Learning and transfer of working memory gating policies

Apoorva Bhandari[a], David Badre[a, b]

[a]Cognitive, Linguistics, & Psychological Sciences, Brown University, Providence, USA

[b]Brown Institute for Brain Science, Brown University, Providence, USA

Corresponding author:

Apoorva Bhandari,

Cognitive, Linguistics, & Psychological Science,

Brown University, Providence, USA.

Phone: 401-863-5197

E-mail: apoorva_bhandari@brown.edu

Abstract

A central question in cognitive neuroscience concerns the knowledge structures and neural representations that support flexible behavior and rapid generalization in novel environments. Previous research has focused almost exclusively on abstract rules that leverage shared structure in stimulus-response-outcome (S-R-O) mappings between tasks as the basis of such task knowledge. Here we provide evidence that working memory (WM) gating policies, required for internal control of WM during a task, leverage a form of task structure separate from S-R-O relationships and are re-used in novel situations. In two experiments, we report specific evidence for the transfer of selective WM gating policies across changes in task context. We show that this transfer is not tied to S-R-O rules, suggesting that gating policies may be a separate component of task knowledge. Collectively, our results highlight the importance of WM gating policies as a key component of the task knowledge that supports flexible behavior and task generalization.

Introduction

Humans display remarkable cognitive flexibility in novel task environments (McClelland, 2009). Faced with a new task, our initial performance is slow and error-prone, but improves rapidly, often achieving asymptotic levels within just a few trials (Ackerman, 1988; Bhandari & Duncan, 2014; Ruge & Wolfensteller, 2010; Wolfensteller & Ruge, 2011). Such rapid learning relies, in part, on abstract task knowledge transferred from prior experience with other tasks. To the extent that such knowledge is relevant to the new task, it constrains the space of possibilities that must be explored, thus speeding up learning (Botvinick, Niv, & Barto, 2009; Cole, Etzel, Zacks, Schneider, & Braver, 2011; Collins & Frank, 2013; Gershman & Niv, 2010).

What form does such abstract task knowledge take? The vast majority of prior studies seeking to address this question have focused on rules, task-sets, or stimulus-response (S-R-O) mappings as the basis of task knowledge. In these frameworks, abstract rules allow us to generalize prior knowledge and thus constrain the (usually) very large space of stimulus-response-outcome contingencies afforded by a novel task environment (Badre, Kayser, & D'Esposito, 2010). Such rules can both be instructed (Cohen-Kdoshay & Meiran, 2007, 2009; Cole, Bagic, Kass, & Schneider, 2010; Meiran, Pereg, Kessler, Cole, & Braver, 2015; Ruge & Wolfensteller, 2010) or transferred from prior experiences (Cole et al., 2011; Collins & Frank, 2013) to rapidly enable successful behavior in novel environments.

The implementation of a task, however, requires more than just the knowledge of stimulus-response-outcome contingencies. Even the simplest, everyday task environments have dynamical structure, with events unfolding in time (Radvansky &

Zacks, 2014). In a dynamic task environment, then, one must also learn an internal control policy or 'task model' for the moment-by-moment control of internal cognitive processing that is aligned with the structure of task dynamics (Bhandari & Duncan, 2014; Duncan et al., 2008). Such implementational control policies are not typically communicated via instruction and must be discovered and implemented "on the fly", through task experience.

In this paper, we ask whether control policies are themselves a form of knowledge that, like rules, can be transferred to novel environments. Just like different real-world tasks often share stimulus-response-outcome contingencies, they also share other forms dynamic structure (Botvinick, Weinstein, Solway, & Barto, 2015; Schank & Abelson, 1977). Such shared structure affords an opportunity for generalization of internal control policies. Instead of learning new control policies from scratch, humans may build repertoires of internal control policies that are re-used in novel tasks.

We operationalize this question within the domain of working memory (WM) control. WM control has been extensively analyzed within the *gating framework,* in which access to WM is controlled by a set of input and output gates (Chatham & Badre, 2015; O'Reilly & Frank, 2006; Todd, Niv, & Cohen, 2009). The contents of WM can be selectively updated by operating an input gate that determines whether stimulus information can enter WM. Similarly, operating a selective output gate allows WM to selectively influence downstream. Learning to perform a task, therefore, involves learning a *gating policy* for operating input and output gates in a moment-by-moment, task-appropriate manner (Frank & Badre, 2012). In the context of WM, a gating policy can be thought of as a type of control policy.

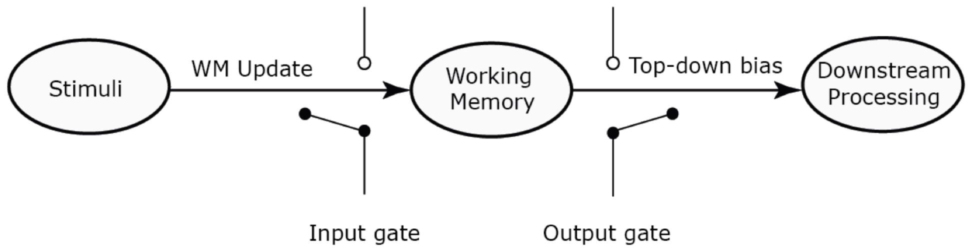

Figure 1: Simple model of working memory control within the gating framework. Access to WM is controlled via the operation of an input gate that determines whether a stimulus is updated into WM. On the other side, an output gate controls whether or not information within WM can influence on behavior. Two broad classes of policies can be distinguished for selective use of WM – ones that achieve control via selective input-gating, and ones that rely on selection via output-gating. Adapted from Hazy, Frank, and O'Reilly R (2007).

We adopt the 2nd order WM control task employed by Chatham, Frank, and Badre (2014). In their task, participants saw a sequence of three items on every trial, one of which specified a context. The context signaled which of the other two items in the sequence was the target item. Critically, there were two kinds of task structures – 'context first' (CF) trials, on which the first item in the sequence was the context item, and 'context last' (CL) trials, in which the last item in the sequence was the context item. CF and CL trials afford the use of different WM gating policies. On a CL trial, subjects had to employ a 'selective output-gating policy' that allowed the storage of both lower level items in WM (a non-selective input-gating operation), and the retrieval of the target item for guiding response selection (a selective output-gating operation). On a CF trial, while a similar selective output-gating policy could be employed, a more efficient 'selective input-gating policy' was possible. Such a policy would enable proactive coding of the contextual cue in WM, followed by selective input-gating of only the relevant lower-level item contingent on context. This allows a reduction in both, WM load, and interference from the competing non-target during response selection. Indeed, Chatham et al. (2014) presented evidence that CF and CL trials are treated differently and that well-trained subjects employ selective input-gating policies on CF trials to improve performance relative to CL trials on which the selective output-gating policy is required.

In the context of this WM control task, we ask whether selective gating policies learned in one task environment are transferred to a novel task environment. For instance, subjects exposed to an environment with only CL trials would learn a selective output-gating policy. Would this policy transfer to a new block with CF trials? In Experiment 1 we demonstrate a pattern of transfer effects that support the hypothesis that a previously learned gating policy influences initial behavior in a novel setting. We replicate these findings in Experiment 2. In addition, we provide evidence that transferred gating policies are dissociable from task rules and have a much larger influence on subsequent behavior. We interpret these findings as evidence that internal control policies comprise an important form of structural task knowledge that supports behavior in novel situations.

Experiment 1

Methods

Participants.

