Distractor-specific control adaptation in multidimensional environments

Davide Gheza & Wouter Kool

Department of Psychological and Brain Sciences,

Washington University in St. Louis

Acknowledgements:

We would like to thank Michael Freund, Roselyne Chauvin, Todd Braver, and Julie Bugg for helpful conversations. We also would like to thank the members of the Control and Decision Making Lab for their advice and assistance, including Eric Gruber and Jack Dolgin for their contributions to early versions of the task. This work was supported by a Multi-University Research Initiative grant (ONR/DoD N00014-23-1-2792) to WK.

Abstract

Goal-directed behavior requires humans to constantly manage and switch between multiple, independent, and conflicting, sources of information. Conventional cognitive control tasks, however, only feature one task and one source of distraction. Therefore, it is unclear how control is allocated in multidimensional environments. To address this question, we developed a multi-dimensional task-set interference paradigm, in which people need to manage distraction from three independent dimensions. We use it to test whether people adapt to prior conflict by enhancing task-relevant information or suppressing task-irrelevant information. Three experiments provided strong evidence for the latter hypothesis. Moreover, control adaptation was highly dimension specific. Conflict from a given dimension only affected processing of that same dimension on subsequent trials, with no evidence for generalization. A new neural network model shows that our results can only be simulated when including multiple independent conflict-detector units. Our results call for an update to classic models of cognitive control, and their neurocomputational underpinnings.

Introduction

From the ever-present lure of phone notifications to the cacophony of urban traffic, daily life requires us to manage many competing sources of information. This poses a challenge to our ability for *attentional control*. To achieve goals, we have to attend to sources that are relevant for our task, but we also have to determine how to process other sources. For example, at a talk, you should focus on the speaker, but complete suppression of the surroundings may lead you to miss the insightful observation your neighbor whispers in your ear. Even though we face a constant need for multidimensional attentional control, decades of research on this topic has mostly considered situations with a single source of distraction. Little is known about how attentional control is deployed in more complex environments. Here, we combine a new multidimensional attentional control task with computational modeling to address this question.

Psychologists study attentional control using response interference paradigms. In these tasks, one stimulus dimension signals the correct response, while a second one provokes a response tendency that is either congruent or incongruent with the former. For example, in the Stroop task (Stroop, 1935) participants name the color in which a word is printed. Responses on incongruent trials (e.g., “green” printed in red ink) are slower and less accurate, indicating a need for attentional control (Cohen et al., 1990; MacLeod, 1991). In other prominent tasks, conflict is achieved by presenting a target stimulus flanked by distractors (Flanker task; Eriksen & Eriksen, 1974), or by manipulating whether the correct response and the irrelevant spatial location of a stimulus are on the same side (Simon task; Simon & Rudell, 1967).

Data from these tasks demonstrate that humans configure attentional control based on prior demands. Specifically, people become less susceptible to distractors after experiencing conflict, a phenomenon known as the congruency sequence effect (Gratton et al., 1992). This effect is commonly interpreted to reflect an online adjustment of control strategies (Gratton et al., 1992), but its mechanisms are still debated. Some argue that prior conflict triggers increased processing of task-relevant information (Egner & Hirsch, 2005), but others have suggested that it results in suppression of distractor information (Danielmeier et al., 2011; Ritz & Shenhav, 2022). These two forms of control adjustments have been argued to produce similar behavior, because enhanced target processing and suppressed distractor processing should generally increase evidence in favor of the correct response (Botvinick et al., 2001).

Most models of the congruency sequence effect do not commit to a specific form of adaptation. According to the most influential account by Botvinick and colleagues (2001), conflict yields both increased target processing and decreased attention to the distractor. Building on this model, Egner (2008) proposed the existence of two independent conflict-control loops to explain adaptation during a task that combined Stroop and Simon interference. One of these loops detects stimulus-level conflict (i.e. between representations of color) and the other response-level conflict (i.e., between response options). One loop affects target processing while the other suppresses distraction, resulting in a mixture of both. Finally, a class of models proposes that conflict changes the connections between stimulus and task representations (Blais et al., 2007; Verguts & Notebaert, 2008). For example Verguts and Notebaert (2008) proposed a Hebbian-learning model of adaptation, by which increased conflict strengthens associations between task and stimulus representations, also implementing both enhancement and suppression.

Unfortunately, there is a lack evidence to adjudicate between these theories. Classic interference paradigms are unable to address this question, because target and distractor processing are

experimentally entangled. Neuroimaging research has also not provided a definitive answer. Early work suggested conflict results in enhanced processing of task-relevant information in posterior regions (Egner & Hirsch, 2005), but suppression of distractors activity has also been reported (Danielmeier et al., 2011). This heterogeneity in findings may derive from the complex or imprecise relationships between neuroimaging techniques and cognitive mechanisms (Ritz & Shenhav, 2022). More recently, Ritz and Shenhav (2022) developed a behavioral task in which target and distractor information varied parametrically, and found that control adaptation affects distractor processing. However, they only used one distractor and one target dimension, and so it remains unclear how control adapts in multidimensional tasks.

