Learning when to learn: Context-specific instruction encoding

Senne Braem^{1,2}, Shengjie Xu¹, Baptist Liefoghe³, & Elger L. Abrahamse^{4,5}

¹Department of Experimental Psychology, Ghent University, Ghent, Belgium

²Department of Experimental and Applied Psychology, Vrije Universiteit Brussel, Brussel, Belgium

³Department of Psychology, Utrecht University, the Netherlands

⁴Department Communication and Cognition, Tilburg University, the Netherlands

⁵Department of Educational Sciences, Atlántico Medio University, Spain

Correspondence

Senne Braem

Department of Experimental Psychology

Henri Dunantlaan 2

9000 Gent

BELGIUM

senne.braem@ugent.be

Acknowledgements: We would like to thank Tom Verguts for valuable comments on a previous draft of the manuscript. This work was supported by an FWO - Research Foundation Flanders grant (12K6319N) and an ERC Starting grant (under European Research Council's Horizon 2020 programme, grant 852570) awarded to SB.

Abstract

Humans are efficient at learning new behaviors through verbal instructions, bypassing the need for trial-and-error learning, but requiring a certain degree of cognitive flexibility. Recent theories suggest that abstract functions of cognition, like cognitive flexibility, can be (up)regulated by contextual features in our environment with which they are associated. Here, we tested whether people show a more flexible control mode in contexts where one frequently learns new task instructions. Specifically, participants had to encode and execute a series of new task instructions, and the likelihood of encountering a new (relative to old) task instruction varied in function of whether the instruction was presented above versus below the center of the screen (Experiment 1, $n = 40$), or in a blue versus green font (Experiment 2, $n = 40$). We measured the level of interference from retained new task instructions on secondary task performance to index cognitive flexibility, and show enhanced flexibility (i.e., enhanced receptibility for new task instructions) for contexts in which the likelihood of encountering new task instructions was relatively high. Similarly, participants were also faster to encode new instructions in these contexts. A third, online study ($n = 77$) replicated the main finding of Experiment 2, whose context manipulation (i.e., font color) was most effective. Together, these findings suggest that cognitive flexibility can bind to, and be signaled by, contextual features in our environment. Such ‘learning when to learn’ provides a (self-)regulating mechanism for the cognitive processes that support the encoding (and execution) of new task instructions.

Keywords: Cognitive control; Context-sensitivity; Instructions

Humans share the intriguing ability of quickly acquiring new behaviors through verbal instructions. This ability is believed to depend on our capacity to form symbolic representations (Deacon, 1997), and is increasingly recognized as an important cognitive control function (e.g., Cole, Braver, & Meiran, 2017). In the assumption that cognitive control functions are organized hierarchically (Badre, 2008), the ability to flexibly (re)configure the cognitive system on the basis of new task instructions could arguably be placed on top of this hierarchy (Abrahamse et al., 2016). However, being on “top” of top-down control functions naturally poses the question about what drives (and thus controls) this encoding of task instructions? How do we know when to encode an instruction, and how strongly to do so?

Recent reviews and theories argue that cognitive control functions are regulated by “simple” associative learning or reinforcement learning mechanisms (Abrahamse et al., 2016; Braem & Egner, 2018; Chiu & Egner, 2019; Egner, 2014; Eisenreich et al., 2017; Lieder et al., 2018; Verguts & Notebaert, 2008). By learning that certain control functions are applied more (or more valuable) in certain environments than others, these control functions become associated in memory with contextual features that characterize those environments. These theories are in part inspired by the observation that basic – and often task-irrelevant – contextual stimulus features can trigger cognitive-control processes with which they have become associated. For example, Mayr and Bryck (2007) demonstrated that conflicting task rules induced larger interference effects when the tasks were presented on the same location on the screen (on different points in time), as opposed to when they were presented on different locations on the screen, even though screen location was irrelevant to the task. In a similar vein, Crump and colleagues (2006) showed that a location where mostly incongruent Stroop stimuli were presented, automatically resulted in smaller interference effects as opposed to a location where mostly congruent stimuli were presented (i.e., the context-specific proportion congruency effect). This finding suggests that the ability to deal with cognitive conflict can be up- or down-regulated by contextual features in our environment (e.g., Bugg & Crump, 2012; Chiu & Egner, 2019; but see, Braem et al., 2019a; Schmidt, 2019).

Here, we investigate whether the ability to encode new task instructions, as a marker of cognitive flexibility, can be modulated by the context in which these instructions need to be learned. To this end,

we adopted a paradigm where the strength by which, and flexibility with which, newly instructed actions are encoded can be evaluated (Liefoghe et al., 2012, 2013; Meiran et al., 2015). In this task, each run starts with the presentation of instructions for the ‘inducer task’, where two words are assigned to left and right response buttons on the basis of their identity (e.g., “if ‘goat’ press left, if ‘wine’ press right”). Participants are instructed to only respond to the word identity when the word is presented in a particular font color (e.g., green). Crucially, in the retention interval between the offset of the instructions and their actual application, several trials of a ‘diagnostic’ task are presented, which include the same words presented in another font (e.g., white). While the instructions for the inducer task change on each run, the instructions of the diagnostic task remain constant throughout the experiment: “for words presented in upright font, press left; for words presented in italic font, press right”. Although the identity of the words is irrelevant for the diagnostic task, it has been shown to impact behavior. For example, the word “goat” presented in italic would require a right button response in the diagnostic task, but a left button response in the pending inducer task. Because of this, performance is slowed down in the diagnostic task. This observation was coined the instruction-based congruency effect (for reviews, see Brass, Liefoghe, Braem, & De Houwer, 2017; Meiran, Liefoghe, & De Houwer, 2017), and it offers an index of the strength or intensity with which new instructions have been encoded into a task-set (e.g., Braem et al., 2019b; Whitehead & Egner, 2018; for similar investigations into more long-term intentions, see Gollwitzer, 1999; Goschke & Kuhl, 1993). For example, the instruction-based congruency effect has been found to increase whenever the likelihood increases that one will also have to execute the inducer task (Whitehead & Egner, 2018), and it is positively related to the efficiency with which the inducer task is performed (both within and between subjects, Braem et al., 2019b).