85 adult, right-handed participants (34 males, 51 females; age-range: 18-30, M = 21.4, SD = 2.7) from the Providence, RI area were recruited to take part in a computer-based behavioral experiment. We endeavored to collect between 18-20 participants in each of four groups based on approximate effect sizes suggested by pilot data. 1

participant was excluded for prior neurological injuries, 3 were excluded as they were on psychoactive medication. 5 participants were excluded because of low performance (<70% accuracy) on the task. This left 76 participants (30 males, 46 females; age-range: 18-29, M = 21.3, SD = 2.6). Participants were randomly assigned to four groups of 19 each and there were no group differences in age or gender ratio ($p > 0.250$). We subsequently recruited another 19 participants as an additional control group (8 males, 11 females; age-range: 18-30, M = 21, SD = 4.1). All participants had normal or corrected-to-normal vision, and no reported neurological or psychological disorders. All participants gave informed, written consent as approved by the Human Research Protections Office of Brown University, and they were compensated for their participation.

Apparatus.

Experiments were conducted on a computer running the Mac OSX operating system. The stimulus delivery program was written in MATLAB 2013b, using the Psychophysics Toolbox. Responses were collected via a standard keyboard. All analyses were carried out in MATLAB 2013b and SPSS 22.

Task and Experiment Design.

Participants were instructed to perform a 2nd order working memory control task (Figure 2). On each trial, they saw a sequence of three items on the computer screen: a number (11 or 53), a letter (A or G), and a symbol (π or ⦿). The number served as a higher-level contextual cue, which specified the lower level items (letter or symbol) that would be the target on each trial. The relationships between the contextual cues and the lower level items were specified by the rule trees shown in Figure 2A. Participants were asked to memorize these rule trees before the task began, and were given an opportunity to review the trees for as long as they wanted at the beginning of each block.

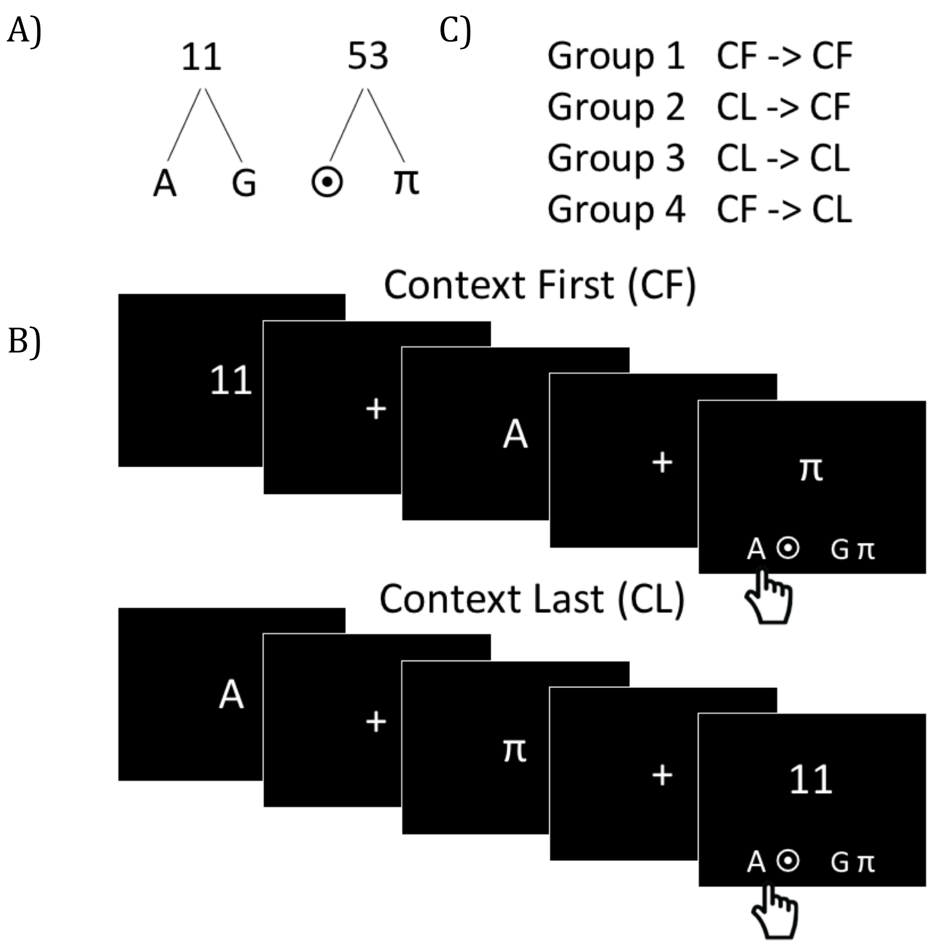

Figure 2: Experiment 1 task and experiment design. (A) Rule trees linking contextual items to lower-level items. Participants memorized the rule trees prior to beginning the task and were given an opportunity to review at the beginning of each block. (B) Sample context first (CF) and context last (CL) trials. The contextual item (number) indicated which of the other two items in the sequence is the target item. Each trial concluded with a response panel that consisted of two pairs of items. Participants indicated with a key press which pair contained the correct target item (denoted by a hand with a tick mark). (C) Participants were assigned to one of four experimental groups. Each group completed two blocks ('training' and 'test'), whose trial structure was determined by group membership.

To illustrate a trial, consider the sequence 11...A…π (Figure 2B). As per the rule trees, 11 indicates that either of the letters, A or G, would be a target on that trial. Therefore, in the above example sequence, A, and not π, is the target. Each trial concluded with a response panel (presented at the same time as the last item in the sequence) that consisted of two pairs of items at the lower left or right of the screen.

Participants had to indicate whether the target item for that trial appeared as part of the left pair (left key) or the right pair (right key). In the example, the left key would be the correct response as it contained the 'A' target. The location of the letters and symbols in the target panel was randomized trial-to-trial such that all items appeared with equal frequency on the left and right across all trials. Further, 50% of the trials were 'incongruent' in that the lower-level items which appeared in the sequence were associated with different response keys in the response panel, and participants could not perform accurately without attending to the contextual cue. Participants were instructed to respond as fast as possible, while being accurate. Response panels were randomized so that on half the trials the left key was the correct response.

Apart from the position of the contextual item in the sequence (either first or last), the position of the lower level items was randomized and balanced between first, last, or middle position in the three item sequence, such that there were equal numbers of each possible trial type in the block. All stimuli were presented in white on a black background. The first two items in the sequence were each presented for 300 ms, while the final item (along with the response panel) was self-timed with response required within a window of 3000 ms. The inter-stimulus interval was jittered between 600 ms and

1600 ms, while the inter-trial interval was fixed at 500 ms. A fixation cross was displayed during each interval.

Two different trial structures were employed (Figure 2B). In the 'context first' (CF) structure, the contextual cue (the number) appeared at the beginning of the sequence. In the 'context last' (CL) structure, the context cue appeared at the end of the sequence. CF and CL trial structures afforded the use of different WM gating policies: CL trials required a 'selective output-gating policy', whereas efficient behavior in CF trials could be supported by a 'selective input-gating policy'. Instructions were worded to be neutral with respect to trial-structure and examples of both trial structures were provided to participants during instruction (see Appendix A). They were also explicitly told that they might see both kinds of trial structures during the experiment. Participants were not given any prior practice on either trial structure.