There are several distinct ways in which control adaptation can be implemented in multidimensional environments. If it affects distractors, there are two possibilities. Distractor adaptation could act at a global level, or it could be managed individually for each source of irrelevant information. In the case of global distractor adaptation, conflict from any source reduces all subsequent distractor processing. This algorithm, which is part of classic models of conflict monitoring (Botvinick et al., 2001; Egner, 2008), is efficient because it requires just one conflict monitor (or one for each conflict "type"; Egner, 2008). In contrast, in the case of distractor-specific adaptation, conflict from a given source only reduces subsequent processing of the same distractor. This is less efficient but more optimal. For instance, it could suppress one (previously incongruent) distractor, while enhancing another (previously congruent) one. In addition, prior conflict may trigger target processing, simplifying the issue of how to allocate control.

Interestingly, some research suggests that control adaptation is conflict-type specific (Braem et al., 2014). For example, when participants perform a task with two types of interference (e.g., Simon and Stroop), conflict from one source of distraction only reduces susceptibility to that distraction. These findings are revealing but limited, because these two sources yield types of conflict that require different forms of adaptation. For example, resolving interference stemming from the side on which the stimulus is presented (e.g., response conflict in a Simon task) requires qualitatively different changes in attentional control than adapting to conflict from word reading (stimulus conflict in a Stroop). It remains unknown how conflict is resolved when multiple distractors generate conflict of the same type (i.e., at the same representational level).

To address these questions, we developed a multi-dimensional task-set interference paradigm, based on a task by Liston and colleagues (2006). On each trial, participants attend to one of four stimulus dimensions, with the remaining dimensions acting as distractors. Each of these three distractors are either compatible or incompatible with the correct response, resulting in a parametric manipulation of congruency. Like prior tasks, participants need to attend to a single cued dimension. However, the automaticity of processing non-cued dimensions is not driven by long-term training or exposure (as the case with word reading in the Stroop task), but because each dimension periodically signals the correct response (Hazeltine et al., 2011; Monsell et al., 2003). Critically, the interference generated from each dimension plays out at the same (stimulus) level.

The task allows us to measure not only the parametric effect of conflict on performance, but also the individual contribution of each distractor dimension. Most importantly, we can use it to measure how control adapts in multidimensional environments. Because each dimension periodically acts as a distractor, the task allows us to determine how conflict from one dimension triggers attentional control adaptation towards itself and all other dimensions. This provides a particularly granular insight into how humans configure attentional control.

Across three task versions, we find striking evidence for dimension-specific control adaptation. Whenever a dimension generated conflict, it only affected subsequent processing of the same dimension. This is not compatible with a control mechanism that enhances task relevant information or that suppresses distraction globally. Rather, this suggests that conflict-driven control adaptation operates in parallel on task-irrelevant information pathways (Lindsay & Jacoby, 1994). Using computational modeling, we show that these results are incompatible with classic models, in which a single monitoring unit detects conflict at the response level (where evidence is aggregated across dimensions). Instead, we find that our results can be captured by including multiple independent conflict detectors for each dimension. This work improves our understanding of how people regulate attentional control in complex environments and provide techniques for studying it.

Results

We conducted three experiments, each with a different variant of the MULtidimensional Task-set Interference paradigm (MULTI; Figure 1A). In Experiment 1, each trial included two stimuli, presented side-by-side. These stimuli differed across four dimensions: their color (blue or orange), shape (oval or rectangle), edge (fully bounded or dashed), and motion of dots inside the object (upward or downward). A centrally located cue dictated which dimension was relevant on a given trial (C for color, S for shape, E for edge, or M for motion). For each of the four cues, participants were told to select the object with a particular target feature (left or right key press). For example, a participant could have been instructed to find the blue object if the task cue was C, the oval when the cue was S, the object with the fully bounded edge if the cue was E, and the object with the upwards moving dots if it was M. Target features were randomly assigned for each participant and remained stable throughout the entire session. The sequence of cues was pseudorandomly generated, cycling through each of the four tasks in random order with each cue repeating between 3 to 5 times (Figure 1C).

Because features are randomly assigned to the stimuli, trials naturally involve different levels of congruency (Figure 1B). Any of the three non-cued dimensions could be congruent if it primed the correct response (i.e., the target features of the cued and a non-cued dimension are on the same side), or incongruent if it primed the opposite response. Therefore, the level of congruency varied from 0 (fully incongruent) to 3 (fully congruent). By independently varying the congruence of each non-cued dimension, the task allows us to comprehensively investigate congruency sequence effects both within and across non-cued dimensions.

We provide access to all data, analysis scripts, and simulation code, on the OSF (https://osf.io/j9fxc/).

