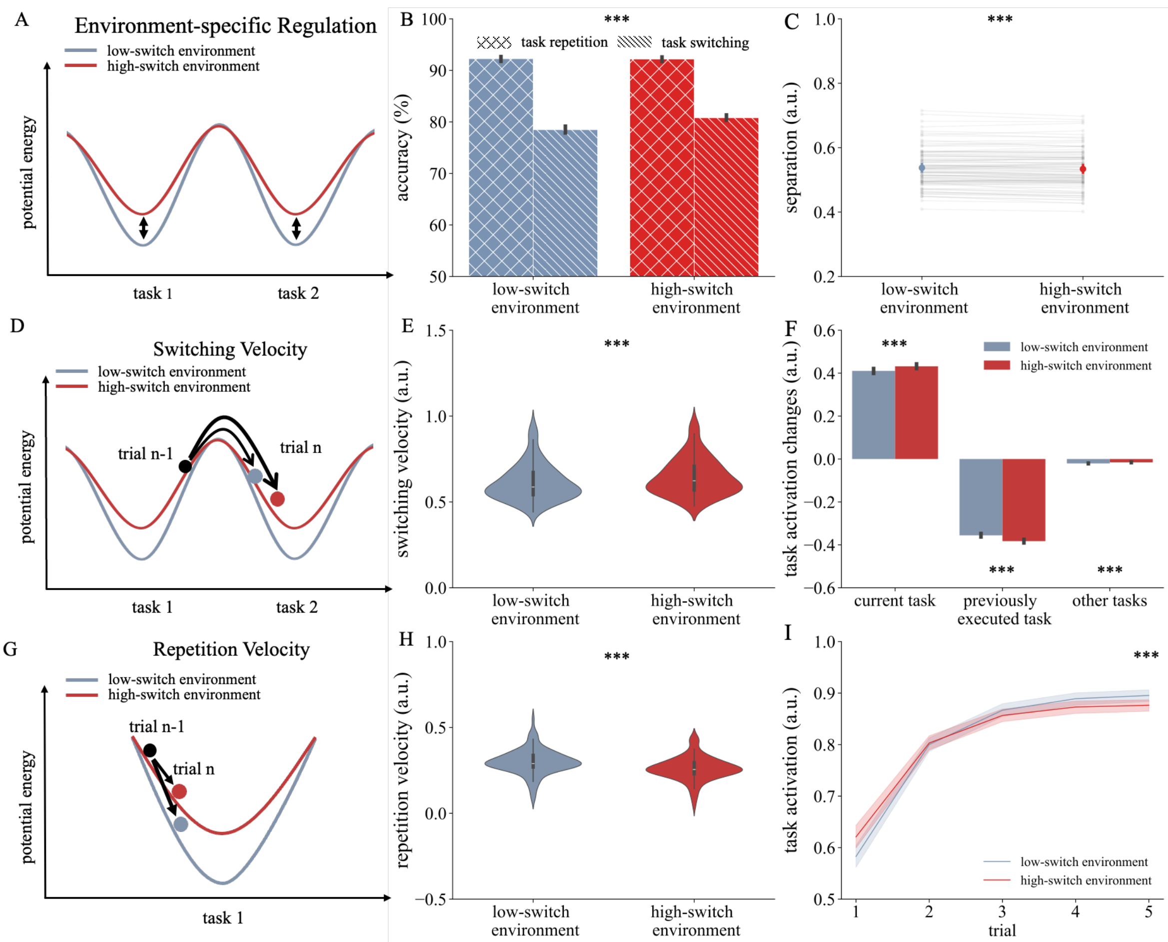

Fig. 5. Environment-specific control regulation in Study 3. (**A**) A visual illustration of environment-specific weight-based control regulation. LCD was trained on both high- and low-switch condition, indicated by

different environment cues. Therefore, LCD was able to adjust its attractor settings based on environment cues, showing shallower attractors when cued by a low-switch environment cue but deeper attractors when cued by a high-switch environment cue. (**B**-**C**) Behavioral accuracy and task representation separation results in study 2 when LCD was tested on 50% task switching when cued by different environment cues. (**D**) A visual illustration of switching velocity difference induced by environment-specific control regulation. During task switching, LCD moves faster (shown by a longer trajectory between a black and a red dot) in the task space when cued by a high-switch environment cue compared to a low-switch environment cue. (**E**) Switching velocity results in the testing phase when cued by different environment cues. (**F**) Task activation changes on current task, previously executed task and other tasks during task switching when cued by different environment cues. (**G**) A visual illustration of repetition velocity difference induced by environment-specific control regulation. In a repetition trial, LCD moved closer to a task attractor when cued by a low-switch environment cue, shown by a longer trajectory between a black point and a light blue point. (**H**) Repetition velocity results in the testing phase cued by different environment cues. (**I**) Task activation accumulation over task repetitions in the testing phase when cued by different environment cues. Error bars and shade area stand for the 95% confidence intervals. In violin plots, white dots in boxplots indicate medians and upper and lower hinges of the boxplot correspond to the first and third quartiles. *** represents $p < .001$.

Learning task-specific cues signaling varying needs for cognitive flexibility

Finally, having established that LCD can learn to apply different control parameters for high- versus slow-switch environments, we investigated whether and how LCD can also adjust its control strategy in response to *task identity* features that were predictive of different needs for cognitive flexibility (i.e., high-switch and low-switch conditions). In line with behavioral findings, we found task-specific switch costs, showing smaller switch costs on high-switch tasks (Fig. 6B). Similarly, we again observed that LCD represented high-switch tasks closer together than low-switch tasks (Fig. 6C). Thus, LCD can learn to regulate cognitive control strategies in accordance to varying needs for cognitive flexibility predicted by *task identity* features, consistent with recent work on human task-specific control regulation.

To study how LCD learned to be more flexible on high-switch versus low-switch tasks, we analyzed its velocity data. Despite the fact that LCD exhibited task-specific cognitive flexibility and represented high-switch tasks closer together, but in contrast to our previous simulation studies, there was no difference in switching velocity, when facing low-switch vs. high-switch tasks (Fig. 6E). Similarly, when analyzing changes in task activation during task switching (Fig. 6F), we did not observe a higher increase in activation of the cued high-switch tasks in switching trials, nor changes in activation of the other tasks. Instead, we only observed a smaller reduction in activations of previously executed tasks. However, the repetition velocity and the task accumulation results were consistent with previous findings, showing LCD had a faster repetition velocity (Fig. 6H), and converged to higher task activation (Fig. 6I, trial 5) on low-switch tasks compared to high-switch tasks.