Participants were randomly assigned to one of four groups. Each group carried out two, forty-eight trial blocks of the task: a 'training' block and a 'test' block (Figure 2C). Within each block, the trial structure (CF or CL) was always fully predictable, providing an opportunity to learn an efficient gating policy. For two groups, the trial structure remained the same between blocks (CF -> CF or CL -> CL). For the other two groups, the trial structure switched (CL -> CF or CF -> CL). Importantly, the task rules (Figure 2A) always remained the same across blocks. Therefore, the design allowed us to dissociate transfer effects from any effects of rule learning.

We subsequently ran an additional group of participants (CF2 -> CF) to control for the effects of WM load in the CL -> CF group. In the training block of the CL -> CF group, two items must be maintained in WM memory on each trial, making the task more difficult. Therefore, any negative transfer effects observed may be due to the fatigue or motivation loss induced by the higher load or the difficulty of the task rather than the change in trial structure. To control for this, in the CF2 -> CF group, the training task consisted of rules in which each contextual item was associated context with four lower-level items. On each trial, context was presented first, followed by two pairs of items, one of which was relevant. Participants were required to keep both relevant lower-level items in WM, as either of them could be the target item. Therefore, the CF2 -> CF served as a control with equivalent WM load during training as the CL -> CF group, but no change in trial structure.

Data Analysis

Analyses utilized response accuracy and response time (RT) measures. All RT analyses focused only on trials where participants responded correctly. In addition, we discarded trials in which participants responded in less than 250 ms. RTs were log-transformed prior to statistical analyses. For analyses of transfer effects, we computed mean RT and response accuracy separately for eight trial bins of six trials each, motivated by our prediction that transfer would most affect the initial trials in a task.

Table 1: Experiment 1 & 2 summary of performance data[a]

Trial structure	*Accuracy (% correct)*	*Response time (ms)*
	Experiment 1	
CF	94.4 ± 0.08	921 ± 245
CL	83.8 ± 0.10	1269 ± 233
Overall	89.1 ± 0.09	1095 ± 239
	Experiment 2	
CF	92.2 ± 0.09	881 ± 294
CL	90.6 ± 0.08	1250 ± 236
Overall	91.4 ± 0.08	1066 ± 267

[a]Summary statistics are means and standard deviations.

Results

Overall, participants performed well at the task (Table 1). As seen in previous studies (Chatham et al., 2014), CF was performed more accurately (94.4% vs. 83.8%) and efficiently (921 ms vs. 1269 ms) than CL (*ps* < 0.001), providing evidence that participants take advantage of the trial structure to selectively input gate lower level items. Note that this was despite the fact that the context stimulus was displayed for longer in the CL condition than in the CF condition. Summary statistics of performance for each group and each block are shown in Table 2.

Transfer effects.

Our primary question concerns the influence of previously learned gating policies on later performance. To this end, we focus on transfer effects from prior experience with a trial structure on performance with the same or new trial structure in subsequent blocks.

Table 2: Experiment 1 & 2 group-wise summary of performance data[a]

Group	*Training block* Accuracy (% correct)	*Training block* Response time (ms)	*Test block* Accuracy (% correct)	*Test block* Response time (ms)
		Experiment 1		
CF ➔ CF	92.0 ± 0.12	971 ± 262	97.4 ± 0.03	907 ± 209
CL ➔ CF	91.6 ± 0.08	1304 ± 228	95.3 ± 0.07	910 ± 213
CL ➔ CL	88.3 ± 0.14	1317 ± 204	94.5 ± 0.05	1266 ± 185
CF ➔ CL	92.7 ± 0.08	896 ± 263	89.7 ± 0.08	1188 ± 279
CL2 ➔ CF	92.6 ± 0.06	1182 ± 228	94.5 ± 0.09	806 ± 204
		Experiment 2		
		Same rules for each block		
CF ➔ CF	94.2 ± 0.07	905 ± 251	96.1 ± 0.09	839 ± 211
CL ➔ CF	83.8 ± 0.12	1299 ± 293	90.1 ± 0.11	883 ± 351
CL ➔ CL	94.1 ± 0.05	1175 ± 166	95.0 ± 0.03	1208 ± 154
CF ➔ CL	87.2 ± 0.05	928 ± 394	90.5 ± 0.07	1367 ± 262
		Different rules for each block		
CF ➔ CF	95.4 ± 0.04	820 ± 149	95.0 ± 0.05	806 ± 140
CL ➔ CF	89.8 ± 0.10	1202 ± 238	87.2 ± 0.15	983 ± 377

[a]Summary statistics are means and standard deviations.

To examine negative transfer effects, we compared performance data from *test blocks* when the trial structure changed, CL -> CF and CF -> CL, with the corresponding *training blocks* with the same trial structure. Log-transformed RT data from the blocks where the structure changed were analysed using two separate mixed-model ANOVAs (for CF and CL trial structures) with trial bin (1-8) as a within-subject factor and training history as a between-subject factor. Note that the training history factor has two levels – ‘no prior training’ and ‘prior training with different structure’. This reflects a comparison between the training block in the no-change group (CF -> CF, for example) and the test block in the structure-change group (CL -> CF, for example)).

For the CF condition (Figure 3A), the ANOVA revealed a main effect of trial bin [F (7, 252) = 7.82, $p < 0.001$, $\eta_p^2 = 0.178$] and a trial bin x training history interaction [F (7, 252) = 3.15, $p = 0.003$, $\eta_p^2 = 0.081$], with participants in the CL -> CF group showing a trend to being slower only in the 1st trial bin in a *t*-test [t (18) = 2.75, $p = 0.01$, $d = 0.89$] (Figure 3A). It is possible that the effect of prior training with the CL trial structure on CF response times was driven by differences in working memory load (CL train structures produce a higher WM load) or task difficulty during training. To rule out this possibility we considered performance of the CF2 -> CF group. In the CF2 training block, context first was always presented first, but participants were required to maintain two lower-level items in WM rather than just one (see methods). Therefore, the task produced an equivalent WM load during training as the CL -> CF group, but no change in trial structure. Indeed, response times in the CF2 training block were significantly lower than in a CF training block [two-sample *t*-test: $t(36) = 3.69$, $p < 0.001$, $d =$], but were barely faster than a CL training block [two-sample *t*-test: $t(36) = -2.0$, $p = 0.054$, d =]. To test whether the negative transfer effect observed in the CL -> CF group could be accounted for by WM load or difficulty, we ran another mixed-model ANOVA comparing the test blocks of CF2 -> CF group and the CL -> CF group. The .

For the CL condition (Figure 3B), the main effect of trial bin was significant [F (7, 252) = 3.13, $p = 0.003$, $\eta_p^2 = 0.08$] reflecting a quadratic speeding or RT over trial bin [F (1, 36) =14.02, $p = 0.001$, $\eta_p^2 = 0.280$], but the trial bin x training history interaction was non-significant [F (7, 252) = 0.94, $p > 0.250$, $\eta_p^2 = 0.026$]. Similar analyses of accuracy data for the CF and CL conditions did not reveal any significant effects.