To study whether contextual features in our environment can help shift us towards a more flexible mode to encode new task instructions, we adjusted the above described paradigm by applying two topographically different context manipulations which had the same functional rule (see Figure 1): in one context the instructions of the inducer task were more likely (75%) to be new (i.e., they had not been seen before), whereas in the other context they were not (25%). Specifically, we presented the instructions for the inducer task, as well as the target for the inducer task, either above or below a central

fixation cross (Experiment 1), or in a blue or green font (Experiments 2 and 3). The diagnostic trials, however, were always presented centrally in a white font. This way, we could investigate whether the context in which the inducer task was instructed (i.e., stimulus location or stimulus font color) would modulate the automatic interference that the inducer task instructions elicit in the diagnostic task. Namely, if people can engage in a more, versus less, flexible mode depending on the context, we would expect a larger impact of the instructed task-rules when the instruction was presented in a context where people more frequently encounter new task instructions. The logic of this is as follows: In a context with mostly new instructions, contextual features signal the cognitive system to enhance the encoding of new information, such that the task instructions are imprinted faster and more efficiently – resulting in more potent stimulus-response representations that elicit stronger interference during the diagnostic task (than is the case for task instructions presented in the context with mostly repeating instructions). Similarly, contextual features that signal a more flexible mode might promote the activation of multiple task representations in working memory leading to more interference between the inducer and diagnostic task (e.g., Braem & Egner, 2018; Dreisbach & Fröber, 2018; Goschke, 2003; Goschke & Bolte, 2014; Musslick & Cohen, 2022).

Experiments 1 and 2 are reported together because their design and analyses were highly identical except for the contextual feature determining the likelihood of encountering new task instructions – and because we had no a priori expectations about which context manipulation would be more effective. Although we first discuss the results of Experiments 1 and 2 separately for the sake of completeness, the across-experiment analyses and subsequent replication of findings in Experiment 3 are considered the main tests of our hypothesis that instructions presented in contexts where mostly new tasks appear should result in bigger instruction-based congruency effects.

Method

Participants. 82 students took part in return for course credits or 10€. Our previous study using a highly similar design and analysis pipeline (but different research question) used 20 to 32 subjects per experiment (Braem et al., 2017). Here, we recruited 40 students for Experiment 1 (range = 17-24 years,

25 women and 15 men, 38 right-handed), and 42 students for Experiment 2 (range = 17-21 years, 33 women and 7 men, 32 right-handed). We tested two additional subjects in Experiment 2, because the performance of two participants in Experiment 2 was below chance level (i.e., 50%) on the diagnostic task, and these subjects were thus removed from further analyses. This way, both Experiments had 40 subjects eligible for analysis, and we had 80% statistical power to detect an effect size of $\eta_p^2 = .05$ (as computed with G*Power 3.1.9.2; Faul, Erdfelder, Lang, & Buchner, 2007).

In Experiment 3, we report a replication of Experiment 2. For this experiment, we tested students from Ghent University in exchange for credits (age range = 17-21 years, 72 women and 8 men). Different from Experiments 1 and 2, Experiment 3 was administered online. Two participants were removed from data analyses due to low accuracy (< 55%) either in diagnostic or inducer trials. One more participant was removed because a severe data missing issue, with 10 % of trials being missing. The final sample included 77 participants eligible for further statistical analyses, which was roughly double the sample size of Experiment 1 and 2.

Stimuli and material. The stimuli consisted of 106 highly frequent four-letter Dutch nouns which were selected using the SUBTLEX-NL database (Keuleers, Brysbaert, & New, 2010; 98 of which were also used in Braem et al., 2017). These stimuli were randomly combined into 53 unique pairs of words for each participant separately, of which one pair was assigned to the practice run and the other 52 to the twelve blocks of experimental runs (blocks were interspersed with self-paced brakes). Four of these 52 pairs were randomly assigned to be “old” pairs, and the 48 remaining pairs were “new” pairs. Each old pair was presented twelve times throughout the experiment, while each new pair was presented only once. For practical reasons, ten instead of twelve blocks were presented in Experiment 2 and 3 (i.e., the old pairs were presented ten times, and 8 new pairs were never presented).

The words were always presented in an Arial 24-point font. However, the exact presentation of the words further depended on whether they were presented during diagnostic or inducer trials. Specifically, in the diagnostic task, words were always presented centrally in white, but could either be presented in an italic or upright font which determined the correct response. Importantly, the congruency of the words

was randomized in order, but balanced within run. The instructions for the inducer task always consisted of three lines, and were presented above or below the fixation cross in white on a black background in Experiment 1, or centrally in green or blue in Experiment 2 (see Figure 1B). The location of these instructions corresponded to the location on which the stimulus for the inducer task would be presented (halfway between the fixation cross and the edge of the screen) in Experiment 1, or the font color in which the inducer task would be presented in Experiment 2. The words during the inducer task in Experiment 1 were always presented in green.

Crucially, the contextual feature (location in Experiment 1; or font color in Experiment 2) was predictive of whether an instruction would be an old or new pair (75% / 25%). Counterbalanced across participants, one context was more likely to contain a new instruction (i.e., “mostly new context”), whereas the other context was more likely to contain old instructions (i.e., “mostly old context”). Specifically, each block consisted of eight runs presented in a random order, of which three used new instructions presented in a mostly new context, three used old instructions presented in a mostly old context, one was a new instruction in a mostly old context, and one was an old instruction in a mostly new context (see Figure 1A). As mentioned above, these four “old” instructions would reoccur across blocks, whereas new instructions were presented only once throughout the entire experiment. A run could consist of either 4, 8, 12, or 16 diagnostic trials (one of each was presented per context type per block). Stimuli were presented on a CRT monitor located 60 cm away from the eyes using Tscope software (Stevens, Lammertyn, Verbruggen, & Vandierendonk, 2006). The left and right response key were the letters F and J, respectively, on a standard QWERTY keyboard. Experiment 3 was programmed in jsPsych (<https://www.jspsych.org/>), and administered online.