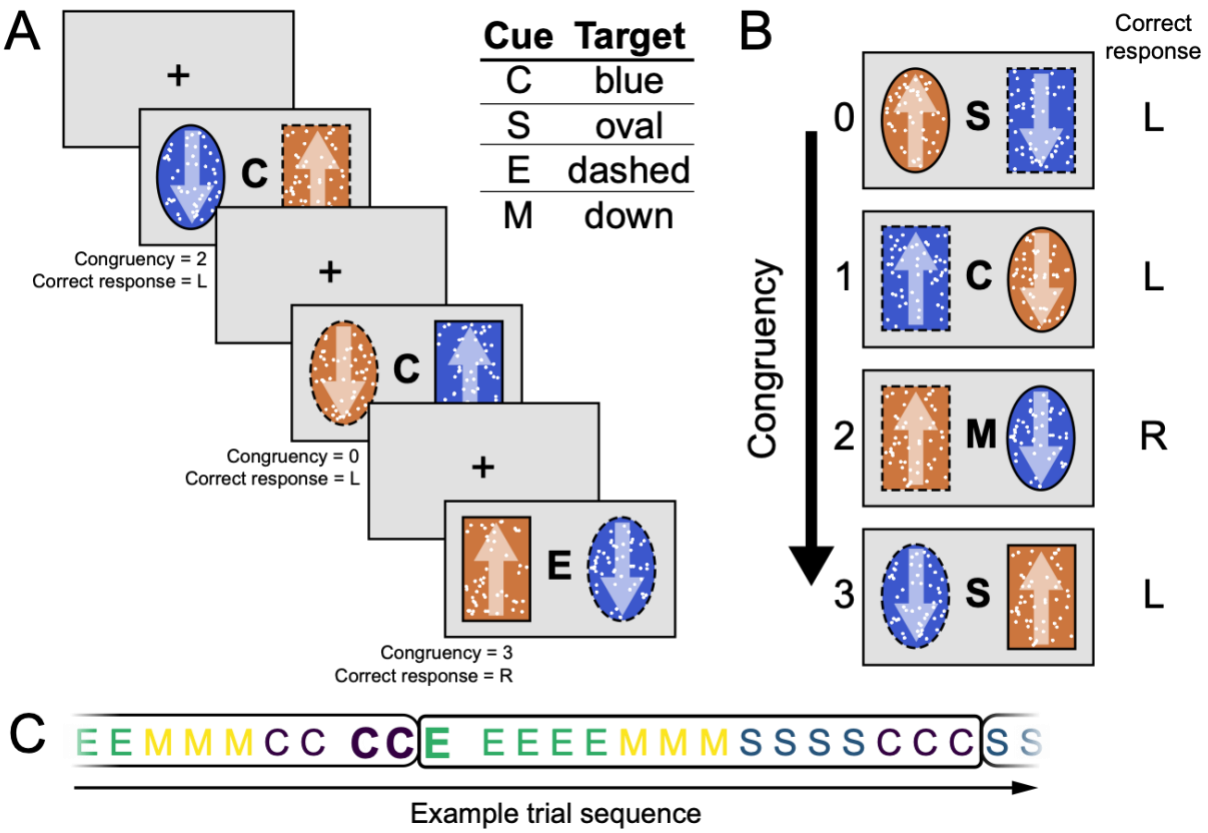

Figure 1. MULti-dimensional Task-set Interference task. a, Sequence of three example trials. On each trial, two objects are presented side-by-side, that differ across four dimensions: the color (orange or blue), the shape (oval or rectangular), the edge (solid or dashed), and the direction of the dot motion (up or down; symbolized by white arrows, not shown in the experiment). Participants are instructed to pay attention to only one 'task' dimension cued by the letter between the stimuli. For each dimension, participants need to find a target feature, as instructed at the beginning of the experiment, and select the corresponding object (left or right response). In the example reported here, the target features are the blue color, the oval shape, the dashed edge, and the downward motion. All stimuli remain on screen until a response is given. Response deadline is set to 1.5s, and the inter trial interval to 1s, during which a white fixation cross is presented. **b,** Each non-cued dimension can be congruent when it primes the same side as where the cued dimension's target feature is located. Therefore, congruency is parametrically manipulated from level 0 (all three non-cued dimensions prime the wrong side) to level 3 (all three non-cued dimensions prime the correct side). **c,** Example trial sequence in which the three example trials in panel **a**, represented by the cue letters C C E in bold, are embedded. Sequences are pseudo-randomly generated, with cued dimensions repeating 3 to 5 times, before a switch trial occurs. All dimensions are cued within 4 switches (encircled sequences), in random order. We refer to a sequence of repeat trials as 'block' (color coded sequences).

The number of congruent dimensions predicts task performance

The data from Experiment 1 revealed the MULTI imposes parametric control demands on participants. We found that the number of congruent dimensions predicted both response time and accuracy (Figure 2). Participants were fastest when the stimulus was fully congruent, and became slower with each decrease in the number of congruent dimensions (M_3 = 649 ms, M_2 = 666 ms, M_1 = 685 ms, M_0 = 704 ms). Analogously, participants showed highest accuracy for fully congruent stimuli, which decreased with each increase in the number of congruent dimensions (M_3 = 98.6%, M_2 = 97.4%, M_1 = 95.3%, M_0 = 91.5%). For both measures, a generalized Bayesian linear regression (see Methods) revealed a negative linear effect of congruency on response time

($ERs_{congruency<0}$ = Inf), which was corroborated by comparison of scores on adjacent congruence levels (ERs = Inf).