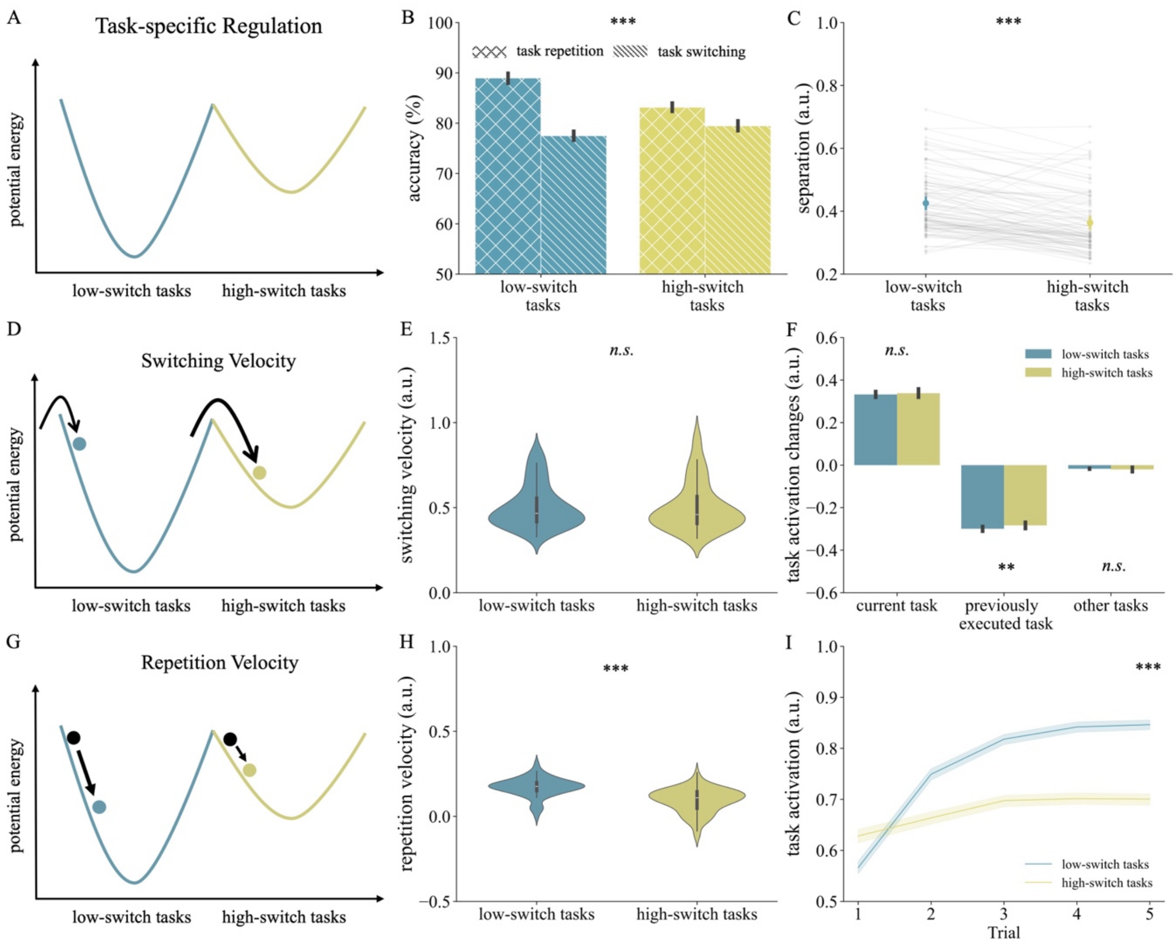

Fig. 6. Task-specific control regulations in Study 4. (A) A visual illustration of task-specific weight-based

control regulation. LCD was trained on two groups of tasks where two tasks appeared more in repetition trials (low-switch) while another two tasks appeared more in switching trials (high-switch). Therefore, LCD developed deeper attractors for low-switch tasks but shallower ones for high-switch tasks (**B**-**C**) Behavioral accuracy and task representation separation results when LCD was tested on 50% task switching with low-switch versus high-switch tasks. (**D**) A visual illustration of switching velocity difference induced by task-specific control regulation. LCD is supposed to move faster in the task space when it performs a high-switch task compared to a low-switch task on a switching trial. (**E**) Switching velocity results on low-switch and high-switch tasks in the testing phase with low-switch versus high-switch tasks. (**F**) Task activation changes on current task, previously executed task and other tasks during task switching in the testing phase with low-switch versus high-switch tasks. (**G**) A visual illustration of repetition velocity difference induced by task-specific control regulation. In a repetition trial, the task-specific model moves closer to a task attractor when it performs a low-switch task. (**H**) Repetition velocity results in the testing phase with low-switch versus high-switch tasks. (**I**) Task activation accumulation over task repetitions in the testing phase with low-switch versus high-switch tasks. Error bars and shade area stand for the 95% confidence intervals. In violin plots, white dots in boxplots indicate medians and upper and lower hinges of the boxplot correspond to the first and third quartiles. *n.s.* indicates non-significancy, ** represents $p < .01$, *** represents $p < .001$.

Together, these results seemingly indicate that the observed task-specific control regulation at the behavioral level may not be driven by a strategic adjustment of the switching velocity, but only the repetition velocity in response to low-switch and high-switch tasks. However, different from the environment cues which were presented throughout the whole experiment in Study 3, *task identity* features do not suggest an optimal global strategy, but rather operate on a trial-by-trial basis in Study 4. When looking at switching trials where a high-switch task was presented, the measures of switching velocity were averaged across conditions where the previous trial was a high-switch or low-switch task, thereby eliminating the potential after-effects of switching velocity induced by the two types of tasks (Fig. 7A). Instead, in repetition trials where a low-switch (versus high-switch) task was presented, we know the previous task was another low-switch task, so the resulting repetition velocities could also be the product of the previous trial(s). Therefore, LCD probably did not use learned task-control knowledge to deal with task switching within the same trial. Instead, it may apply the learned knowledge after receiving a low-switch or high-switch task in a previous trial (Jia et al., 2024) (Fig. 7B-C).