Figure 1. Block and trial procedure. **A.** Each block consisted of eight runs with 4, 8, 12, or 16 trials of the diagnostic task. Mostly new contexts indicated new instructions on 75% of the runs (i.e., green) and mostly old contexts featured new instructions on only 25% of the runs (i.e., blue). The diagnostic task was the same independent of run type. **B-D.** Participants were instructed about the inducer task at the start of each run (**B**, dashed lines), after which they had to perform the diagnostic task (**C**, full lines), followed by one trial of the inducer task (**D**, dotted lines). In the diagnostic task, participants had to respond to the orientation of the font type, while the meaning of the word was irrelevant. The inducer task instructions (and target) could appear above or below the fixation cross (Experiment 1) or in blue or green (Experiment 2), and this contextual feature was associated to the likelihood (75% / 25%) of the instruction being new or old (i.e., recurring). The background of the actual experiment was black, and the black font was white.

Procedure. Participants read the general instructions, after they were seated behind the computer in a dimly lit room and filled in the informed consent form. The instructions explained the general procedure as well as the specific response rules for the diagnostic task. These response rules were the same for all participants: participants had to press left when the stimulus was presented in upright font and right when presented in italic. The instructions further emphasized to use the location information

(Experiment 1) or color font (Experiment 2) of the inducer task instructions to respond as fast as possible to the inducer task. After a first example run with 12 diagnostic trials, the participants read the instructions once more, and completed all experimental blocks of eight runs. Note that the first experimental block was also excluded from analyses because the difference between old versus new instruction pairs could only become apparent from the second block onwards.

Each run started with the inducer task instructions, which were presented on screen above or below the fixation cross (Experiment 1), or in green or blue (Experiment 2), consistent with the inducer task location/color (Figure 1B). In Experiment 1, the first line mentioned whether the task would be presented above or below the fixation cross with the location printed in uppercase letters (i.e., the Dutch translation of "The task will be presented ABOVE"), under which the next two lines indicated the run-specific response mapping (e.g., the Dutch translation of "if wine, press right"). In Experiment 2, instructions were presented centrally but the first line would state whether the task is, and will be presented in green or blue, again denoting the color in uppercase letters. The order of the lines indicating the left or right button response was randomized across runs in both experiments. These instructions remained on the screen for 20 seconds or until the participant pressed the spacebar. After the instructions and a 750 ms interval, the first stimulus of the diagnostic task appeared (Figure 1C). Each stimulus of the diagnostic task was presented until response or for 2000 ms. When the participant responded incorrectly, a red screen appeared for 200 ms. In Experiment 3, participants saw the words "incorrect" or "too slow" (in Dutch), rather than the red screen, to ensure experiment design details were self-explanatory (because Experiment 3 was ran online, and there was no experimenter present to elaborate on instructions concerning the feedback screen). In between trials, there was a random intertrial interval that was between 500 and 1000 ms long. Following the final diagnostic trial of each run, the same intertrial interval was presented after which the fixation cross turned green for another 500 milliseconds. Finally, the inducer task stimulus appeared (Figure 1D) which was also presented in green (or sometimes blue in Experiment 2) and remained on screen until response or for 2000 ms. A new run started 1500 ms after the participant performed the inducer task. The fixation cross remained on screen throughout the entire run.

Analyses. The raw data can be found online on the Open Science Framework (<https://osf.io/tzemg/>) [the experiment code and analysis scripts will be uploaded upon publication]. In line with other studies using this paradigm, the main analyses focused on the diagnostic trials where a *new* instruction was shown. The *old* instructions were considered “filler” items that helped us to create the two contexts, within we could study the implementation of new instructions. Moreover, this purpose aside, we also had no hypothesis regarding the impact of context on instruction-based congruency effects for old instructions. The repeated presentation of these instructions (more for one context than the other), could be hypothesized to result in stronger task representations and thus more interference with the diagnostic task. However, people also had more training in “shielding off” those instructions in previous diagnostic runs, which could be argued to result smaller instruction-based congruency effects.

We used the exact same outlier criteria and analysis pipeline as in our previous study that used a highly similar design (Braem et al., 2017). Specifically, to avoid post-error or post-feedback slowing effects, we removed all trials following an error. Also trials that were part of a run on which participants incorrectly responded to the eventual inducer trial were excluded from analyses, because we cannot ensure that participants successfully encoded the inducer task on those runs. For the primary set of analyses on reaction time (RT), we only considered correct RTs slower than 200 ms. Per experiment, median RT and mean accuracy results were analyzed statistically in repeated-measures analyses of variance (rANOVA) with the two within-subject factors for Congruency (2; congruent versus incongruent) and Context (2; mostly new versus mostly new context). Although not of main interest (see above), we also report the same set of analyses on the old, recurring instructions. Finally, separate follow-up rANOVAs are reported where Experiment is an extra between-subjects factor (2; Experiment 1: location versus Experiment 2: color). All rANOVA and follow-up tests were two-tailed. For the main interaction of interest, we also computed a Bayes factor that signifies the evidence for the alternative hypothesis (i.e., that the congruency effect would be larger for mostly new contexts versus mostly old contexts) against the evidence for the null hypothesis, using the open source statistical program JASP Version 0.13.1 (Love et al., 2015) with a Cauchy prior width of $\sqrt{2}/2$ (i.e., default setting in JASP).

In a secondary set of analyses, we also analysed the encoding times (i.e., the time participants took to read and memorize an instruction), and the inducer task RTs and accuracy. Due to an unfortunate oversight in the programming of Experiment 3, we did not record encoding times and thus could not perform this analysis in Experiment 3. In general, note that these analyses were based on a much smaller set of data points per subject, and were therefore considered underpowered and of secondary importance. However, due to the interesting nature of these measurements, we also investigated the impact of our contextual manipulations there, albeit only across both experiments to obtain sufficient statistical power. Similar to the analysis of the diagnostic task performance, instruction encoding times were only considered for runs where participants were correct on the inducer trial, and only correct RTs slower than 200 ms were included in the analyses. Median RT and mean accuracy analyses were compared using two-tailed t-tests or repeated-measures analyses of variance (rANOVA) with the within-subject factor context and between-subjects factor context type.

Results

Experiment 1

The mean accuracy in Experiment 1 was 95.0% ($SD = 3.6\%$; 79% - 99%) on the diagnostic task and 90.9% ($SD = 7.0\%$; 72% - 100%) on the inducer task.