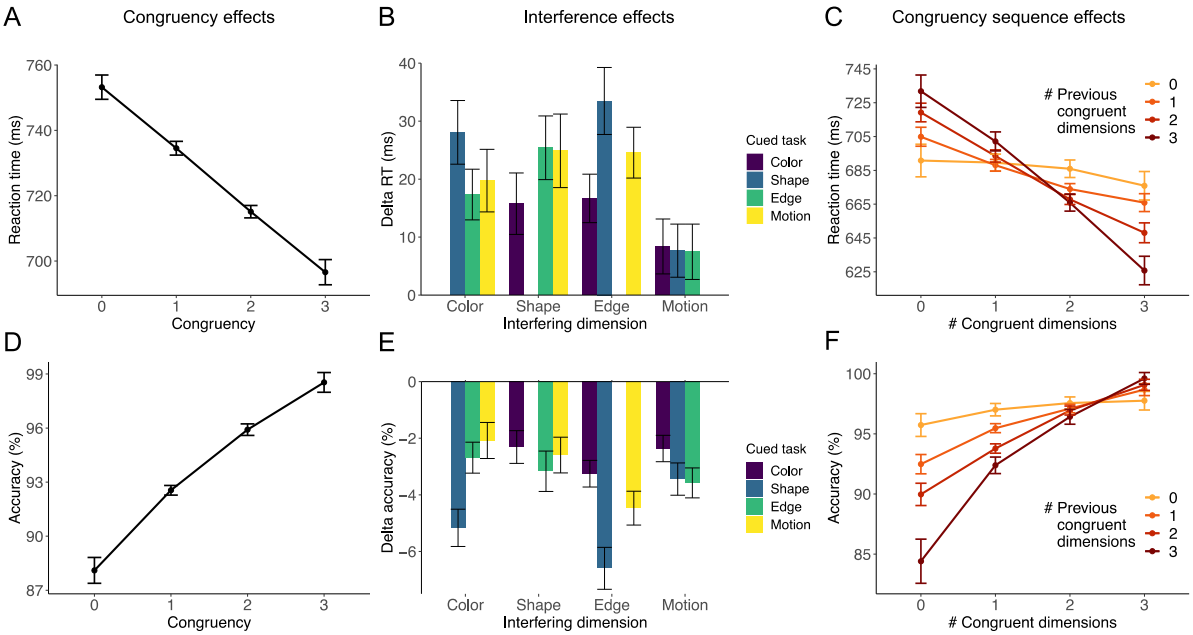

Fig. 2. Experiment 1 (*N* = 104). RT results in top row, accuracy results in bottom row. **a,d,** Parametric congruency effects. Participants responded faster and more accurately when the congruency level increased. **b,e,** Each non-cued dimension (x axis) generated a reliable interference effect (performance difference between trials in which the non-cued dimension is incongruent vs. congruent) on any cued task (color coded). **c,f,** Parametric congruency sequence effects: The effect of congruency on the current trial (x axis) is moderated by the congruency level of the preceding trial (color coded). Participants are less sensitive to the current congruency level after experiencing a low level of congruency in the previous trial, as indicated by flatter lines (lighter colors). Steeper lines (darker colors) reflect both stronger facilitation of the correct response by congruent non-cued dimensions, as well as stronger interference from incongruent non-cued dimensions; both facilitation and inhibition effects are more pronounced following highly congruent trials, reflecting a relaxed attentional control. Error bars represent within-subject 95% confidence intervals.

Each stimulus dimension produces interference

A key benefit of MULTI is that it allows measurement of interference effects (the difference between congruent and incongruent trials) for each non-cued dimension (Figure 2BC). All non-cued dimensions produced interference for both response time (color: BF = 2.7×10^{18}, shape: BF = 6.8×10^{17}, edge: BF = 1.4×10^{22}, motion: BF = 2.0×10^{5}), and accuracy (color: BF = 1.1×10^{22}, shape: BF = 3.0×10^{21}, edge: BF = 4.4×10^{27}, motion: BF = 6.8×10^{20}). In fact, we observed reliable interference from any non-cued dimension on any cued task (Figure 2D) for both RT (min BF = 26 , max BF = 2.8×10^{16}) and accuracy (min BF = 4.6×10^{5}, max BF = 2.7×10^{24}).

These results indicate that the task produces interference from multiple distractor dimensions. However, one alternative explanation is that each participant was susceptible to interference from only one randomly selected dimension, and indifferent to the others. At the group level, this would yield reliable interference for each dimension. Our results are inconsistent with this hypothesis, because on average between 3 and 4 non-cued dimensions induced a positive interference effect (RT: M = 3.38, SD = 0.74; accuracy: M = 3.85, SD = 0.36). If only one dimension provided

interference per participant, we would expect 2.5 positive effects on average (1 from the single interfering dimension, and 1.5 from 3 effects hovering around zero). Indeed, the average number of positive interference effects was higher than this value for both RT (BF = 5.1×10^{39}) and accuracy (BF = 3.9×10^{96}).

Unlike classic tasks, distractors in the MULTI are not more automatically processed than targets. This is intentional, because each dimension is both a target and a distractor. Therefore, it is possible that a dimension only produces interference right after it was cued (Wylie & Allport, 2000). Contrary to this hypothesis, the Supplementary Information reports an analysis that demonstrates that interference effects last for several blocks of trials (see Supplementary Figures 1 and 2) after the last switch (for all Experiments reported here).

Conflict adaptation is modulated by the degree of prior congruency

Having validated our task, we tested how attentional control is regulated during multidimensional interference. Specifically, we measured susceptibility towards distraction as a function of degree of congruence on the previous trial. Classic interference paradigms show a congruency sequence effect, where participants show a reduced interference after an incongruent trial (Kerns et al., 2004). In the MULTI, however, congruence varies parametrically, which allows us to test whether the congruency sequence effect covaries with prior congruency level.

A generalized Bayesian linear regression (see Methods) revealed very strong evidence for an interaction between current and previous congruency on RT (ER = Inf). The marginal means in Figure 2C show that the effect of current congruency was strongest after fully congruent trials, and then parametrically declined with prior congruency. Indeed, we observed only modest evidence for an effect of congruency when the previous trial was fully incongruent (ER = 11.2, M_{00} = 682 ms, M_{03} = 677 ms), but there was very strong evidence for an effect of congruency after fully congruent trials (ER = Inf, mean M_{30} = 729 ms, M_{33} = 624 ms). Compared to trials following fully incongruent trials, trials following fully congruent trials showed strong interference when current congruency was low (leading to longer RTs; $M_{30\text{-}00}$ = 47 ms) and strong facilitation when current congruency was high (leading to shorter RTs; $M_{33\text{-}03}$ = -53 ms).