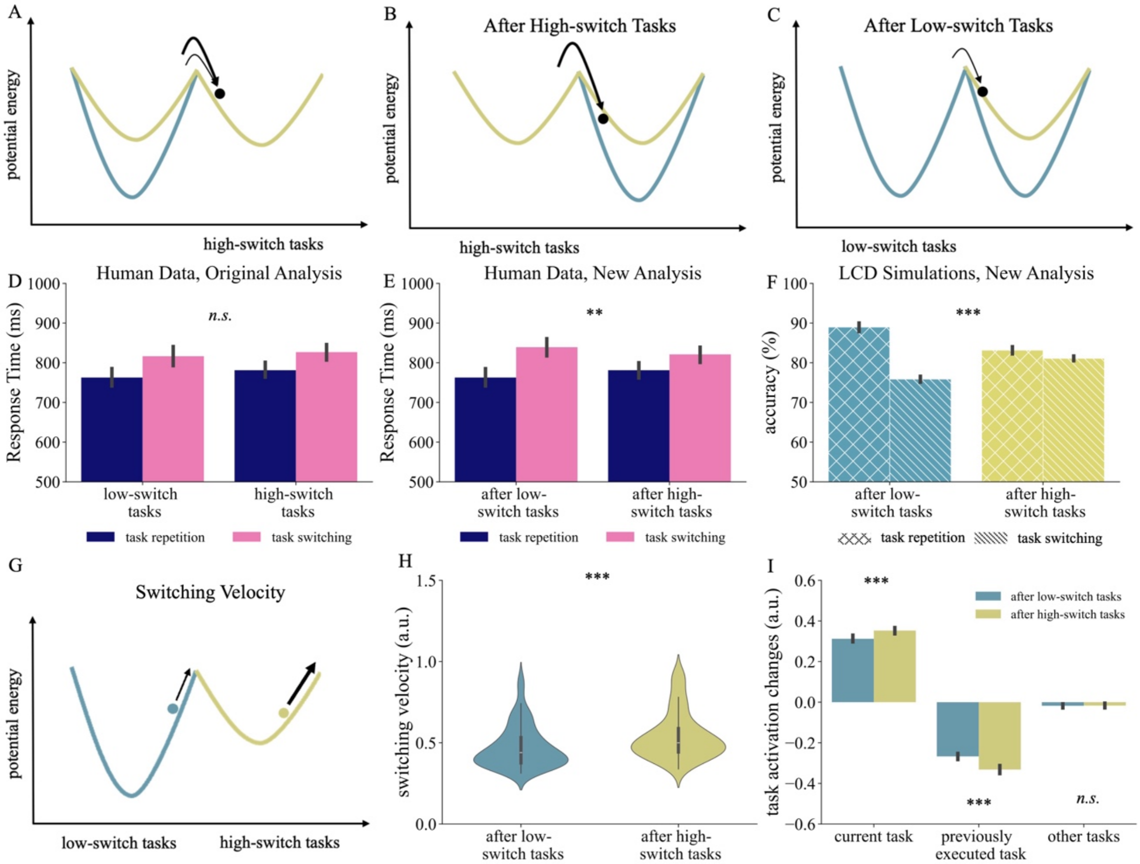

Fig. 7. Task-specific switching velocity after receiving low-switch and high-switch tasks in Study 4. (**A**) Unlike environment-specific switching velocity, task-specific switching velocity is influenced by which type of task is presented in a previous trial. Directly looking at switching trials, where for example a high-switch task is presented, averages the task-specific switching velocity induced by low- and high-switch tasks in previous trials, hence eliminating the switching velocity difference between low-switch and high-switch tasks. To observe the task-specific switching velocity, one should measure the effect after receiving high-switch tasks (**B**) or low-switch tasks (**C**). (**D-E**) Human behavioral results of difference analyses on an existing human task-specific experiment dataset by looking at task-specific effect in trials presenting low- or high-switch tasks (**D**) or following low- or high-switch tasks (**E**) respectively. (**F**) LCD behavioral accuracy in trials following low- or

high-switch tasks. (**G**) A visual illustration of the new measure of switching velocity in task-specific control regulation. Instead of looking at switching to each type of task, we now look at the velocity of switching away from each type of task. (**H**) Switching velocity results in the testing phase, after receiving low- and high-switch tasks. (**I**) Task activation changes on current task, previously executed task and other tasks during task switching in the testing phase, after receiving low- and high-switch tasks. Error bars stand for the 95% confidence intervals. *n.s.* indicates non-significancy, *** represents $p < .001$.

To directly test whether LCD adjusted its switching velocity only after receiving high-switch or low-switch tasks, we analyzed LCD's behavioral accuracy on trials *after* high-switch or low-switch tasks on previous trials (Fig. 7F). The results again showed task-specific switch costs, with lower switch costs after receiving high-switch tasks. Next, we computed switching velocity after low-switch and high-switch tasks. This time, we did observe a higher switching velocity on switching trials *after* high-switch tasks (Fig. 7H). In addition, the decomposition of this new switching velocity result pointed to a higher activation increase of the cued tasks as well as a higher reduction in activation of the previously executed tasks on switching trials following high-switch tasks (Fig. 7I). Notably, there was still no difference in reduction in activations of the other tasks, further suggesting that this learned task-specific control regulation was applied in a very task-specific trial-by-trial manner.

These results predict that learned task-specific requirements for cognitive flexibility are more pronounced after receiving predictive task identity features. To further investigate these new predictions, we reanalyzed an earlier human dataset (Siqi-Liu & Egner, 2020). When directly looking at the task type (high-switch versus low-switch), there is a significant main effect of trial type, $F(1,101) = 203.87$, $p < .001$, $\eta_p^2 = .669$, but no main effect of task type or their interaction $F(1,101) < 2.57$, $p > .112$, $\eta_p^2 < .025$. However, when looking after biased tasks, the results showed the critical interaction of previous trial type and trial type, $F(1,101) = 8.77$, $p = .004$, $\eta_p^2 = .080$. A smaller switch cost was observed in trials *after* switch-predictive tasks (Fig. 7E, switch cost difference *M* = -36.39, 95% confidence interval CI = [-60.48, -12.30]), but not on them (Fig. 7D, *M* = -8.11, 95% (CI) = [-18.70, 2.47]).

In sum, although able to learn from both environmental and task identity features predictive of different levels of task switching, LCD applied this learned context-control knowledge in different ways, depending on the characteristics of the predictive contextual information, i.e., global environment vs. trial-by-trial task information.

Discussion

Over the past decades, many cognitive control studies yielded diverse insights into how people tradeoff cognitive flexibility versus stability depending on task characteristics. Here, we show how classic markers of these opposing states can be understood through the lens of both fast but fleeting activation-based adaptation, and slower but more enduring weight-based adaptation. Together, these findings shed an important, new light on the interpretation and implementation of (context-specific) cognitive flexibility versus stability. Our LCD model is also able to learn and associate these different task activation adjustment strategies to co-occurring contextual features, such as the identity of the environment or task in which agents are subjected to different degrees of task switching. These results help ground contemporary theories of cognitive control and offer a useful framework to understand old and new debates in the field – several of which we will briefly discuss below.