New task instructions. An analysis of the reaction times of the diagnostic task indicated no significant effects of context or congruency, all $F_s(1, 39) < 1$, but a marginally significant two-way interaction between context and congruency, $F(1, 39) = 2.89, p = .097, \eta_p^2 = .069$ (Figure 2). Follow-up tests indicated a marginally significant congruency effect (7 ms) when the task instructions were presented in the mostly new location, $t(39) = 1.963, p = .057$, but not in the mostly old location (-3 ms), $t(39) = -.642, p = .524$. The Bayes Factor for this interaction between congruency and context was 1.20, suggesting that these data are 1.20 more likely to be observed under the alternative hypothesis. This can be seen as inconclusive evidence for neither the null or alternative hypothesis (Jeffreys, 1961).

The error rates analysis showed a significant congruency effect (2.6 %), $F(1, 39) = 8.75$, $p = .005$, $\eta_p^2 = .183$, indicating more accurate responses on congruent trials. The main effect of context and its interaction with congruency did not reach significance, both $F_s(1, 39) < 1$.

Old task instructions. A similar reaction-time analysis on the old, recurring task instructions indicated a marginally significant effect of congruency (8 ms), $F(1, 39) = 3.18$, $p = .082$, $\eta_p^2 = .075$, but no main effect of context or interaction between context and congruency, both $F_s(1, 39) < 1$ (see Figure 3).

The error rates analysis showed a significant congruency effect (3.4 %), $F(1, 39) = 15.66$, $p < .001$, $\eta_p^2 = .287$, indicating more accurate responses on congruent trials, and a main effect of context, $F(1, 39) = 9.92$, $p = .003$, $\eta_p^2 = .203$, indicating more errors when the instructions were presented on the mostly old location. The interaction between context and congruency did not reach significance, $F(1, 39) = 1.20$, $p = .28$, $\eta_p^2 = .030$.

Figure 2. Mean median reaction times and mean error rates for the diagnostic task in runs where new instructions were presented, in Experiments 1 and 2. The data are depicted for each congruency and context condition

separately. As reported in the results sections, the instruction-based congruency effect for new instructions was larger when the task instructions were presented in a mostly new location (Experiment 1) or font color (Experiment 2), in the reaction times (in the across-experiment analyses and in Experiment 2), and trending in the error rates in Experiment 2. All error bars are ± 1 standard error of the mean (SEM).

Experiment 2

Participants had a mean accuracy of 94.0% ($SD = 3.5\%$; 83% - 99%) on the diagnostic task and 90.4% ($SD = 7.2\%$; 65% - 100%) on the inducer task.

New task instructions. The reaction times of the diagnostic task indicated no significant effects of context or congruency, both $F_s(1, 39) < 1$. However, the two-way interaction between context and congruency was significant, $F(1, 39) = 6.86, p = .012, \eta_p^2 = .150$ (Figure 2). In line with Experiment 1, this indicated that a significant congruency effect (13 ms) was observed when the task instructions were presented in the mostly new font color, $t(39) = 4.890, p < .001$, but not in the mostly old font color (-8 ms), $t(39) = -1.131, p = .265$. The Bayes Factor for the interaction was 6.70, suggesting that these data are 6.70 more likely to be observed under the alternative hypothesis. This is considered substantial evidence for the alternative hypothesis (Jeffreys, 1961).

The error rate analysis did not show a significant congruency effect, $F(1, 39) = 1.55, p = .220, \eta_p^2 = .038$. Similarly, the main effect of context did not reach significance, $F(1, 39) < 1$. However, the two-way interaction between context and congruency was marginally significant, $F(1, 39) = 3.92, p = .055, \eta_p^2 = .091$. Similar to the RT analysis, this showed that a significant congruency effect (1.8%) was observed when the task instructions were presented in the mostly new font color, $t(39) = 3.473, p = .001$, but not when in the mostly old font color (-0.1%), $t(39) = .136, p = .892$.

Old task instructions. We also performed the same reaction time analysis on the old, recurring task instructions. This analysis indicated a significant effect of congruency (9 ms), $F(1, 39) = 6.93, p = .012, \eta_p^2 = .151$, but no main effect of context, $F_s(1, 39) < 1$, or interaction between context and congruency, $F(1, 39) = 2.06, p = .159, \eta_p^2 = .050$ (see Figure 3). If anything, the numeric direction of this non-

significant interaction suggested a larger congruency effect when the instructions were presented in the mostly new font color.

The error rates analysis showed a significant congruency effect (2.3 %), $F(1, 39) = 11.23, p = .002$, $\eta_p^2 = .224$, indicating more accurate responses on congruent trials, but did not show a main effect of context (as Experiment 1 did), $F(1, 39) < 1$. Similarly, the interaction between context and congruency was also not significant, $F(1, 39) < 1$.

Figure 3. Mean median reaction times and mean error rates for the diagnostic task in runs where old instructions were presented, in Experiments 1 and 2. The data are depicted for each congruency and context condition separately. As reported in the results sections, the instruction-based congruency effect did not significantly differ depending on which context the instructions were presented in. All error bars are ± 1 standard error of the mean (SEM).

Across experiment analyses

Diagnostic task – New task instructions. An across-experiment analysis of the reaction times for new task instructions corroborated the above conclusions. Namely, there were no significant effects of

context or congruency, both $F_s(1, 78) < 1$, but the two-way interaction between context and congruency was significant, $F(1, 78) = 9.73$, $p = .003$, $\eta_p^2 = .111$. Specifically, congruent trials (560 ms) were significantly faster responded to than incongruent trials (571 ms), when the task instruction was presented in a context where mostly new task instructions were presented, $F(1, 78) = 20.79$, $p < .001$, $\eta_p^2 = .210$, but not when it was presented in a mostly old context, $F(1, 78) = 1.68$, $p = .199$, $\eta_p^2 = .021$. If anything, the latter congruency effect was numerically reversed (571 ms versus 566 ms). None of the effects interacted with the factor Experiment, all $F_s(1, 78) < 1.29$, $p_s > .258$, suggesting these results did not differ between experiments. The Bayes Factor for this interaction between congruency and context was 20.65, suggesting that these data are 20.65 more likely to be observed under the alternative hypothesis. This is considered strong evidence for the alternative hypothesis (Jeffreys, 1961).