Analogous congruency sequence effects were observed for accuracy (Figure 2F). There was strong evidence for an interaction between current and previous congruency (ER = Inf), with larger congruency effects as previous congruency increases. Here, there was a strong effect of congruency even when the previous trial was fully incongruent (ER = Inf, M_{00} = 95.9%, M_{03} = 98.3%), but this effect was larger after fully congruent trials (ER = Inf, mean M_{30} = 84.6%, M_{33} = 98.8%). Mirroring RTs, we also found interference and facilitation effects on accuracy ($M_{30\text{-}00}$ = -11.3%, $M_{33\text{-}03}$ = 0.5%).

The Supplementary Information describes a series of analyses that demonstrate that these adaptation effects persist over multiple trials, but that they do not survive a task switch (e.g., switching from the color to the shape task). That is, after a task switch, the attentional state ‘resets’, with prior congruency not affecting behavior.

Conflict adaptation is selectively driven within and not across dimensions

Even though these results establish parametric conflict adaptation, its mechanism remain unclear. One possibility is that conflict is measured across dimensions, and then applied to all distractors or to the target dimension. Alternatively, conflict may be estimated separately for each non-cued dimension, and then resolved selectively within them.

To address this question, we fully leveraged the MULTI, estimating trial-to-trial conflict adaptation within and across non-cued dimensions. For each participant, we computed average adaptation effects for each of the sixteen possible combinations of non-cued dimensions. Some of these combinations assess within-dimension adaptation (e.g., the effect of previous color congruency on current color sensitivity), while the other assess across-dimension adaptation (e.g., the effect of previous shape congruency on current edge sensitivity). For each combination of "current" and "previous" dimension, we computed the adaptation effect as the difference between the interference effects after congruent and incongruent values on the previous dimension. Larger adaptation effects indicate an increase in sensitivity to a non-cued dimension after a trial with congruence on (the same or a different) non-cued dimension.

Our results showed a highly specific pattern of congruency sequence effects (Figure 3), with strong adaptation effects for within-dimension sequences (on the diagonal of the heatmap on panels A and C) and no adaptation effects for across-dimension sequences (off-diagonal). This pattern was observed for RTs (Figure 3AB) and accuracy (Figure 3CD). To test this, we compared each within-dimension adaptation effect to an average across-dimension adaptation effect. For reaction times, we observed strong evidence for larger within-dimension compared to across-dimension adaptation effects ($BF_{color} = 3\times10^{30}$, $BF_{shape} = 6\times10^{14}$, $BF_{edge} = 5\times10^{9}$, $BF_{motion} = 217$). When considering across-dimension effects by themselves, we found some evidence favoring the null hypothesis, suggesting absence of any adaptation effect ($BF0_{color} = 8.8$, $BF0_{shape} = 1.8$, $BF0_{edge} = 4.9$, $BF0_{motion} = 9.0$). Accuracy rates also showed larger within- compared to across-dimension adaptation effects ($BF_{color} = 3.9\times10^{18}$, $BF_{shape} = 2.9\times10^{7}$, $BF_{edge} = 1.1\times10^{10}$, $BF_{motion} = 669$). Again, when considering the across-dimension effects by themselves, we found some evidence favoring the null hypothesis, suggesting absence of adaptation effect for color and motion ($BF0_{color} = 9.20$, $BF0_{shape} = 0.34$, $BF01_{edge} = 0.06$, $BF0_{motion} = 5.95$).