Recent years have seen an emerging interest in using attractor dynamics (Hopfield, 1982; Rolls et al., 2008; Yang et al., 2019) to understand the computational mechanisms of cognitive control (Grahek et al., 2024; Haykin et al., 2012; Jaffe et al., 2023; Khona & Fiete, 2022; Musslick & Bizyaeva, 2024; Musslick & Cohen, 2021; Ueltzhöffer et al., 2015). However, an open question has been whether and how humans learn these different control states. LCD promotes our understanding of not only how task attractors are shaped in the first place, but also how different task attractors are learned in response to varying environmental demands. One critical contribution is the distinction between activation-based and weight-based adaptations on task performance. Here, we show how LCD can both differentially benefit from accumulated task activation in low- versus high-switch environments, as well as learn to adaptively shape the depth of its attractor landscapes.

This cooperation between activation-based and weight-based adaptations is reminiscent of the recent distinction between in-weight learning versus in-context learning as manifested in recurrent neural networks (Wang et al., 2018) and transformers (Vaswani et al., 2017). Any system that learns in weights (e.g., a neural network) but also shows sufficiently complex and sustained activation dynamics (e.g., a recurrent network or a transformer), can sculpt its weights so that it can also learn in activation space (in-context learning). Specifically, during (weight) training such networks learn how to use their activations from previous time steps to apply them to the current time step (Russin et al., 2024). In this way, large language models (LLM), for example, learn to extract the relevant information from the recent word stream for next-word predictions, resulting in incredible capacities to meet challenges in new tasks. Similarly, activation dynamics in LCD make sure that the model gradually moves toward the right attractor (which can be called in-context adaptation); where the direction and speed of movement have been shaped by in-weight learning.

Our model also has implications for prevalent discussions in the broader literature on cognitive control. For example, many theoretical discussions on task switching have focused on the origin of switch costs, where researchers have dissociated between an interference view and reconfiguration view (Kiesel et al., 2010; Vandierendonck et al., 2010). The interference view holds that the persistent activation of an old and irrelevant task representation (task-set inertia) results in switch costs (Allport et al., 1994; Evans et al., 2015; Logan & Bundesen, 2003; Meuter & Allport, 1999; Philipp et al., 2007; Waszak et al., 2003;

Wylie & Allport, 2000). In contrast, the reconfiguration view proposes that switch costs reflect a process of activating new task representations (reconfiguring) during task switching compared to task repetition (Jost et al., 2008; Meiran, 1996; Monsell & Mizon, 2006; Rogers & Monsell, 1995; Rubinstein et al., 2001; Verbruggen et al., 2007; Weaver et al., 2014). In this terminology, LCD's process of reconfiguring to a new task involves moving away from previous, irrelevant tasks and toward the currently relevant ones in task space. From this perspective, both processes are two sides of the same coin, where task switching involves a constant balancing between different tasks.

The past two decades have also witnessed a debate over whether markers of cognitive control, such as switch costs, reflect cognitive control processes at all, or learning instead (Mayr & Kliegl, 2003; Mordkoff, 2012; Rothermund et al., 2005; Schmidt et al., 2020), often attributing control-related effects to episodic memory (Frings et al., 2020; Giesen & Rothermund, 2014; Hommel, 2004) or contingency learning processes (Hommel et al., 2004; Mayr et al., 2003; Schmidt et al., 2020; Schmidt & Besner, 2008). Here, in line with recent theories that emphasize the learning of control settings (Abrahamse et al., 2016; Botvinick & Cohen, 2014; Braem & Egner, 2018; Lieder et al., 2018; Verguts & Notebaert, 2009), we show how cognitive control involves the learning of such control settings but crucially in tight interaction with episodic memory as implemented via task-based activation. Indeed, LCD even learns abstract environment-control relations in response to changing environment-specific demands for cognitive flexibility versus stability. Therefore, LCD provides a unified framework to integrate not only learning and control, but also the respective contributions of episodic memory and contingency learning (Frings et al., 2023) through activation-based and weight-based changes, respectively.

Another long-standing debate in the field is whether states of cognitive flexibility versus stability should be considered domain-general or domain-specific. While some studies suggest that training task switching may promote performance on other related tasks that require cognitive flexibility (Johann & Karbach, 2020; Karbach & Kray, 2009; Minear & Shah, 2008), others find limited evidence for such transfer (Aben et al., 2019; Ganesan et al., 2024; Kassai et al., 2019; Melby-Lervåg et al., 2016; Sabah et al., 2019, 2021; Simons et al., 2016; Uddin, 2021). Our model takes a clear stance in that it predicts that learned control states are highly task- and context-specific, thereby rendering it unlikely that transfer would extend far beyond the current task (components). This is consistent with recent models and observations that learned states of cognitive flexibility and stability can be very task- and environment-specific (Fröber et al., 2022; Nack & Yu-Chin, 2024; Siqi-Liu & Egner, 2020, 2023; Xu et al., 2024). Namely, we document how different states of cognitive flexibility versus stability can be learned, but, in doing so, will necessarily bind to the task identities of broader environments in which these different attractor landscapes are learned. From this perspective, generalization can only occur for whole task sets or its disentangled components (Frankland & Greene, 2020; Franklin & Frank, 2020; Hupkes et al., 2020; Ito et al., 2022). This model-based interpretation has the benefit of allowing for concrete predictions about the extent to which tasks and task environments can benefit from one another.

LCD also shows clear links to, and takes inspiration from, different biological processes. First, activation-based adaptations have been related to the activation dynamics of neurons (via neural spikes), whereas changes in weights are thought to depend on pre-or post-synaptic membrane configurations (e.g., based on AMPA receptors). Interestingly, some have also proposed an intermediate activity-silent working memory system that maintains information via short-term synaptic plasticity instead of persistent neural activity (Masse et al., 2020; Stokes, 2015). We speculate that weight changes may initially reside in activity-silent working memory, but then gradually transfer to membrane configurations when repeatedly triggered. Second, the gating parameters inherent in any LSTM (including LCD), are plausibly supported by striato-cortical loops systems (Braver & Cohen, 1999, 2000; Chatham & Badre, 2015; Ede & Nobre, 2023; Hazy et al., 2007). Finally, attractors, as described here, may be sculpted in cortex, whereas thalamus integrates inputs from basal ganglia and cerebellum to adapt their depth to construct a dynamically evolving attractor landscape (Houk & Wise, 1995; Shine, 2021).