Finally, the across-experiment error rates analysis showed a significant congruency effect, $F(1, 78) = 9.73$, $p = .003$, $\eta_p^2 = .111$, indicating more accurate responses on congruent trials. The main effect of context and its interaction with congruency did not reach significance, both $F_s(1, 78) < 1$. However, the three-way interaction between Experiment, context, and congruency was significant, $F(1, 78) = 4.45$, $p = .038$, $\eta_p^2 = .054$, consistent with the observation that a (marginally significant) two-way interaction between context and congruency was only observed in Experiment 2, but not Experiment 1.

Diagnostic task – Old task instructions. The across-experiment analysis of reaction times in the diagnostic task for old task instructions showed a significant main effect of congruency, $F(1, 78) = 9.88$, $p = .002$, $\eta_p^2 = .112$ (congruent = 554 ms; incongruent 562 ms), which did not interact with context $F(1, 78) = 0.95$, $p = .333$, $\eta_p^2 = .012$. There was no significant effect of context, $F(1, 78) = 0.04$, $p = .849$, $\eta_p^2 = .000$, and none of the effects interacted with the factor Experiment, all $F_s(1, 78) < 1.64$, $p_s > .203$, suggesting these results did not differ between experiments.

A similar analysis in the error rates showed a significant congruency effect, $F(1, 78) = 26.73$, $p < .001$, $\eta_p^2 = .256$, indicating more accurate responses on congruent trials (96.2 %) than incongruent trials (93.4 %), as well as a main effect of context, $F(1, 78) = 4.97$, $p = .029$, $\eta_p^2 = .060$, showing better performance when the instructions were presented in a new (95.1 %), as opposed to an old context

(94.5 %). Importantly, the factors congruency and context did not interact, $F(1, 78) = 1.51, p = .223, \eta_p^2 = .019$, nor did any of the effects interact with the factor Experiment, all $F_s(1, 78) < 2.26, p_s > .136$.

Figure 4. Mean median instruction encoding times for Experiments 1 and 2. Reaction times are depicted for each novelty and context condition separately. As reported in the results sections, an across-experiment analysis revealed that participants were faster to encode new task instructions when the instruction was presented in a mostly new (as opposed to old) location (Experiment 1) or font color (Experiment 2). All error bars are ± 1 standard error of the mean (SEM).

Instruction encoding times. Across both experiments, a rANOVA with context (mostly new versus mostly old) and instruction novelty (new versus old) showed a main effect of instruction novelty, $F(1, 78) = 7.48, p = .008, \eta_p^2 = .088$, as well as an interaction with context, $F(1, 78) = 6.82, p = .011, \eta_p^2 = .080$, but no main effect of context, $F(1, 78) < 1$. When tested for new and old instructions separately, the encoding times for new instructions were significantly faster in contexts where mostly new instructions were presented (3152 ms) than contexts where mostly old instructions were presented (3369 ms), $F(1, 78) = 5.53, p = .021, \eta_p^2 = .066$ (see Figure 4). In contrast, the encoding times for old instructions were not significantly faster in contexts where mostly old instructions were presented (2832 ms) than contexts where mostly new instructions were presented (3016 ms), $F(1, 78) = 2.63, p = .109, \eta_p^2 = .033$. These effects did not further interact with Experiment, all $F_s(1, 78) < 1$.

Inducer task. Across both experiments, a rANOVA with context (mostly new versus mostly old) and instruction novelty (new versus old) showed a main effect of instruction novelty, with better performance for old, recurring instructions, in both RTs (610 ms versus 625 ms), $F(1, 78) = 3.93, p =$

.051, $\eta_p^2 = .048$, and accuracy (91.8% versus 89.7%), $F(1, 78) = 5.92$, $p = .017$, $\eta_p^2 = .071$. None of the other main effects or interactions reached significance, all $F_s(1, 78) < 2.24$, $ps > .138$, $\eta_p^2s < .028$.

Notably, while instruction context affected instruction encoding times and the interference effect in the diagnostic task, we found no evidence for an effect on the actual inducer task. Possibly, in line with the reasoning and analysis that was used in our previous study (Braem et al., 2017), this contextual modulation can only be observed after sufficient training by looking at the second half of the experiment. Therefore, although post-hoc in the present study, we also performed the same analysis by only focusing on second half of each experiment. Indeed, when focusing on the second half only, we did observe an interaction between instruction novelty and context, $F(1, 78) = 16.05$, $p < .001$, $\eta_p^2 = .171$ (and a marginally significant main effect of instruction novelty in RTs, $F(1, 78) = 3.21$, $p = .077$, $\eta_p^2 = .040$). None of the other main effects or interactions reached significance, all $F_s(1, 78) < 1.01$, $ps > .318$, $\eta_p^2s < .014$. A similar analysis on the error rates showed no significant effects, all $F_s(1, 78) < 1.96$, $ps > .165$, $\eta_p^2s < .026$. When we further unpacked this interaction effect in the RTs, this seemed to be driven by both faster RTs for the execution of new instructions in mostly new contexts as opposed to mostly old contexts (606 ms versus 639 ms), $F(1, 78) = 3.87$, $p = .052$, $\eta_p^2 = .047$, as well as faster RTs for the execution of old instructions in mostly old contexts as opposed to mostly new contexts (577 ms versus 622 ms), $F(1, 78) = 9.35$, $p = .003$, $\eta_p^2 = .107$.

Experiment 3 (online replication of Experiment 2)

Aiming to replicate the main findings of Experiment 2 in an online study (as Experiment 2 showed the strongest results of both the above experiments), we programmed the experiment anew and collected data online from 77 new subjects. Apart from the feedback screen which now displayed the words “too slow” or “incorrect” (rather than colouring the whole screen red as in Experiments 1 and 2), and the unfortunate oversight that we forgot to program the registration of encoding times, Experiment 3 was virtually identical to Experiment 2. Similar to Experiment 2, participants had a mean accuracy of 94.4% ($SD = 3.2\%$; 84.5% - 99%) in the diagnostic task and 90.0% ($SD = 8.9\%$; 57.5% - 100%) in the inducer task.