To examine positive transfer effects, we compared the training and test blocks RTs from participants in the CF -> CF and CL -> CL groups using the same trial bin x training history ANOVA. Data from the CF -> CF group were analysed with a repeated measures ANOVA with trial bin (8 levels: 1-8) and block (2 levels: training or test) as within subject factors. The main effect of trial bin was non-significant [F (7, 126) = 1.04, $p > 0.250$, $\eta_p^2 = 0.055$], while the main effect of block was marginal [F (1, 18) = 4.55, p = 0.047, $\eta_p^2 = 0.202$]. On the other hand, the trial bin x block interaction was significant [F (7, 126) = 3.64, p = 0.001, $\eta_p^2 = 0.168$]. Bonferroni-corrected (α = 0.006), post-hoc paired t-tests revealed that this interaction was driven by faster RTs in the first trial bin in the test block compared to the training block [t (18) = 3.31, p = 0.004, d = 0.76]. A similar analysis of the CL -> CL group data revealed a main effect of trial bin [F (7, 126) = 2.62, p = 0.015, $\eta_p^2 = 0.127$], while the main effect of block [F (7, 126) = 1.76, p = 0.100, $\eta_p^2 = 0.089$], and the trial bin x block interaction were non-significant [F (7, 126) = 1.04, $p > 0.250$, $\eta_p^2 = 0.054$].

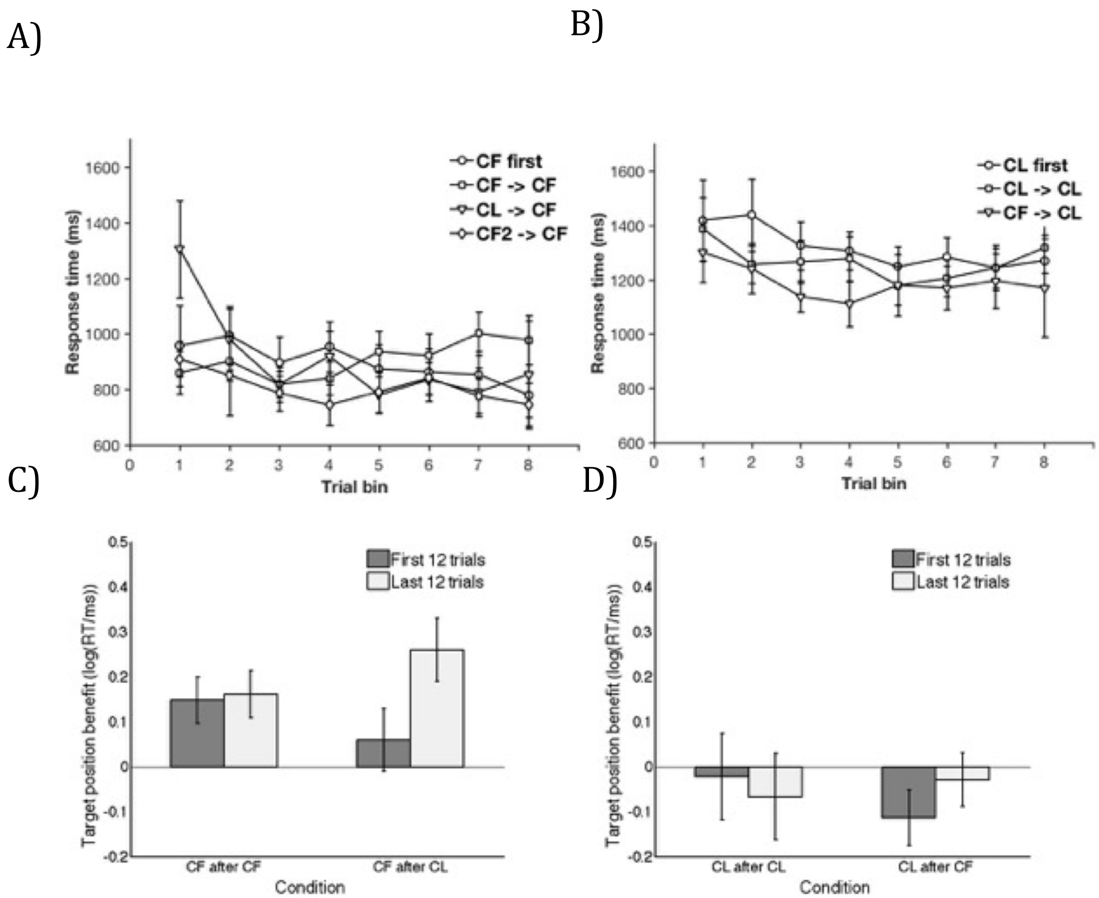

Figure 3: (A) and (B) - Transfer effects for CF (A) and CL (B) trial structures. Mean RT for each 6-trial bin is plotted as a function of training history - no prior training (circles), prior training with same structure (squares), or different trial structure (triangles). Error bars reflect 95% confidence intervals. (C) and (D) - Target position effects for CF (C) and CL (D) trial structures. Mean target position benefit in RT (target-middle RT – target-last RT) plotted for early (first 12 trials – dark grey bars) and late (last 12 trials – light grey bars) in test blocks of CF -> CF and CL -> CF conditions (C) and CL -> CL and CF -> CL conditions (D). RTs were log-transformed. Error bars reflect 95% confidence intervals.

ANOVA of the accuracy data in the CF -> CF group showed a marginally non-significant main effect of block [F (1, 18) = 4.56, p = 0.056, η_p^2 = 0.188]. All other effects were non-significant [$p > 0.1$]. For the CL -> CL group, again, the main effect of block was marginally non-significant [F (1, 18) = 4.3, p = 0.053, η_p^2 = 0.193] and the main effect of trial bin was significant [F (7, 126) = 3.27, p = 0.003, η_p^2 = 0.54]. The interaction was non-significant [F (7, 126) = 1.83, p = 0.088, η_p^2 = 0.09].

In summary, we found evidence of asymmetric transfer effects. Participants with prior experience with the CL trial structure responded more slowly on early trials of a subsequent CF block compared to those who had no prior experience (negative transfer). Participants with prior experience with the CF trial structure were also faster on the early trials of a subsequent CF block (positive transfer). Participants in both the CL->CL and CF-CF groups were also marginally more accurate in test block compared to the training block. On the other hand, prior training with the CF trial structure had no effect on performance in a subsequent CL block.

Target position effects.

In the previous section, we found evidence of both positive (CF -> CF) and negative transfer (CL -> CF) effects. These transfer effects are consistent with the learning and transfer of gating policies. On this account, participants learn a gating policy in the training block which they transfer to the test block. Initial performance in the test block, is supported by the transferred policy. Participants in the CF -> CF group learn a selective input-gating policy in the initial CF training block thus enabling efficient performance in the subsequent CF test block. Participants in the CL -> CF group, on the other hand, learn a selective output-gating policy in the training block. This policy continues to support performance in the changed circumstances of the CF test block, thus reducing initial learning demands, but at the cost of slower responses. In contrast, in the CF -> CL group, a transferred selective input-gating policy would not support performance in the changed circumstances, thus immediately requiring the deployment of a new policy. This account predicts that initial behaviour in the CL -> CF test block is under the control of a selective output-gating policy, while initial behaviour in the CF -> CF test block is under the control of a selective input-gating policy, despite both the block being identical.

A feature of the task design allows us to test these predictions. CF trials come in two forms, those in which the target item is in the second position (target middle), and those in which the target item is in the last position (target last). The final response decision requires knowledge of the target item, and therefore cannot be initiated or prepared for until target selection has taken place. Critically, the timing of target selection is contingent on the kind of gating policy that one employs. Under a selective output-gating policy, target selection does not proceed until the final item is presented. Therefore, no benefit can be gained on target-middle trials on which target information is available early. On the other hand, under a selective input-gating policy, target selection occurs at the level of input-gating. Therefore, on target-middle CF trials, target selection can occur once the second item in the sequence is processed, and one can begin preparing for the upcoming response decision; for example, setting up an attentional set to search for the target item in the response panel. Thus, under an input-gating policy only, a *target position benefit* can be expected on target-middle vs. target-last trials. No such benefit would follow from an output-gating policy.