LCD also holds potential for several extensions. We focused on accuracy, but switch costs are often observed in both accuracy and response time (RT) data. In general, empirical effects can be observed in accuracy, RT, or some combination, depending on where an agent resides on the speed-accuracy tradeoff continuum (Heitz, 2014; Woodworth, 1899). Future work can model RT so that our model can be compared more comprehensively to human data. A second and related extension is fitting the model to behavioral data. Given their complexity, neural networks are notoriously difficult to fit to data, and it is a classic modeling trade-off to make models expressive or instead statistically tractable. Nevertheless, promising new avenues allow for a more thorough statistical treatment of neural networks (Ji-An et al., 2023; Radev et al., 2020). Third, as just mentioned, efficient generalization across tasks may require a compositional (e.g., hierarchical) task structure (Liu & Frank, 2022). LCD did not focus on simulating these hierarchical task structures, which pertain mostly to the first learning step of our model. However, incorporating this aspect of task learning will be an important extension to allow for more fine-grained predictions on the generalization of control strategies.

In conclusion, LCD provides a unified framework to study how people find optimal states of cognitive flexibility versus stability when switching between tasks. We simulated several findings from the literature and proposed a clear computational mechanism of cognitive control. Inspired by this framework, we discussed new insights to longstanding debates regarding control functions in regulating human behavior. Given the ubiquity of task switching in (modern) daily life, we believe that sculpting adaptive attractors as modeled in our LCD model, reflects a crucial aspect of cognitive control.

References

Aben, B., Iseni, B., Van den Bussche, E., & Verguts, T. (2019). Persistent modification of cognitive control through attention training. *Quarterly Journal of Experimental Psychology*, *72*(3), 413–423. https://doi.org/10.1177/1747021818757979

Abrahamse, E., Braem, S., Notebaert, W., & Verguts, T. (2016). Grounding cognitive control in associative learning. *Psychological Bulletin*, *142*(7), 693–728. https://doi.org/10.1037/bul0000047

Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. *Attention and Performance 15: Conscious and Nonconscious Information Processing. Attention and Performance Series*, 421–452.

Botvinick, M. M., & Cohen, J. D. (2014). The Computational and Neural Basis of Cognitive Control: Charted Territory and New Frontiers. *Cognitive Science*, *38*(6), 1249–1285. https://doi.org/10.1111/cogs.12126

Braem, S. (2017). Conditioning task switching behavior. *Cognition*, *166*, 272–276. https://doi.org/10.1016/j.cognition.2017.05.037

Braem, S., Bugg, J. M., Schmidt, J. R., Crump, M. J. C., Weissman, D. H., Notebaert, W., & Egner, T. (2019). Measuring Adaptive Control in Conflict Tasks. *Trends in Cognitive Sciences*, *23*(9), 769–783. https://doi.org/10.1016/j.tics.2019.07.002

Braem, S., Chai, M., Held, L. K., & Xu, S. (2024). One cannot simply 'be flexible’: Regulating control parameters requires learning. *Current Opinion in Behavioral Sciences*, *55*, 101347. https://doi.org/10.1016/j.cobeha.2023.101347

Braem, S., & Egner, T. (2018). Getting a Grip on Cognitive Flexibility. *Current Directions in Psychological Science*, *27*(6), 470–476. https://doi.org/10.1177/0963721418787475

Braver, T. S., & Cohen, J. D. (1999). Chapter 19 Dopamine, cognitive control, and schizophrenia: The gating model. In J. A. Reggia, E. Ruppin, & D. Glanzman (Eds.), *Progress in Brain Research* (Vol. 121, pp. 327–349). Elsevier. https://doi.org/10.1016/S0079-6123(08)63082-4

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. *Attention and Performance*, *18*, 712–737.

Brosowsky, N. P., & Egner, T. (2021). Appealing to the cognitive miser: Using demand avoidance to modulate cognitive flexibility in cued and voluntary task switching. *Journal of Experimental Psychology: Human Perception and Performance*, *47*(10), 1329–1347. https://doi.org/10.1037/xhp0000942

Brown, J. W., Reynolds, J. R., & Braver, T. S. (2007). A computational model of fractionated conflict-control mechanisms in task-switching. *Cognitive Psychology*, *55*(1), 37–85. https://doi.org/10.1016/j.cogpsych.2006.09.005

Chatham, C. H., & Badre, D. (2015). Multiple gates on working memory. *Current Opinion in Behavioral Sciences*, *1*, 23–31. https://doi.org/10.1016/j.cobeha.2014.08.001

Cools, R. (2019). Chemistry of the Adaptive Mind: Lessons from Dopamine. *Neuron*, *104*(1), 113–131. https://doi.org/10.1016/j.neuron.2019.09.035

Crone, E. A., & Steinbeis, N. (2017). Neural Perspectives on Cognitive Control Development during Childhood and Adolescence. *Trends in Cognitive Sciences*, *21*(3), 205–215. https://doi.org/10.1016/j.tics.2017.01.003

Dreisbach, G., & Fröber, K. (2019). On How to Be Flexible (or Not): Modulation of the Stability-Flexibility Balance. *Current Directions in Psychological Science*, *28*(1), 3–9. https://doi.org/10.1177/0963721418800030

Dreisbach, G., & Haider, H. (2006). Preparatory adjustment of cognitive control in the task switching paradigm. *Psychonomic Bulletin & Review*, *13*(2), 334–338. https://doi.org/10.3758/BF03193853

Dreisbach, G., & Mendl, J. (2024). Flexibility as a matter of context, effort, and ability: Evidence from the task-switching paradigm. *Current Opinion in Behavioral Sciences*, *55*, 101348. https://doi.org/10.1016/j.cobeha.2023.101348

Ede, F. van, & Nobre, A. C. (2023). Turning Attention Inside Out: How Working Memory Serves Behavior. *Annual Review of Psychology*, *74*(Volume 74, 2023), 137–165. https://doi.org/10.1146/annurev-psych-021422-041757

Egner, T. (2023). Principles of cognitive control over task focus and task switching. *Nature Reviews Psychology*, 1–13. https://doi.org/10.1038/s44159-023-00234-4

Eppinger, B., Goschke, T., & Musslick, S. (2021). Meta-control: From psychology to computational neuroscience. *Cognitive, Affective, & Behavioral Neuroscience*, *21*(3), 447–452. https://doi.org/10.3758/s13415-021-00919-4