New task instructions. The rANOVA results of reaction time revealed a significant main effect of congruency ($F(1, 76) = 43.88, p < .001, \eta_p^2 = .366$), indicating that participants had faster responses on congruent trials (i.e., 21 ms). Different from Experiment 2, the interaction between context and congruency was not significant, $F(1, 76) = 0.74, p = .393, \eta_p^2 = .010$. The Bayes Factor for this interaction was 0.29, suggesting that these data are 3.44 ($= 1 / 0.29$) more likely to be observed under the null hypothesis (see Figure 5). However, also different from Experiment 2, there was a main effect of context ($F(1, 76) = 11.85, p < .001, \eta_p^2 = .135$), in that participants were faster in the mostly new context (i.e., 12 ms).

The rANOVA results of error rates replicated those of Experiment 2, in that they revealed no main effect of context ($F(1, 76) = 2.75, p = .101, \eta_p^2 = .035$), but a significant congruency effect (-2.8%, $F(1, 76) = 51.04, p < .001, \eta_p^2 = .402$). Crucial to our hypothesis, there was a significant interaction between context and congruency, $F(1, 76) = 8.37, p < .001, \eta_p^2 = .099$, indicating that the congruency effect was smaller in the mostly old context (see Figure 5). The Bayes factor of this interaction was 11.67, providing strong evidence for the alternative hypothesis. The follow-up tests revealed that a significant congruency effect was observed in both the mostly new context (3.7%), $t(76) = 9.53, p < .001$, and the mostly old context (1.8%), $t(76) = 3.00, p < .001$.

Figure 5. Mean median reaction times and mean error rates for the diagnostic task in runs where new instructions were presented, in Experiments 3. The data are depicted for each congruency and context condition separately. As reported in the results sections, the instruction-based congruency effect for new instructions was larger when the task instructions were presented in a mostly new font color, but only in the error rates, replacing the error rate pattern of Experiment 2. Different from Experiment 2, participants in Experiment 3 showed an overall reaction time benefit in the mostly new versus mostly old context when retaining new task instructions. All error bars are ± 1 standard error of the mean (SEM).

Old task instructions. We performed the same rANOVA to RT and error rate data from trials with old task instructions. The results revealed a significant congruency effect (22 ms), $F(1, 76) = 47.86$, $p < .001$, $\eta_p^2 = .386$, indicating the response time was also smaller in incongruent trials with old instructions. The main effect of context and its interaction with congruency did not reach significance, both $F_s(1, 76) < 1$. The error rate results also revealed a significant congruency effect (3.3%), $F(1, 76) = 52.17$, $p < .001$, $\eta_p^2 = .407$, but no main effect of context or interaction with congruency, both $F_s(1, 76) < 1$.

Inducer task. We performed a rANOVA with context (mostly new versus mostly old) and instruction novelty (new versus old) on RT and error rate data in the inducer task. The RT results revealed that

participants were faster on trials with old, recurring task instructions (662 ms versus 638 ms), as suggested by a significant main effect of instruction novelty $F(1, 76) = 8.04, p = .006, \eta_p^2 = .096$. Although participants had more accurate responses on old instructions (90.2% versus 91.6%), this difference did not reach significance, $F(1, 76) = 2.68, p = .106, \eta_p^2 = .034$. The remaining effects were not significant, with all $F_s(1, 76) < 1$. Much like the across-experiment analyses, we also performed the same rANOVA to the second half of RT and error rate data. However, none of the effects reached significance, all $F_s(1, 76) < 2.59, p_s > .0112, \eta_p^2 < .033$.

Discussion

The present study was set up to determine whether the uniquely human ability to quickly transform a verbal instruction into an action plan, is one that can bind to, and be regulated by, the context in which it was learned. To this end, we presented instructions in a context where new instructions were presented more versus less often, and measured their automatic impact on behavior depending on this context. The results suggest that people indeed form more potent task representations when instructions are presented in contexts with mostly new instructions, relative to contexts with mostly recurring instructions, as measured by the automatic impact of these instructions on an irrelevant, “diagnostic” task. Across-experiment analyses of the two lab experiments further showed that people were also faster to encode a new instruction, when presented in a mostly new context. These findings are consistent with a number of previous observations indicating that abstract goals or cognitive control functions can bind to contextual features in our environment, allowing for the recruitment of these functions upon contextual demand (Abrahamse et al., 2016). A more recent line of studies also extended this observation to the domain of “cognitive flexibility” – the ability to quickly reconfigure our task sets to meet new goals (for a review, see Braem & Egner, 2018). For example, in a design where specific task items were associated to more task switches than others, Chiu and Egner (2017) demonstrated that the efficiency with which one switches between tasks (as measured by the task switch cost) is modulated by the degree to which the current item is associated to “the act of task switching”. More recently, Whitehead, Pfeuffer, and

Egner (2020) showed this learning of stimulus-flexibility associations can even occur after a single instance of experiencing a given stimulus under task repetition versus task switching conditions. The present study further extends this idea that cognitive flexibility can be regulated by the context, by showing context-specific cognitive flexibility under conditions where people do not just have to alternate between two or three tasks (Chiu and Egner, 2017; Whitehead et al., 2020), but between many new tasks, consistent with the idea of regulating a general updating threshold (Braem & Egner, 2018; Dreisbach & Fröber, 2018; Goschke, 2003; Goschke & Bolte, 2014). Potentially, this use of many new tasks in our paradigm also helped people get in a flexible mode faster, as previous training studies have shown that task variability can help in training cognitive flexibility (e.g., Karbach & Kray, 2009; Sabah, Dolk, Meiran, & Dreisbach, 2019).