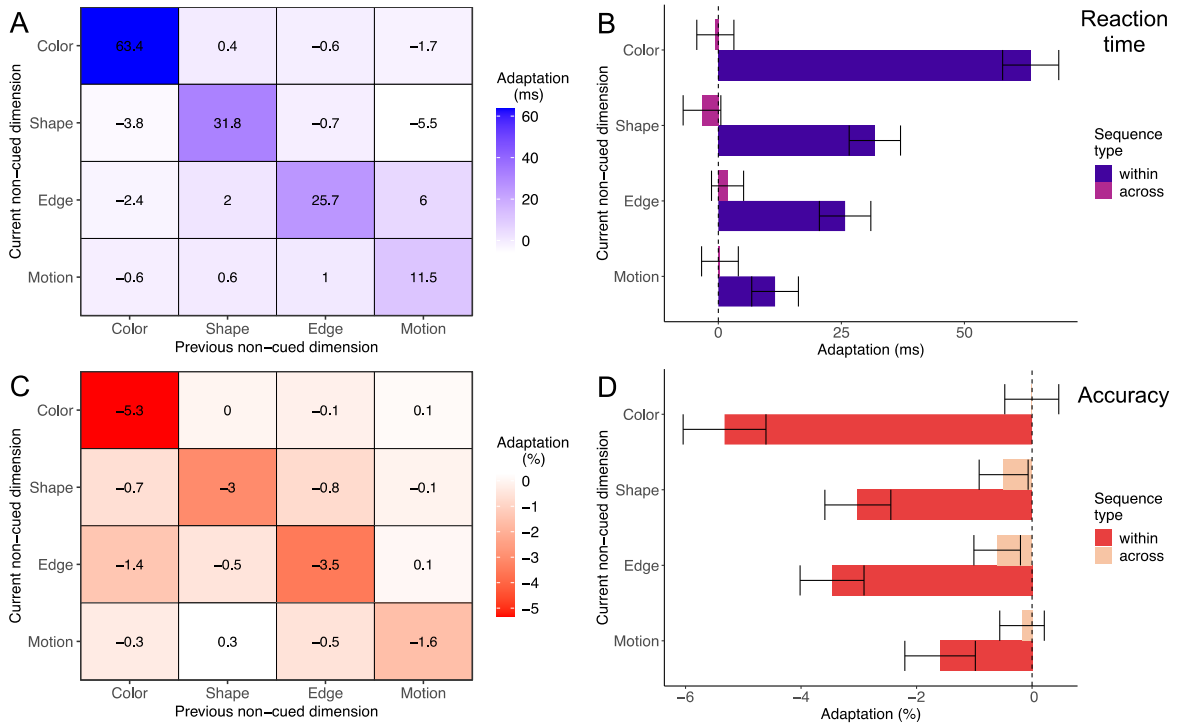

Fig. 3. Distractor-specific conflict adaptation in Experiment 1. RT results in top row, accuracy results in bottom row. **a,c,** The heatmaps represent the trial-to-trial conflict adaptation effects measured in reaction time and accuracy, for each of the 16 possible sequences of previous and current non-cued dimensions. Diagonal cells correspond to within-dimension effects (e.g., the effect of previous color congruency on current color sensitivity), while off-diagonal cells correspond to across-dimension effects (e.g., the effect of previous color congruency on current shape sensitivity). The adaptation effects are defined as the difference between the current dimension's interference effect after a congruent value on the previous dimension compared to after an incongruent value on the previous dimension. Adaptation effects were observed selectively within dimensions (diagonal), reflecting high specificity in control adaptation: Sensitivity to a current non-cued dimension only depends on its own congruency history. **b,d,** Average adaptation affects for sequences across dimensions (lighter colors, off-diagonal in heatmaps) compared to adaptation effects within dimensions (darker colors, on-diagonal cells in heatmaps). We found strong evidence for larger within-dimensions compared to across-dimensions effects. When considering across-dimensions adaptation effects by themselves, we found moderate evidence for absence of effects for color, edge, and motion dimensions. Error bars represent within-subject 95% confidence intervals.

This trial-average approach isolates adaptation effects for each combination of non-cued dimension, but it cannot account for cumulative effects of multiple non-cued dimensions on any given trial, as suggested by the parametric CSE results. In the Supplementary Materials, we report a series of Bayesian multilevel generalized linear models to explain within- and across-dimension adaptation effects at the single trial level. Their results were analogous to what is reported here.

Note that it is possible that each participant only applied within-dimension adaptation for one dimension, which then revealed itself as reliable within-dimension adaptation for all of them at the group level. To rule this out, we computed participant's average number of positive within-dimension adaptation effects for both RT (M = 3.38, SD = 0.68) and accuracy (M = 3.37, SD = 0.71), and found strong evidence that these numbers were larger than the 2.5 positive adaptation effects one would expect under the alternative explanation (RT: BF = 5.5×10^{45}, Accuracy: BF = 3.6×10^{42}).

Ruling out the role of eye movements

Because the MULTI presents two stimuli side-by-side, an alternative explanation is that our results are best explained by eye movements towards the stimulus with most target features, and not by attentional control. To rule out this explanation, participants in Experiment 2 (*N* = 104) completed a version of the MULTI in which each trial only presented a single stimulus, eliminating eye movement between objects (Figure 4A). The single central stimulus was constructed by randomly sampling one feature for each of the four dimensions, with the task cue presented in its middle. Participants reported whether the target feature of the cued dimension was present. We report the most important results here, but the Supplementary Information contains a more comprehensive description.

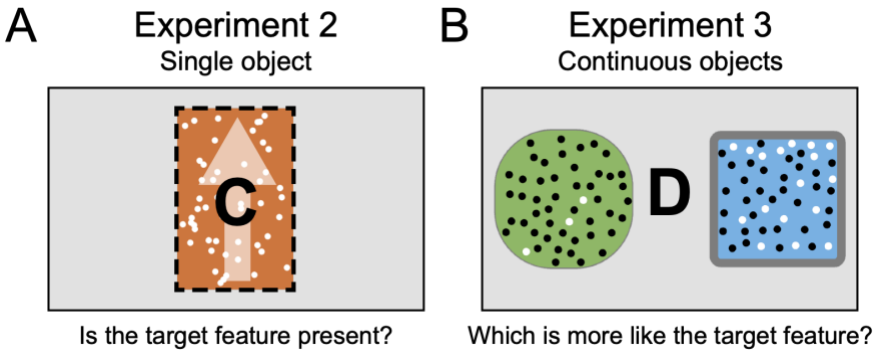

Fig. 4. Design of Experiments 2 and 3. a, In Experiment 2 a single central stimulus was presented at the center of the screen, consisting of a random sample of each of the four dimensions used in Experiment 1. The task cue was presented in the middle of the stimulus. At the beginning of the experiment, participants were given a target feature for each of the four dimensions. On each trial, they report whether the target feature of the cued dimension was present or absent. **b,** In Experiment 3, stimuli varied continuously on each dimension. Two objects are presented side-by-side, flanking a cue letter. Across trials, the two objects continuously differed in color (blue to green), shape (circle to square), edge (thick to thin), and the ratio of black to white dots (dots were static unlike Experiment 1 and 2). At the beginning of the experiment, participants were instructed to find a prototypical target feature, for each dimension (e.g., the more circular shaped). On each trial, participants reported which object was closer to the instructed feature.