Evans, L. H., Herron, J. E., & Wilding, E. L. (2015). Direct Real-Time Neural Evidence for Task-Set Inertia. *Psychological Science*, *26*(3), 284–290. https://doi.org/10.1177/0956797614561799

Frankland, S. M., & Greene, J. D. (2020). Concepts and Compositionality: In Search of the Brain's Language of Thought. *Annual Review of Psychology*, *71*(Volume 71, 2020), 273–303. https://doi.org/10.1146/annurev-psych-122216-011829

Franklin, N. T., & Frank, M. J. (2020). Generalizing to generalize: Humans flexibly switch between compositional and conjunctive structures during reinforcement learning. *PLOS Computational Biology*, *16*(4), e1007720. https://doi.org/10.1371/journal.pcbi.1007720

Frings, C., Foerster, A., Moeller, B., Pastötter, B., & Pfister, R. (2023). The relation between learning and stimulus-response binding. *Psychological Review*. https://doi.org/10.1037/rev0000449

Frings, C., Hommel, B., Koch, I., Rothermund, K., Dignath, D., Giesen, C., Kiesel, A., Kunde, W., Mayr, S., Moeller, B., Möller, M., Pfister, R., & Philipp, A. (2020). Binding and Retrieval in Action Control (BRAC). *Trends in Cognitive Sciences*, *24*(5), 375–387. https://doi.org/10.1016/j.tics.2020.02.004

Fröber, K., & Dreisbach, G. (2017). Keep flexible – Keep switching! The influence of forced task switching on voluntary task switching. *Cognition*, *162*, 48–53. https://doi.org/10.1016/j.cognition.2017.01.024

Fröber, K., Jurczyk, V., & Dreisbach, G. (2022). Keep flexible—Keep switching? Boundary conditions of the influence of forced task switching on voluntary task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *48*(9), 1249–1262. https://doi.org/10.1037/xlm0001104

Ganesan, K., Thompson, A., Smid, C. R., Cañigueral, R., Li, Y., Revill, G., Puetz, V., Bernhardt, B. C., Dosenbach, N. U. F., Kievit, R., & Steinbeis, N. (2024). Cognitive control training with domain-general response inhibition does not change children's brains or behavior. *Nature Neuroscience*, *27*(7), 1364–1375. https://doi.org/10.1038/s41593-024-01672-w

Garner, K. G., & Dux, P. E. (2023). Knowledge generalization and the costs of multitasking. *Nature Reviews Neuroscience*, *24*(2), Article 2. https://doi.org/10.1038/s41583-022-00653-x

Giesen, C., & Rothermund, K. (2014). You Better Stop! Binding "Stop" Tags to Irrelevant Stimulus Features. *Quarterly Journal of Experimental Psychology*, *67*(4), 809–832. https://doi.org/10.1080/17470218.2013.834372

Gilbert, S. J., & Shallice, T. (2002). Task Switching: A PDP Model. *Cognitive Psychology*, *44*(3), 297–337. https://doi.org/10.1006/cogp.2001.0770

Goschke, T. (2013). Volition in Action: Intentions, Control Dilemmas, and the Dynamic Regulation of Cognitive Control. In W. Prinz, M. Beisert, & A. Herwig (Eds.), *Action Science* (pp. 409–434). The MIT Press. https://doi.org/10.7551/mitpress/9780262018555.003.0024

Goschke, T. (2014). Dysfunctions of decision-making and cognitive control as transdiagnostic mechanisms of mental disorders: Advances, gaps, and needs in current research. *International Journal of Methods in Psychiatric Research*, *23*(S1), 41–57. https://doi.org/10.1002/mpr.1410

Grahek, I., Leng, X., Musslick, S., & Shenhav, A. (2024). *Control adjustment costs limit goal flexibility: Empirical evidence and a computational account* (p. 2023.08.22.554296). bioRxiv. https://doi.org/10.1101/2023.08.22.554296

Haykin, S., Fatemi, M., Setoodeh, P., & Xue, Y. (2012). Cognitive Control. *Proceedings of the IEEE*, *100*(12), 3156–3169. Proceedings of the IEEE. https://doi.org/10.1109/JPROC.2012.2215773

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1601–1613. https://doi.org/10.1098/rstb.2007.2055

Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience*, *8*. https://doi.org/10.3389/fnins.2014.00150

Held, L. K., Vermeylen, L., Dignath, D., Notebaert, W., Krebs, R. M., & Braem, S. (2024). Reinforcement learning of adaptive control strategies. *Communications Psychology*, *2*(1), 1–13. https://doi.org/10.1038/s44271-024-00055-y

Held, L. K., Vermeylen, L., Krebs, R. M., Notebaert, W., & Braem, S. (2024). The Effects of Reinforcing Task Alternations on Voluntary Task Selection. *Collabra: Psychology*, *10*(1), 120329. https://doi.org/10.1525/collabra.120329

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780. Neural Computation. https://doi.org/10.1162/neco.1997.9.8.1735

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences*, *8*(11), 494–500. https://doi.org/10.1016/j.tics.2004.08.007

Hommel, B. (2015). Between Persistence and Flexibility. In *Advances in Motivation Science* (Vol. 2, pp. 33–67). Elsevier. https://doi.org/10.1016/bs.adms.2015.04.003

Hommel, B., & Colzato, L. S. (2017). The social transmission of metacontrol policies: Mechanisms underlying the interpersonal transfer of persistence and flexibility. *Neuroscience & Biobehavioral Reviews*, *81*, 43–58. https://doi.org/10.1016/j.neubiorev.2017.01.009

Hommel, B., Proctor, R. W., & Vu, K.-P. L. (2004). A feature-integration account of sequential effects in the Simon task. *Psychological Research*, *68*(1), 1–17. https://doi.org/10.1007/s00426-003-0132-y

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558. https://doi.org/10.1073/pnas.79.8.2554

Houk, J. C., & Wise, S. P. (1995). Feature Article: Distributed Modular Architectures Linking Basal Ganglia, Cerebellum, and Cerebral Cortex: Their Role in Planning and Controlling Action. *Cerebral Cortex*, *5*(2), 95–110. https://doi.org/10.1093/cercor/5.2.95

Hummos, A., Wang, B. A., Drammis, S., Halassa, M. M., & Pleger, B. (2022). Thalamic regulation of frontal interactions in human cognitive flexibility. *PLOS Computational Biology*, *18*(9), e1010500. https://doi.org/10.1371/journal.pcbi.1010500

Hupkes, D., Dankers, V., Mul, M., & Bruni, E. (2020). Compositionality Decomposed: How do Neural Networks Generalise? *Journal of Artificial Intelligence Research*, *67*, 757–795. https://doi.org/10.1613/jair.1.11674

Ito, T., Klinger, T., Schultz, D., Murray, J., Cole, M., & Rigotti, M. (2022). Compositional generalization through abstract representations in human and artificial neural networks. *Advances in Neural Information Processing Systems*, *35*, 32225–32239.