We believe the present findings bring evidence to the idea that we can flexibly “learn when to learn” (Wang, 2021) or, in other words, when to allocate resources for the purpose of encoding task instructions. Pereg and colleagues (2019) already demonstrated that the ability to encode new task instructions is one that can be trained (i.e., as shown by improved inducer task performance after working memory training). However, the current findings suggest that this ability can also bind to contextual features in one’s environment. Namely, by presenting new instructions mostly on a particular location or in a specific font color, people potentially associated the act of encoding new task instructions to these specific features, as suggested by the stronger task representations (as measured by the interference effects) for new tasks that were instructed in these contexts. We further believe this process can be relatively automatic. This would also be consistent with our observation that participants were actually faster to encode new task instructions in mostly new contexts, without the cost of reduced task performance on the inducer task (and even improved performance in the second half of the experiment). In contrast, Cole, Patrick, Meiran, and Braver (2018) have shown that participants usually show reduced task performance when they take less time to encode task instructions.

Although we had no prior hypothesis about which context manipulation would be most effective, the data suggest that, in this specific task, the color context cue (Experiments 2 and 3) was more effective than the location context cue (Experiment 1). Specifically, the evidence was inconclusive in the location

cue experiment, according to the Bayesian analyses, and the across-experiment analyses showed that there was a between-experiment difference in the congruency-by-cue interaction in the error rates, again favoring the color cue manipulation over the location cue manipulation. Our post-hoc interpretation of this finding is that, in this task, font color was presumably more difficult to ignore when processing word meaning to encode the instructions. However, we should be cautious when interpreting this finding, as this experiment difference was not predicted and the evidence for it not very strong.

An alternative interpretation could be that people actually experienced more difficulty with encoding new task instructions in contexts where they already associated other tasks with (i.e., the four recurring tasks). Although the different tasks did not use interfering stimuli, it is possible that contexts that are already strongly associated to certain tasks, require the subject to break down those event files first before new tasks can be encoded (e.g., Hommel et al., 2001). This way, rather than contexts being more or less associated to the encoding of new instructions, they would be associated to the higher or lower requirement of having to break down (or shield off) old task instructions. However, we believe this interpretation is less plausible, because this would also predict bigger interference effects (i.e., stronger task representations) on mostly old contexts (versus mostly new contexts), when studying diagnostic task performance for old instructions. This was not case (if anything, the numerical trend suggested the opposite pattern).

We did not make explicit predictions for the old instructions. However, as outlined in the previous paragraph, mostly old contexts may have primed the specific old instruction with which they were observed before, and thereby facilitated the processing of old instructions. Interestingly, one could also have expected increased interference effects on mostly new contexts for old instructions: If mostly new contexts broadly triggered the need to encode new task instructions this could arguably have facilitated the processing of old instructions, too. As illustrated by these two opposing predictions, it is not straightforward to predict how old instructions could or should have been affected by our manipulation, given that they can rely on both the current implementation strategy, as well as the memory trace of their previous instruction presentation and action execution. A conclusive interaction in either direction could have arbitrated between one of the two post-hoc explanations above, but our data showed no

significant interaction in any of the three experiments, leaving it unclear whether either or both mechanisms were at play.

It is noteworthy that the congruency effects for old instructions were not reliably larger than those for new instructions. Similarly, in a previous study where we compared interference effects from new instructions versus overtrained representations (Braem, Deltomme, & Liefoghe, 2019), we even observed smaller congruency effects for the latter. This raises the question why a repeated instruction, typically thought to result in more or stronger memories, did not show more interference with the diagnostic task. We believe the reason for this may have to do with the type of memory representations they rely on. Specifically, while new instructions are thought to require the explicit implementation to carry out the instruction (Wenke et al., 2009; Whitehead & Egner, 2018), and the active retaining of procedural information in working memory (e.g., Abrahamse et al., 2022; Brass et al., 2017; Liefoghe et al., 2012; Meiran & Cohen-Kadosh, 2012), old instructions can typically rely on long-term memory traces and episodic memories of previous action implementation. It has been argued that only the latter type of memory representations allows for context-sensitive behavior (e.g., Braem et al., 2017), which makes them more applicable for inducer-task-specific activations, and could be a reason why they do not necessarily result in stronger interference effects with the diagnostic task.

Finally, while we observed the stronger interference effect from new task instructions in the mostly new context in both the original across-experiment analyses of the two lab-based experiments, as well as the online replication attempt of Experiment 2 (i.e., Experiment 3), the effects were more pronounced in reaction times in the former, and only observed in the error rates in the latter. Instead of showing a stronger interference effect in the mostly new context, participants in Experiment 3 were overall faster to retain new task instructions in this context. Although this difference was not expected, we offer two post-hoc interpretations as to why this difference might have occurred. First, participants showed a comparable accuracy across Experiments 2 and 3, but were slower in the latter, online experiment, suggesting they used a slightly different (speed-accuracy) strategy while performing the task. The original Experiments 1 and 2 were executed in the lab, where the instructions were also repeated in person, and participants could ask for further explanation (before starting the experiment). Moreover,

the feedback screen for responding too slow or incorrect was also much more salient in Experiment 2 (i.e., the whole screen was colored red for 200 ms). Therefore, participants were likely less focused on speed in Experiment 3, and therefore showed less of a modulation of the interference effect in reaction times. Second, the participants in Experiment 3 showed a general reaction time benefit for retaining new instructions in the mostly new context. This suggests that participants were more optimal in retaining the new task instructions in the mostly new context, as this overall reaction time cost has been associated to the dual task load of retaining the task instruction (Abrahamse et al., 2022). Although speculative, the relaxed emphasis on speed possibly allowed them to rely on this more optimal strategy to retain new task instructions – a strategy which they associated to the mostly new context.

In conclusion, we believe the current study demonstrates how we can flexibly regulate (the need for) instruction encoding by relying on contextual information. In line with recent theories on cognitive control (Abrahamse et al., 2016; Braem & Egner, 2018; Chiu & Egner, 2019; Egner, 2014; Eisenreich et al., 2017; Lieder et al., 2018), these findings indicate that abstract, higher-order functions of behavior, like the encoding of new task instructions, are processes that can be regulated by contextual features in our environment. This context-sensitivity of cognitive control functions also fits with the observation that cognitive control functions can be trained, but rarely show transfer (Melby-Lervåg, Redick, & Hulme, 2016; Simons et al., 2016). Although this might seem counterproductive at first, this same context-sensitivity could be what allows cognitive control functions to meet specific developmental needs (Doebel, 2020) and cultural demands (Braem & Hommel, 2019).