Consistent with the account developed above, if participants in the CL -> CF group transfer a selective output-gating policy from the training block, target position benefits should be unavailable on the early trials of the test block. As participants learn a selective input-gating policy with experience, target position benefits would become available later in the block. On the other hand, in the CF -> CF test blocks participants are more likely to transfer a selective input-gating policy (Chatham et al., 2014) and robust target position benefits should be observed on the early trials as well. Finally, target position benefits should not be observed at all in the CL -> CL and CF -> CL test blocks, which thus serve as a negative control.

Target position benefits were computed by subtracting log-transformed RTs on target-middle trials from those on target-last trials and are plotted in Figure 3 (C & D). To begin with, we employed one-sample *t*-tests to examine the presence of target position benefits. As predicted, no target position benefit was found on the early (first 12) trials of the test block in the CL -> CF group [one sample t (18) = 0.89, $p > 0.250$, $d = 0.2$]. On the other hand, a robust target position effect was observed on the late (last 12) trials [one-sample t (18) = 5.62, $p < 0.001$, $d = 1.29$]. A paired *t*-test confirmed a significant effect of trial bin on the size of the target position benefit [t (18) = 2.86, $p = 0.010$, $d = 0.66$].

We next compared target position benefits in the CL -> CF and CF -> CF test blocks. A mixed-model ANOVA with trial bin (2 levels: first 12 trials bin, last 12 trials bin) as a within-subject factor and group (2 levels: CF -> CF, CL -> CF) as a between-subject factor confirmed a significant trial bin x group interaction [F (1, 36) = 4.61, $p = 0.039$, $\eta_p^2 = 0.114$]. As a final control, we examined target position benefits in the CL -> CL test block and found no evidence of a target position benefit. On the early trials, however, there was a slight target position cost [t (18) = -2.6, $p = 0.020$, $d = -0.6$].

In summary, these results provide positive evidence for the transfer of a selective output-gating policy from the training to the test block in the CL -> CF group, and the learning of selective input-gating policy with experience in the test block.

Experiment 2

The results of Experiment 1 provide evidence for the transfer of working memory gating policies. In particular, we observed negative transfer across changes in trial structure, even as S-R rule structure was held constant. In Experiment 2, we attempted to replicate and extend these findings. We carried out a stronger test of the hypothesis that gating policies and S-R rules are separable by independently manipulating prior experience with a rule vs. a trial structure. We replicated the effects of prior experience with a trial structure, and found that changing a rule has little effect on later performance.

Methods

Participants.

117 adult, right-handed participants (45 males, 72 females; age-range: 18-30, M = 21.6, SD = 2.6) from the Providence, RI area were recruited to take part in the computer-based experiment. We endeavored to collect between 18-20 participants in each of six groups based on approximate effect sizes suggested by pilot data and Experiment 1. 4 participants were excluded as they were on psychoactive medication. 5 participants were excluded because of low performance (<70% accuracy) on the task. This left 108 participants (40 males, 68 females; age-range: 18-30, M = 21.4, SD = 2.6). All remaining participants had normal or corrected-to-normal vision, and no reported neurological or psychological disorders. All participants gave informed, written consent as approved by

the Human Research Protections Office of Brown University, and they were compensated for their participation.

Task and Experiment Design.

Experiment 2 used the same task as Experiment 1. Participants were randomly assigned to one of six experimental groups. Four of these groups performed identical transfer conditions to the four conditions used in Experiment 1. Two additional groups (CF -> CL (different rules) and CL -> CF (different rules), were presented with novel rules at the beginning of each block. The new rules utilized the same categories (numbers for context, letters and symbols for lower-level items) but a novel set of numbers, letters and symbols. The assignment of each rule set to the training and test block was counterbalanced across participants and groups.

Results

Overall performance was strong and comparable between the two experiments (Table 1). Summary statistics for all measures are displayed in Table 2. Results are presented in two sections. In the first section, we replicate the transfer effects and target position effects from Experiment 1. In the final section, we turn to a comparison of the same-rule and different-rule groups.

Transfer effects and target position effects.

Behavioural results from the same rule conditions are plotted in Figure 5 and the results of the ANOVAs are presented in Table 3. In summary, we replicated the asymmetric negative transfer effects from Experiment 1, finding a significant trial bin x training history interaction [F (7, 238) = 4.99, $p < 0.001$, $\eta_p^2 = 0.128$] on RTs in the CF condition (Figure 5A), but not the CL condition (Figure 5B). We also obtained an asymmetric positive transfer effect, with a significant trial bin x block interaction on response accuracies in the CF -> CF, [F (7, 238) = 4.90, p < 0.001, $\eta_p^2 = 0.214$], but not the CL -> CL conditions.

We also replicated the pattern of target position effects observed in Experiment 1. A non-significant trend for a target position cost was observed on the early (first 12) trials of the test block in the CL -> CF group [one sample t (17) = -1.92, p = 0.072, d = -0.45]. A significant target position benefit was observed on the late (last 12) trials [one-sample t (17) = 2.83, p = 0.012, d = 0.67]. A paired t-test confirmed a significant effect of trial bin on the size of the target position benefit [t (18) = 2.95, p = 0.009, d = 0.70]. ANOVAs (full detail provided in Table 3) confirmed that a significant trial bin (first 12 vs. last 12 trials) x training history interaction was present for the CF condition, [F (7, 238) = 4.55, p = 0.04, $\eta_p^2 = 0.118$] but not for the CL condition [F (7, 238) = 1.55, p = 0.222, η_p^2 = 0.044].