Jaffe, P. I., Poldrack, R. A., Schafer, R. J., & Bissett, P. G. (2023). Modelling human behaviour in cognitive tasks with latent dynamical systems. *Nature Human Behaviour*, 1–15. https://doi.org/10.1038/s41562-022-01510-8

Jia, S., Liu, D., Song, W., Beste, C., Colzato, L., & Hommel, B. (2024). Tracing conflict-induced cognitive-control adjustments over time using aperiodic EEG activity. *Cerebral Cortex*, *34*(5), bhae185. https://doi.org/10.1093/cercor/bhae185

Ji-An, L., Benna, M. K., & Mattar, M. G. (2023). *Automatic Discovery of Cognitive Strategies with Tiny Recurrent Neural Networks* (p. 2023.04.12.536629). bioRxiv. https://doi.org/10.1101/2023.04.12.536629

Johann, V. E., & Karbach, J. (2020). Effects of game-based and standard executive control training on cognitive and academic abilities in elementary school children. *Developmental Science*, *23*(4), e12866. https://doi.org/10.1111/desc.12866

Jost, K., Mayr, U., & Rösler, F. (2008). Is task switching nothing but cue priming? Evidence from ERPs. *Cognitive, Affective, & Behavioral Neuroscience*, *8*(1), 74–84. https://doi.org/10.3758/CABN.8.1.74

Karbach, J., & Kray, J. (2009). How useful is executive control training? Age differences in near and far transfer of task-switching training. *Developmental Science*, *12*(6), 978–990. https://doi.org/10.1111/j.1467-7687.2009.00846.x

Kassai, R., Futo, J., Demetrovics, Z., & Takacs, Z. K. (2019). A meta-analysis of the experimental evidence on the near- and far-transfer effects among children's executive function skills. *Psychological Bulletin*, *145*(2), 165–188. https://doi.org/10.1037/bul0000180

Khona, M., & Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, *23*(12), 744–766. https://doi.org/10.1038/s41583-022-00642-0

Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching—A review. *Psychological Bulletin*, *136*(5), 849–874. https://doi.org/10.1037/a0019842

Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLOS Computational Biology*, *14*(4), e1006043. https://doi.org/10.1371/journal.pcbi.1006043

Liu, R. G., & Frank, M. J. (2022). Hierarchical clustering optimizes the tradeoff between compositionality and expressivity of task structures for flexible reinforcement learning. *Artificial Intelligence*, *312*, 103770. https://doi.org/10.1016/j.artint.2022.103770

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, *29*(3), 575–599. https://doi.org/10.1037/0096-1523.29.3.575

Masse, N. Y., Rosen, M. C., & Freedman, D. J. (2020). Reevaluating the Role of Persistent Neural Activity in Short-Term Memory. *Trends in Cognitive Sciences*, *24*(3), 242–258. https://doi.org/10.1016/j.tics.2019.12.014

Mayr, U., Awh, E., & Laurey, P. (2003). Conflict adaptation effects in the absence of executive control. *Nature Neuroscience*, *6*(5), 450–452. https://doi.org/10.1038/nn1051

Mayr, U., & Kliegl, R. (2003). Differential effects of cue changes and task changes on task-set selection costs. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(3), 362–372. https://doi.org/10.1037/0278-7393.29.3.362

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1423–1442. https://doi.org/10.1037/0278-7393.22.6.1423

Melby-Lervåg, M., Redick, T. S., & Hulme, C. (2016). Working Memory Training Does Not Improve Performance on Measures of Intelligence or Other Measures of "Far Transfer": Evidence From a Meta-Analytic Review. *Perspectives on Psychological Science*, *11*(4), 512–534. https://doi.org/10.1177/1745691616635612

Mendl, J., & Dreisbach, G. (2024). Is task switching avoided to save effort or time? Shorter intertrial durations following task switches increase the willingness to switch tasks. *Motivation Science*. https://doi.org/10.1037/mot0000347

Meuter, R. F. I., & Allport, A. (1999). Bilingual Language Switching in Naming: Asymmetrical Costs of Language Selection. *Journal of Memory and Language*, *40*(1), 25–40. https://doi.org/10.1006/jmla.1998.2602

Minear, M., & Shah, P. (2008). Training and transfer effects in task switching. *Memory & Cognition*, *36*(8), 1470–1483. https://doi.org/10.3758/MC.336.8.1470

Mittelstädt, V., Miller, J., & Kiesel, A. (2018). Trading off switch costs and stimulus availability benefits: An investigation of voluntary task-switching behavior in a predictable dynamic multitasking environment. *Memory & Cognition*, *46*(5), 699–715. https://doi.org/10.3758/s13421-018-0802-z

Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology. Human Perception and Performance*, *32*(3), 493–516. https://doi.org/10.1037/0096-1523.32.3.493

Mordkoff, J. T. (2012). Observation: Three reasons to avoid having half of the trials be congruent in a four-alternative forced-choice experiment on sequential modulation. *Psychonomic Bulletin & Review*, *19*(4), 750–757. https://doi.org/10.3758/s13423-012-0257-3

Musslick, S., & Bizyaeva, A. (2024). Examining cognitive flexibility and stability through the lens of dynamical systems. *Current Opinion in Behavioral Sciences*, *57*, 101375. https://doi.org/10.1016/j.cobeha.2024.101375

Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, *25*(9), 757–775. https://doi.org/10.1016/j.tics.2021.06.001

Nack, C. A., & Yu-Chin, C. (2024). Learned switch readiness via concurrent activation of task sets: Evidence from task specificity and memory load. *Memory & Cognition*. https://doi.org/10.3758/s13421-024-01560-5

O’Reilly, R. C., & Frank, M. J. (2006). Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and Basal Ganglia. *Neural Computation*, *18*(2), 283–328. https://doi.org/10.1162/089976606775093909

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain* (pp. xxvii, 504). The MIT Press.