References

- Abrahamse, E., Braem, S., De Houwer, J., & Liefoghe, B. (2022). Tenacious instructions: How to dismantle newly instructed task rules?. *Journal of Experimental Psychology: General*. In press.
- Abrahamse, E. L., Braem, S., Notebaert, W., & Verguts, T. (2016). Grounding cognitive control in associative learning. *Psychological Bulletin*. in press.
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5), 193-200.
- Braem, S., Bugg, J. M., Schmidt, J. R., Crump, M. J., Weissman, D. H., Notebaert, W., & Egner, T. (2019a). Measuring adaptive control in conflict tasks. *Trends in Cognitive Sciences*, 23, 769-783.
- Braem, S., Deltomme, B., & Liefoghe, B. (2019b). The instruction-based congruency effect predicts task execution efficiency: Evidence from inter-and intra-individual differences. *Memory & Cognition*, 47(8), 1582-1591.
- Braem, S., & Egner, T. (2018). Getting a grip on cognitive flexibility. *Current Directions in Psychological Science*, 27(6), 470-476.
- Braem, S., & Hommel, B. (2019). Executive functions are cognitive gadgets. *Behavioral and Brain Sciences*, 42, e173-e173.
- Braem, S., Liefoghe, B., De Houwer, J., Brass, M., & Abrahamse, E. L. (2017). There are limits to the effects of task instructions: Making the automatic effects of task instructions context-specific takes practice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(3), 394.
- Brass, M., Liefoghe, B., Braem, S., & De Houwer, J. (2017). Following new task instructions: Evidence for a dissociation between knowing and doing. *Neuroscience & Biobehavioral Reviews*, 81, 16-28.
- Bugg, J. M., & Crump, M. J. (2012). In support of a distinction between voluntary and stimulus-driven control: A review of the literature on proportion congruent effects. *Frontiers in psychology*, 3, 367.
- Chiu, Y. C., & Egner, T. (2019). Cortical and subcortical contributions to context-control learning. *Neuroscience & Biobehavioral Reviews*, 99, 33-41.

- Cole, M. W., Braver, T. S., & Meiran, N. (2017). The task novelty paradox: Flexible control of inflexible neural pathways during rapid instructed task learning. *Neuroscience & Biobehavioral Reviews*, 81, 4-15.
- Cole, M. W., Patrick, L. M., Meiran, N., & Braver, T. S. (2018). A role for proactive control in rapid instructed task learning. *Acta psychologica*, 184, 20-30.
- Crump, M. J., Gong, Z., & Milliken, B. (2006). The context-specific proportion congruent Stroop effect: Location as a contextual cue. *Psychonomic bulletin & review*, 13(2), 316-321.
- Deacon, T. W. (1997). *The symbolic species*. Norton.
- Doebel, S. (2019). Rethinking executive function development. *Perspectives on Psychological Science*.
- 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.
- Egner, T. (2014). Creatures of habit (and control): a multi-level learning perspective on the modulation of congruency effects. *Frontiers in psychology*, 5, 1247.
- Eisenreich, B. R., Akaishi, R., & Hayden, B. Y. (2017). Control without controllers: toward a distributed neuroscience of executive control. *Journal of cognitive neuroscience*, 29(10), 1684-1698.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39, 175-191.
- Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54, 493-503.
- Goschke, T. (2003). Voluntary action and cognitive control from a cognitive neuroscience perspective. *Voluntary action: Brains, minds, and sociality*, 49-85. Oxford University Press.
- Goschke, T., & Bolte, A. (2014). Emotional modulation of control dilemmas: The role of positive affect, reward, and dopamine in cognitive stability and flexibility. *Neuropsychologia*, 62, 403-423.
- Goschke, T., Kuhl, J. (1993). Representation of intentions: Persisting activation in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1211-1226.

- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24(5), 849-87.
- 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.
- Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods*, 42(3), 643-650.
- Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLoS Computational Biology*, 14(4), e1006043.
- Liefooghe, B., De Houwer, J., & Wenke, D. (2013). Instruction-based response activation depends on task preparation. *Psychonomic Bulletin & Review*, 20(3), 481-487.
- Liefooghe, B., Wenke, D., & De Houwer, J. (2012). Instruction-based task-rule congruency effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1325-1335.
- Mayr, U., & Bryck, R. L. (2007). Outsourcing control to the environment: effects of stimulus/response locations on task selection. *Psychological Research*, 71(1), 107-116.
- Meiran, N., Liefooghe, B., & De Houwer, J. (2017). Powerful instructions: Automaticity without practice. *Current Directions in Psychological Science*, 26(6), 509-514.
- 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.
- 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.
- Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, 25(9), 757-775.
- Pereg, M., Shahar, N., & Meiran, N. (2019). Can we learn to learn? The influence of procedural working-memory training on rapid instructed-task-learning. *Psychological research*, 83(1), 132-146.

- 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.
- Simons, D. J., Boot, W. R., Charness, N., Gathercole, S. E., Chabris, C. F., Hambrick, D. Z., & Stine-Morrow, E. A. (2016). Do “brain-training” programs work?. *Psychological Science in the Public Interest*, 17(3), 103-186.
- Schmidt, J. R. (2019). Evidence against conflict monitoring and adaptation: An updated review. *Psychonomic Bulletin & Review*, 26(3), 753-771.
- Stevens, M., Lammertyn, J., Verbruggen, F., & Vandierendonck, A. (2006). Tscope: AC library for programming cognitive experiments on the MS Windows platform. *Behavior research methods*, 38(2), 280-286.
- Verguts, T., & Notebaert, W. (2008). Hebbian learning of cognitive control: dealing with specific and nonspecific adaptation. *Psychological review*, 115(2), 518.
- Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences*, 38, 90-95.
- Whitehead, P. S., & Egner, T. (2018). Frequency of prospective use modulates instructed task-set interference. *Journal of Experimental Psychology: Human Perception and Performance*, 44(12), 1970.
- Whitehead, P. S., Pfeuffer, C. U., & Egner, T. (2020). Memories of control: One-shot episodic learning of item-specific stimulus-control associations. *Cognition*, 199, 104220.