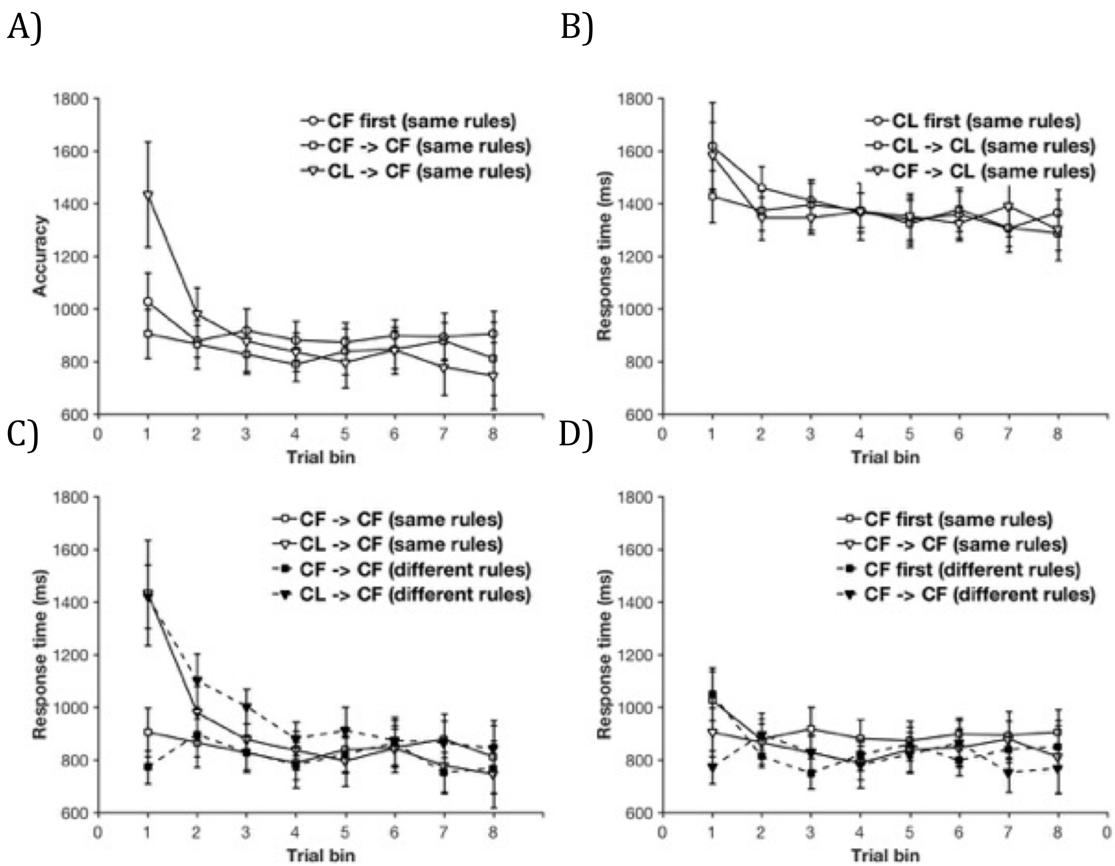

Figure 4: Results from Experiment 2. (A) & (B) – Transfer effects for CF (A) and CL (B) trial structures when rules remain the same. Mean RT for each 6-trial bin is plotted as a function of training history - no prior training (circles), prior training with same structure (squares), or different trial structure (triangles). Transfer effects observed in Experiment 1 are replicated here. (C) & (D) - Comparison of negative (C) and positive (D) transfer effects in the different-rule (open shapes, dashed lines) and corresponding same-rule groups (filled shapes, solid lines). The transfer effects are primarily driven by previous experience with a trial structure rather than with a rule. Error bars reflect 95% confidence intervals.

Table 3: Experiment 2 summary ANOVA results

Statistical test	*Effect*	*df*	*F*	*p*	η_p^2
	Negative transfer Effects				
CF trials response time: Trial Bin x Training History mixed-model ANOVA	Trial bin ME*	(7, 238)	10.85	<0.001	0.242
	Training history ME	(1, 34)	0.38	>0.25	0.011
	Trial bin x training history interaction*	(7, 238)	4.99	<0.001	0.128
CF trials response accuracy: Trial Bin x Training History mixed-model ANOVA	Trial bin ME*	(7, 238)	5.64	<0.001	0.142
	Training history ME	(1, 34)	0.88	>0.250	0.025
	Trial bin x training history interaction	(7, 238)	1.14	>0.250	0.032
CL trials response time: Trial Bin x Training History mixed-model ANOVA	Trial bin ME*	(7, 238)	5.98	<0.001	0.150
	Training history ME	(1, 34)	0.17	>0.250	0.005
	Trial bin x training history interaction	(7, 238)	0.69	>0.250	0.019
CL trials response accuracy: Trial Bin x Training History mixed-model ANOVA	Trial bin ME*	(7, 238)	4.43	<0.001	0.115
	Training history ME	(1, 34)	1.38	0.249	0.039
	Trial bin x training history interaction	(7, 238)	1.24	>0.250	0.035
	Positive transfer Effects				
CF trials response time: Trial Bin x Block repeated-measures ANOVA	Trial bin ME	(7, 126)	1.84	0.085	0.093
	Block ME * (1,17)	(1, 18)	4.51	0.048	0.200
	Trial bin x block	(7, 126)	0.58	>0.250	0.031
CF trials response accuracy: Trial bin x block repeated-measures ANOVA	Trial bin ME*	(7, 126)	2.22	0.037	0.110
	Block ME*	(1, 18)	7.06	0.016	0.268
	Trial bin x block*	(7, 126)	4.90	<0.001	0.214
CL trials response time: Trial bin x Block repeated-measures ANOVA	Trial bin ME*	(7, 126)	2.66	0.014	0.129
	Block ME	(1, 18)	0.52	>0.250	0.028
	Trial bin x block	(7, 126)	0.74	>0.250	0.040
CL trials response accuracy: Trial bin x Block repeated-measures ANOVA	Trial bin ME*	(7, 119)	4.01	0.001	0.190
	Block ME	(1, 18)	0.01	>0.250	0.001
	Trial bin x block	(7, 119)	1.23	>0.250	0.067
	Target position effects				
CF trials response time: Trial bin x Training history mixed-model ANOVA	Trial bin ME*	(1, 34)	9.53	0.004	0.219
	Training history ME*	(1, 34)	6.59	0.015	0.162
	Trial bin x training history interaction*	(1, 34)	4.55	0.040	0.118
CL trials response time: Trial bin x Training history mixed-model ANOVA	Trial bin ME	(1, 34)	3.55	0.068	0.095
	Training history ME*	(1, 34)	0.21	>0.250	0.006
	Trial bin x training history interaction*	(1, 34)	1.55	0.222	0.044

Effect of rule versus trial structure change

Finally, we turn to a comparison of the same-rule and different-rule groups and examine the relative influence of prior experience with rules vs. trial structures on subsequent performance.

Log-transformed RT data from the test blocks of four groups: CF -> CF (same rule), CL -> CF (same rule), CF -> CF (different rule) and CL -> CF (different rule) (Figure 4C and 4D) were analysed with a mixed-model with trial bin (8 levels: 1-8) as a within-subject factor and previous trial structure (2 levels: same trial structure, different trial structure) and previous rule (2 levels: same rule, different rule) as between-subject factors. The ANOVA revealed a significant main effect of trial bin [F (7, 476) = 19.39, $p < 0.001$, $\eta_p^2 = 0.222$] and a trial bin x previous trial structure interaction [F (7, 476) = 13.31, $p < 0.001$, $\eta_p^2 = 0.164$]. Importantly, neither the main effect of the previous rule [F (7, 476) = 0.001, $p > 0.250$, $\eta_p^2 = 0.00$], nor the three-way interaction between trial bin, previous trial structure, and previous rule [F (1, 68) = 0.03, $p > 0.250$, $\eta_p^2 = 0.00$] was significant, suggesting that prior experience with a rule had no influence on subsequent performance. To further confirm that prior experience with a rule did not subtly influence the slope of the RT curves in the CL -> CF test blocks (Figure 6A), we fit power functions to each subject's curve and entered the estimated parameters into an independent-sample t-test comparing the rule change and no rule change group. We found no reliable differences [t (34) = 0.68, $p = 0.208$, $d = 0.23$].

Analysis of the accuracy data, similarly, revealed a significant main effect of previous trial structure [F (1, 68) = 7.75, $p = 0.007$, $\eta_p^2 = 0.102$] and a trial bin x previous trial structure interaction [F (7, 476) = 3.45, $p = 0.001$, $\eta_p^2 = 0.048$]. Again, the main effect of the previous rule was non-significant [F (1, 68) = 0.622, $p > 0.250$, $\eta_p^2 = 0.009$] as was the three-way interaction between trial bin, previous trial structure, and previous rule [F (7, 476) = 0.89, $p > 0.250$, $\eta_p^2 = 0.013$].