Philipp, A. M., Gade, M., & Koch, I. (2007). Inhibitory processes in language switching: Evidence from switching language-defined response sets. *European Journal of Cognitive Psychology*, *19*(3), 395–416. https://doi.org/10.1080/09541440600758812

Radev, S. T., Mertens, U. K., Voss, A., & Köthe, U. (2020). Towards end-to-end likelihood-free inference with convolutional neural networks. *British Journal of Mathematical and Statistical Psychology*, *73*(1), 23–43. https://doi.org/10.1111/bmsp.12159

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*(2), 207–231. https://doi.org/10.1037/0096-3445.124.2.207

Rolls, E. T., Loh, M., Deco, G., & Winterer, G. (2008). Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. *Nature Reviews Neuroscience*, *9*(9), 696–709. https://doi.org/10.1038/nrn2462

Rothermund, K., Wentura, D., & De Houwer, J. (2005). Retrieval of Incidental Stimulus-Response Associations as a Source of Negative Priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(3), 482–495. https://doi.org/10.1037/0278-7393.31.3.482

Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(4), 763–797. https://doi.org/10.1037/0096-1523.27.4.763

Ruel, A., Devine, S., & Eppinger, B. (2021). Resource-rational approach to meta-control problems across the lifespan. *WIREs Cognitive Science*, *12*(5), e1556. https://doi.org/10.1002/wcs.1556

Russin, J., McGrath, S. W., Williams, D. J., & Elber-Dorozko, L. (2024). *From Frege to chatGPT: Compositionality in language, cognition, and deep neural networks* (arXiv:2405.15164). arXiv. https://doi.org/10.48550/arXiv.2405.15164

Sabah, K., Dolk, T., Meiran, N., & Dreisbach, G. (2019). When less is more: Costs and benefits of varied vs. fixed content and structure in short-term task switching training. *Psychological Research*, *83*(7), 1531–1542. https://doi.org/10.1007/s00426-018-1006-7

Sabah, K., Dolk, T., Meiran, N., & Dreisbach, G. (2021). Enhancing task-demands disrupts learning but enhances transfer gains in short-term task-switching training. *Psychological Research*, *85*(4), 1473–1487. https://doi.org/10.1007/s00426-020-01335-y

Schmidt, J. R., & Besner, D. (2008). The Stroop effect: Why proportion congruent has nothing to do with congruency and everything to do with contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(3), 514–523. https://doi.org/10.1037/0278-7393.34.3.514

Schmidt, J. R., Liefooghe, B., & Houwer, J. D. (2020). An Episodic Model of Task Switching Effects: Erasing the Homunculus from Memory. *Journal of Cognition*, *3*(1), Article 1. https://doi.org/10.5334/joc.97

Shine, J. M. (2021). The thalamus integrates the macrosystems of the brain to facilitate complex, adaptive brain network dynamics. *Progress in Neurobiology*, *199*, 101951. https://doi.org/10.1016/j.pneurobio.2020.101951

Simons, D. J., Boot, W. R., Charness, N., Gathercole, S. E., Chabris, C. F., Hambrick, D. Z., & Stine-Morrow, E. A. L. (2016). Do "Brain-Training" Programs Work? *Psychological Science in the Public Interest*, *17*(3), 103–186. https://doi.org/10.1177/1529100616661983

Siqi-Liu, A., & Egner, T. (2020). Contextual Adaptation of Cognitive Flexibility is driven by Task- and Item-Level Learning. *Cognitive, Affective, & Behavioral Neuroscience*, *20*(4), 757–782. https://doi.org/10.3758/s13415-020-00801-9

Siqi-Liu, A., & Egner, T. (2023). Task sets define boundaries of learned cognitive flexibility in list-wide proportion switch manipulations. *Journal of Experimental Psychology. Human Perception and Performance*, *49*(8), 1111–1122. https://doi.org/10.1037/xhp0001138

Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. https://doi.org/10.1016/j.tics.2015.05.004

Stucke, N., & Doebel, S. (2023). *Early Childhood Executive Function Predicts Concurrent and Later Social and Behavioral Outcomes: A Review and Meta-Analysis*. OSF. https://doi.org/10.31234/osf.io/s59ev

Uddin, L. Q. (2021). Cognitive and behavioural flexibility: Neural mechanisms and clinical considerations. *Nature Reviews Neuroscience*, *22*(3), Article 3. https://doi.org/10.1038/s41583-021-00428-w

Ueltzhöffer, K., Armbruster-Genç, D. J. N., & Fiebach, C. J. (2015). Stochastic Dynamics Underlying Cognitive Stability and Flexibility. *PLOS Computational Biology*, *11*(6), e1004331. https://doi.org/10.1371/journal.pcbi.1004331

Vandierendonck, A., Liefooghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin*, *136*(4), 601–626. https://doi.org/10.1037/a0019791

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. ukasz, & Polosukhin, I. (2017). Attention is All you Need. *Advances in Neural Information Processing Systems*, *30*. https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

Verbruggen, F., Liefooghe, B., Vandierendonck, A., & Demanet, J. (2007). Short cue presentations encourage advance task preparation: A recipe to diminish the residual switch cost. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(2), 342–356. https://doi.org/10.1037/0278-7393.33.2.342

Verguts, T., & Notebaert, W. (2009). Adaptation by binding: A learning account of cognitive control. *Trends in Cognitive Sciences*, *13*(6), 252–257. https://doi.org/10.1016/j.tics.2009.02.007

Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences*, *38*, 90–95. https://doi.org/10.1016/j.cobeha.2021.01.002

Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., & Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, *21*(6), 860–868. https://doi.org/10.1038/s41593-018-0147-8

Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus–task bindings in task-shift costs. *Cognitive Psychology*, *46*(4), 361–413. https://doi.org/10.1016/S0010-0285(02)00520-0

Weaver, S. M., Foxe, J. J., Shpaner, M., & Wylie, G. R. (2014). You Can't Always Get What you Want: The Influence of Unexpected Task Constraint on Voluntary Task Switching. *Quarterly Journal of Experimental Psychology*, *67*(11), 2247–2259. https://doi.org/10.1080/17470218.2014.917115

Woodworth, R. S. (1899). Accuracy of voluntary movement. *The Psychological Review: Monograph Supplements*, *3*(3), i–114. https://doi.org/10.1037/h0092992

Wylie, G., & Allport, A. (2000). Task switching and the measurement of "switch costs". *Psychological Research*, *63*(3), 212–233. https://doi.org/10.1007/s004269900003

Xu, S., Simoens, J., Verguts, T., & Braem, S. (2024). Learning where to be flexible: Using environmental cues to regulate cognitive control. *Journal of Experimental Psychology: General*, *153*(2), 328–338. https://doi.org/10.1037/xge0001488

Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., & Wang, X.-J. (2019). Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, *22*(2), 297–306. https://doi.org/10.1038/s41593-018-0310-2