All results replicated in this single-stimulus version of the MULTI (Figure 5). The number of congruent dimensions parametrically modulated task performance in terms of RT and accuracy. Non-cued dimensions produced interference both on the aggregate and when considered separately for each cued dimension. We again observed sequential conflict adaptation effects (Figure 5CF), with susceptibility to current congruency increasing when following more congruent trials (ER_{RT} = Inf, $ER_{accuracy}$ = Inf). Most importantly, conflict was again resolved in a dimension-specific fashion, driven by conflict sequences within and not across dimensions (Figure 6A-D). Bayesian hierarchical regressions provided consistent results at the single trial level (see Supplementary Information).

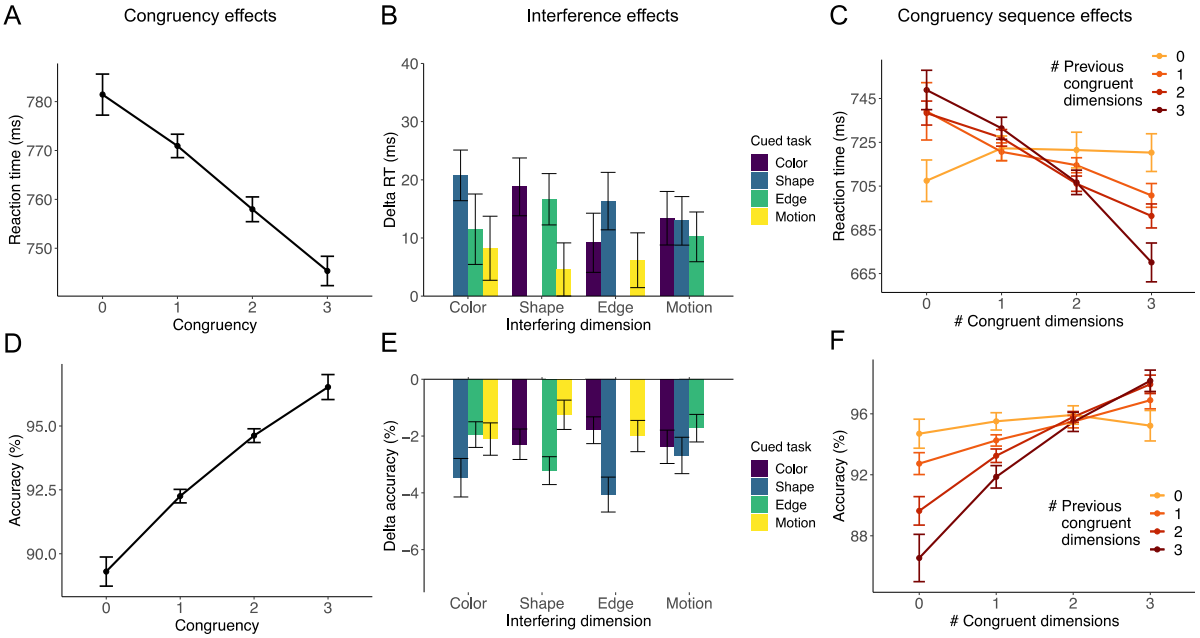

Fig. 5. Experiment 2. (*N* = 104). RT results in top row, accuracy results in bottom row. Compare to Figure 2. **a,d,** Participants showed a parametric congruency effect with faster and more accurate responses when the congruency level was higher. **b,e,** Each non-cued dimension (x axis) generated a reliable interference effect on any cued task (color coded), except for a weak effect of the interfering dimension shape on the motion task (BF = 2.0). **c**,**f**, Parametric congruency sequence effects: Participants are less sensitive to the current congruency (x-axis) after experiencing a low level of congruency in the previous trial, as indicated by flatter lines (lighter colors). When previous congruency increases (darker colors), participants become sensitive to current congruency, with both facilitation (better performance for relatively congruent trials) and interference (worse performance for relatively incongruent trials), reflecting a relaxed attentional control. Error bars represent within-subject 95% confidence intervals.