We also specifically compared positive transfer effects in the same rule and different rule groups. Log-transformed RT data from both training and test blocks of the CF -> CF (same rule), and CF -> CF (different rule) groups were analysed with a mixed-model with block (2 levels: training or test) and trial bin (8 levels: 1-8) as within-subject factors and previous rule (2 levels: same rule, different rule) as a between-subject factor. The ANOVA revealed significant main effects of block [F (1, 238) = 10.07, $p = 0.003$, $\eta_p^2 = 0.229$] and trial bin [F (7, 238) = 3.03, $p = 0.005$, $\eta_p^2 = 0.082$], and a significant block x trial bin interaction [F (7, 238) = 3.20, $p = 0.003$, $\eta_p^2 = 0.086$]. Importantly, the main effect of the previous rule was non-significant [F (1, 34) = 0.31, $p > 0.250$, $\eta_p^2 = 0.009$] while the three-way interaction between trial bin, previous trial structure, and previous rule was marginally non-significant [F (1, 68) = 1.99, $p = 0.068$, $\eta_p^2 = 0.055$]. Overall, this suggests that while prior rule experience may have some influence on subsequent performance, the positive transfer effect is largely driven by prior experience with a trial structure.

Discussion

Psychologists have focused almost exclusively on the relations between stimuli, contexts, responses and outcomes as a framework for understanding cognitive control and our ability to adapt and generalize to novel task environments. In this paper, we developed the hypothesis that internal control policies, required for coordinating cognitive processing during a task, form an essential component of task knowledge independently of the stimulus-response-outcome (S-R-O) rule structure of the task. Thus, we sought to test whether control policies can be learned and transferred to novel task situations separate from S-R-O rules. We tested this hypothesis in the context of a working memory control task, leveraging previous findings that subjects learn different selective gating policies for context first or context last trial structures (Chatham et al., 2014). We argue that the results presented in this paper are clear evidence in support of the hypothesis.

First, the experiments provided evidence of both positive and negative transfer. Participants with prior experience in a CF block responded faster on the initial trials of a subsequent CF block. On the other hand, participants with prior experience with a CL block were slower to respond on the initial trials of a subsequent CF block compared to participants with no prior experience. This effect was not driven simply by WM load experienced during the training block, as participants in the CF2 -> CF group, who were matched for load, did not show such slowing. Crucially, the negative transfer effects were asymmetric, in that they were not obtained in participants carrying out a CL block after having previously experienced a CF block. Thus, simply encountering a novel ordering or a task switch was not a cause of the initial trial slowing. Rather, such an asymmetric effect is consistent with gating policy transfer, as a selective input-gating policy would not support performance on CL trials and so could not be transferred. Therefore, participant's behavior would resemble those with no prior experience.

Second, the pattern of target position effects we obtained provided positive evidence of the transfer of a gating policy. As previously discussed, target position benefits are contingent on utilizing a selective input-gating strategy. Participants who had previously experience the CF trial structure, but not those who had previously experienced only the CL trial structure, showed a target position effect on the initial trials of a subsequent CF block. In other words, the early trial behaviour in the test block consistent with the trial structure of the previously experienced trial structure. In other words, the initial slowing observed in the test block of the CL -> CF group was due to negative transfer of a less efficient gating policy. Toward the end of the block, however, as they gained experience and the negative transfer effects on RT diminished, a robust target position benefit emerged. This result directly supports our contention that participants transferred a selective output-gating policy from the CL to CF block.

A key feature of Experiment 1 is that these learning and transfer effects occurred without any change in the prevailing S-R rule structure. To further test the independence of internal control policies from S-R rules, Experiment 2 independently manipulated prior experience with an S-R rule versus the trial structure and showed that both the negative and positive transfer effects are primarily driven by prior experience with a trial structure, regardless of whether the specific rule repeated or was new.

This observation suggests that S-R-O rules and gating policies may be dissociable constructs. Strikingly, positive transfer was obtained even when participants were required to use new rules, reflecting remarkable generalization. One possible interpretation of this finding is that gating policies are abstract entities, not tied to specific S-R-O mappings. Indeed, gating policies may be thought of as a collection of abstract rules relating a set of gating operations to broad task contexts (see Taatgen, 2013 for a related view). However, we acknowledge that our rule change manipulation was minimal – we only modified individual items in the task, keeping category structure intact.

Therefore, generalization may be a consequence of category knowledge rather than an abstract gating policy. It is plausible that a stronger manipulation of rule or overall task context would produce larger effects of prior rule experience that may interact with gating policy transfer. Indeed, previous studies have demonstrated transfer of task rules (Badre et al., 2010; Cole et al., 2011; Collins & Frank, 2013; Shanks & Darby, 1998). Future work will be necessary to define the limits gating policy generalization and its relationship with task rules.

Instead, we interpret our findings as evidence that internal control policies embody an important form of structural task knowledge that supports task performance in novel situations. Our results add to a growing literature on *structure learning,* extending its principles to the domain of cognitive control policies (Braun, Mehring, & Wolpert, 2010; Collins & Frank, 2013; Gershman, Blei, & Niv, 2010; Huys et al., 2015). Structure learning refers to our ability to identify and leverage invariant structure in natural tasks to improve learning of novel or complex tasks (Botvinick et al., 2009; Botvinick et al., 2015; Gershman & Niv, 2010; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). Tasks that share structure may load overlapping internal cognitive processes or engage them in similar sequential orders, affording an opportunity for re-using learned control policies in novel task settings. Our paper provides the first evidence that working memory gating policies leverage a form of task structure separate from S-R relationships and are re-used across task contexts.

Our notions of selective input v/s output gating policies bear some resemblance with the notions of proactive v/s reactive modes of control postulated within the dual mechanisms of control framework (Braver, 2012). On this view, selective input-gating policies rely on a proactive mode in which a goal-relevant contextual representation is activated and sustained in the anticipation of future control demand. On the other hand, selective output-gating policies rely on a reactive mode in which the goal-relevant contextual representation is transiently activated when the control demand arises. From this perspective, our transfer effects could also be viewed as providing evidence for the persistence of a control mode across different tasks, not tied to the underlying S-R-O rules. Note, however, that the finding of transfer of a selective output-gating policy from the CL to CF task in our experiments strains against the notion of reactive control being transiently triggered, 'on-the-fly', as a result of current control demands. Instead, our transfer effects are more consistent with the transfer of a reactive control *policy,* regardless of the nature of current control demands.

Finally, our results raise new questions about the mechanisms of gating policy learning. Several computational models implement WM gating via cortico-basal ganglia loops and learning via dopaminergic prediction error signals acting on striatal synapses (reviewed in Bhandari, Badre, & Frank, in press; O'Reilly & Frank, 2006). What mechanisms support the learning and generalization of abstract gating policies to novel tasks? One proposal is that generalization relies on indirection mechanisms implemented via output-gating that support the separate representation of 'roles' and their 'fillers' (Kriete, Noelle, Cohen, & O'Reilly, 2013). Alternatively, generalization may rely on independently acquired category knowledge (Seger & Miller, 2010). Finally, generalization may be achieved by modifying domain general properties of control like action-selection decision thresholds, WM input-gating thresholds, or arousal levels.

Understanding the mechanisms supporting broad generalization of gating policies would be critical to explaining the remarkable flexibility of human cognition.

Author contributions

A. Bhandari and D. Badre designed the study. A. Bhandari conducted the experiments and analyzed the data. A. Bhandari and D. Badre wrote the paper.

Acknowledgements

We thank Ryan K. Fugate, Aja Evans, Celia Ford, and Adriane Spiro for assistance with data collection. This work was supported by grants from NINDS (NS065046) and NIMH (MH099078, MH111737) at the NIH, and a MURI award from the Office of Naval Research (N00014-16-1-2832).

References

Ackerman, P. L. (1988). Determinants of Individual-Differences during Skill Acquisition - Cognitive-Abilities and Information-Processing. *Journal of Experimental Psychology-General, 117*(3), 288-318. doi:Doi 10.1037/0096-3445.117.3.288

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

Bhandari, A., Badre, D., & Frank, M. J. (in press). Learning cognitive control. In T. Egner (Ed.), *The Wiley Handbook of Cognitive Control*. Oxford: John Wiley & Sons.

Bhandari, A., & Duncan, J. (2014). Goal neglect and knowledge chunking in the construction of novel behaviour. *Cognition, 130*(1), 11-30. doi:10.1016/j.cognition.2013.08.013

Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition, 113*(3), 262-280. doi:10.1016/j.cognition.2008.08.011

Botvinick, M. M., Weinstein, A., Solway, A., & Barto, A. C. (2015). Reinforcement learning, efficient coding, and the statistics of natural tasks. *Current opinion in behavioral sciences, 5*, 71-77.

Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural brain research, 206*(2), 157-165.

Chatham, C. H., & Badre, D. (2015). Multiple gates on working memory. *Current opinion in behavioral sciences, 1*, 23-31.

Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal output gating during selection from working memory. *Neuron, 81*(4), 930-942. doi:10.1016/j.neuron.2014.01.002

Cohen-Kdoshay, O., & Meiran, N. (2007). The representation of instructions in working memory leads to autonomous response activation: Evidence from the first trials in the flanker paradigm. *Q J Exp Psychol (Hove), 60*(8), 1140-1154. doi:10.1080/17470210600896674

Cohen-Kdoshay, O., & Meiran, N. (2009). The Representation of Instructions Operates Like a Prepared Reflex Flanker Compatibility Effects Found in First Trial Following S-R Instructions. *Experimental Psychology, 56*(2), 128-133. doi:10.1027/1618-3169.56.2.128

Cole, M. W., Bagic, A., Kass, R., & Schneider, W. (2010). Prefrontal Dynamics Underlying Rapid Instructed Task Learning Reverse with Practice. *Journal of Neuroscience, 30*(42), 14245-14254. doi:10.1523/Jneurosci.1662-10.2010

Cole, M. W., Etzel, J. A., Zacks, J. M., Schneider, W., & Braver, T. S. (2011). Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. *Frontiers in human neuroscience, 5*.

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychol Rev, 120*(1), 190.

Duncan, J., Parr, A., Woolgar, A., Thompson, R., Bright, P., Cox, S., . . . Nimmo-Smith, I. (2008). Goal neglect and Spearman's< em> g: Competing parts of a complex task. *Journal of Experimental Psychology: General, 137*(1), 131.

Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. *Cereb Cortex, 22*(3), 509-526. doi:10.1093/cercor/bhr114

Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychol Rev, 117*(1), 197.

Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol, 20*(2), 251-256. doi:10.1016/j.conb.2010.02.008

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. *Philos Trans R Soc Lond B Biol Sci, 362*(1485), 1601-1613. doi:10.1098/rstb.2007.2055

Huys, Q. J., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., . . . Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences, 112*(10), 3098-3103.

Kriete, T., Noelle, D. C., Cohen, J. D., & O'Reilly, R. C. (2013). Indirection and symbol-like processing in the prefrontal cortex and basal ganglia. *Proc Natl Acad Sci U S A, 110*(41), 16390-16395. doi:10.1073/pnas.1303547110

McClelland, J. L. (2009). Is a machine realization of truly human-like intelligence achievable? *Cognitive Computation, 1*(1), 17-21.

Meiran, N., Pereg, M., Kessler, Y., Cole, M. W., & Braver, T. S. (2015). The Power of Instructions: Proactive Configuration of Stimulus-Response Translation. *Journal of Experimental Psychology-Learning Memory and Cognition, 41*(3), 768-786. doi:10.1037/xlm0000063

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 Comput, 18*(2), 283-328. doi:10.1162/089976606775093909

Radvansky, G. A., & Zacks, J. M. (2014). *Event cognition*: Oxford University Press.

Ruge, H., & Wolfensteller, U. (2010). Rapid formation of pragmatic rule representations in the human brain during instruction-based learning. *Cereb Cortex, 20*(7), 1656-1667. doi:10.1093/cercor/bhp228

Schank, R. C., & Abelson, R. P. (1977). Scripts, plans, goals and understanding: An inquiry into human knowledge structures.

Seger, C. A., & Miller, E. K. (2010). Category learning in the brain. *Annu Rev Neurosci, 33*, 203.

Shanks, D. R., & Darby, R. J. (1998). Feature-and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 24*(4), 405.

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychol Rev, 120*(3), 439.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science, 331*(6022), 1279-1285. doi:10.1126/science.1192788

Todd, M. T., Niv, Y., & Cohen, J. D. (2009). *Learning to use working memory in partially observable environments through dopaminergic reinforcement.* Paper presented at the Advances in neural information processing systems.

Wolfensteller, U., & Ruge, H. (2011). On the timescale of stimulus-based action–effect learning. *The Quarterly Journal of Experimental Psychology, 64*(7), 1273-1289.

Appendix A

Verbatim instructions provided to participants for Experiment 1:

"The task you will be doing today involves responding to sequences of items being presented on the computer screen. Each sequence will consist of three items presented one after the other. Your job is to identify a target item from the sequence based on a couple of rules. (Show the rules)

These are the two rules. They are in the form of trees. Each tree has a number at the top and associated with each number are a pair of items. 11 goes with the letters A and G. 53 goes with the symbols, ⊙ (dot) and π (pi). You need to memorize these relationships.

Now, the sequences that you will see will consist of 1 item from this pair (point to the letters pair), 1 item from this pair (point to the symbols pair), and 1 of the two numbers. For example, you might see the sequence A-Pi-11. Or the sequence might be 53-Dot-G. Your job is to identify the target item. The target item is always the item in the sequence that is associated with the number in the sequence.

(Show first example trial.) So, in the sequence A- π -11 – the target item would be A and not Pi, because A is associated with 11, while π is not. Similarly, in the sequence 53 - ⊙ - G, the target would be ⊙. So, the number tells you which of the other two items is the target. Is that clear? (If not, repeat and clarify).

At the end of the sequence, just below the last item in the sequence, you will see a response panel. The response panel will have two items on the left and two items on the right. You have to look for the target item you have identified in this panel. If the target item that you have identified occurs as part of the left pair, you should press button 1. If the target item occurs as part of the right pair, you should press button 2. You should use the first two fingers of your right hand to make your response. You should make your response as quickly as possible without making errors.

The order in which the items are presented in the sequence is not important. Sometimes, the number might appear at the beginning of the sequence and sometimes it might appear at the end. Either way, the rules apply in exactly the same way. The target is always the item that goes with the number.

The items in the sequence each appear for a very short duration. The items will also be presented fairly quickly one after the other. So it is very important that you be alert and pay attention at all times. You will be doing a total of 2 blocks. It is important that you maintain a level of high alertness throughout the entire experiment.

Each block will begin with the rules being presented to you. You can review them for as long as you want. Once you are ready, press space bar to begin the block. There is no feedback, so make sure that you know the rules before you start the block. The rules are the same across all the blocks.

We will begin now. Remember to response as quickly as possible while being accurate."