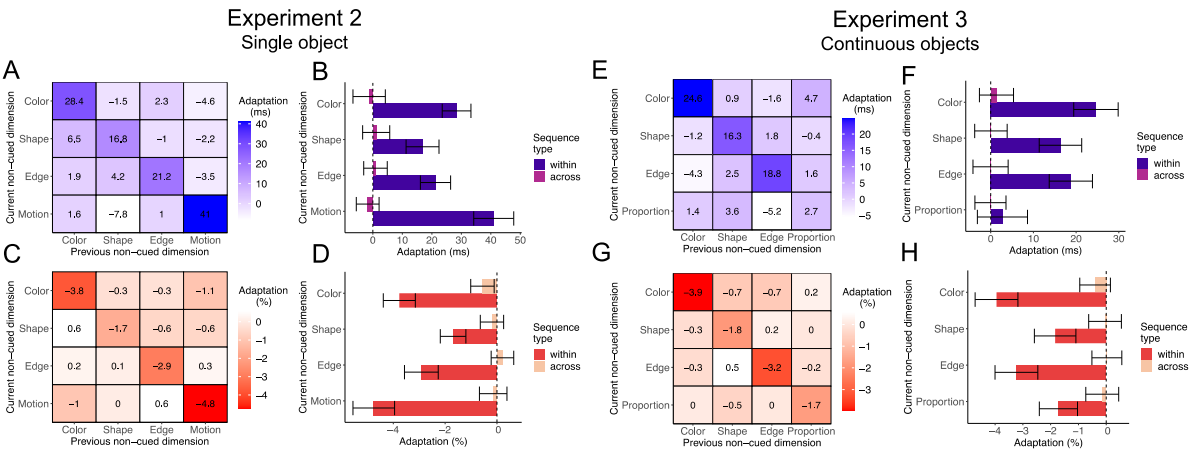

Fig. 6. Distractor-specific conflict adaptation in Experiments 2 and 3. Both experiments replicated the pattern of highly specific congruency sequence effects found in experiment one, with strong within-dimension effects and absence of across-dimensions effects. Results of Experiment 2 (*N* = 104) on the left side: **a,c,** Heatmaps represent the trial-to-trial conflict adaptation effects on reaction times (**a**) and accuracy (**c**), for each of the 16 possible sequences of previous and current non-cued dimensions. Diagonal cells corresponds to within-dimension effects, while off-diagonal cells corresponds to across-dimension effects. Adaptation effects were observed selectively on diagonal cells, reflecting high specificity in control adaptation: Sensitivity to a current non-cued dimension only depends on its own congruency history. Notably, in the task with a single stimulus per trial (Experiment 2), the dimension motion induced the most robust interference effect (Figure 5BE), mirrored by the strongest within-dimension congruency sequence effect. **b, d,** Average across-dimension sequence effects (lighter colors) compared to within-dimension adaptation effects (darker colors). For each dimension, we observed larger evidence for within compared to across dimension adaptation for both RT ($BF_{color} = 3.5\times10^{11}$, $BF_{shape} = 2.3*10\times^{2}$, $BF_{edge} = 5.9\times10^{6}$, $BF_{motion} = 3.8\times10^{15}$) and accuracy ($BF_{color} = 2.3\times10^{9}$, $BF_{shape} = 4771$, $BF_{edge} = 1.2\times10^{9}$, $BF_{motion} = 3.8\times10^{12}$). When we only considered across dimension adaptation effects, we observed some evidence favoring the null hypothesis, suggesting absence of adaptation for RT ($BF0_{color} = 8.31$, $BF0_{shape(off)} = 7.83$, $BF0_{edge} = 8.24$, $BF0_{motion} = 6.21$) and accuracy ($BF0_{color} = 0.25$, $BF0_{shape} = 5.64$, $BF0_{edge} = 5.30$, $BF0_{motion} = 7.60$). Results of Experiment 3 (*N* = 103) on the right side: **e,g,** All effects replicated except for a weak within-dimension congruency sequence effect for the new dimension dot-proportion, on RT selectively. **f,h,** For RT, we observed larger evidence for within compared to across dimension adaptation ($BF_{color} = 1.7\times10^{8}$, $BF_{shape} = 4.2*10\times^{4}$, $BF_{edge} = 5.6\times10^{4}$), except for the new dot proportion dimension ($BF_{proportion} = 0.23$). For accuracy, there was strong evidence for larger within than across adaptation for all dimensions ($BF_{color} = 3.7\times10^{8}$, $BF_{shape} = 150$, $BF_{edge} = 1.4\times10^{7}$, $BF_{proportion} = 45$). Note that this asymmetry in adaptation for dot proportion between measures is consistent with the weak interference effect on RT but not accuracy for this dimension (cf. Figure 7B). We found moderate evidence for a lack of across-dimension adaptation for all dimensions for RT ($BF0_{color} = 7.18$, $BF0_{shape} = 9.16$, $BF0_{edge} = 9.16$, $BF0_{proportion} = 9.16$) and accuracy ($BF0_{color} = 3.27$, $BF0_{shape} = 9.03$, $BF0_{edge} = 9.13$, $BF0_{proportion} = 7.94$). Error bars represent within-subject 95% confidence intervals.

Ruling out associative priming

It is often suggested that congruency sequence effects do not measure dynamics in attentional control, but instead reflect priming of associations between stimulus features or between stimulus and response (Hommel et al., 2004; Mayr et al., 2003; Schmidt & De Houwer, 2011). The MULTI mitigates some of these concerns because responses are not tied to a given feature: each target feature can appear left or right on any trial. Moreover, stimuli are unlikely to repeat from trial-to-trial, because there are sixteen possible combinations of features. Contingency learning is a more serious confound affecting both experiments: among the sixteen potential stimuli, fully congruent and fully incongruent ones occur less frequently. Therefore, stimuli with more extreme congruency levels could be more strongly linked to certain responses.

To rule out any stimulus identity confounds, we ran a third experiment (*n* = 103) in which all features on each dimension were sampled from continuous ranges (Figure 4B). Objects continuously differed in color (blue to green), shape (circle to square), edge (thick to thin), and the ratio of black to white dots. Participants were asked to report which object was closer to the instructed feature (e.g., which is object is more circular?). Because each trial included two unique new features for each dimension, there was no possibility for associative priming and contingency learning to affect behavior. The Supplementary Information report the results in full detail.

Experiment 3 replicated the previous two experiments. On each trial, the level of congruency affected performance for both RT and accuracy (Figure 7). Individual dimensions produced reliable interference effects for both performance measures, although the effect of the new dot proportion dimension was markedly reduced for RT. We once again found strong evidence for conflict adaptation. On the aggregate, susceptibility to current congruency was increased following more congruent trials (ER_{RT} = Inf, $ER_{accuracy}$ = Inf).

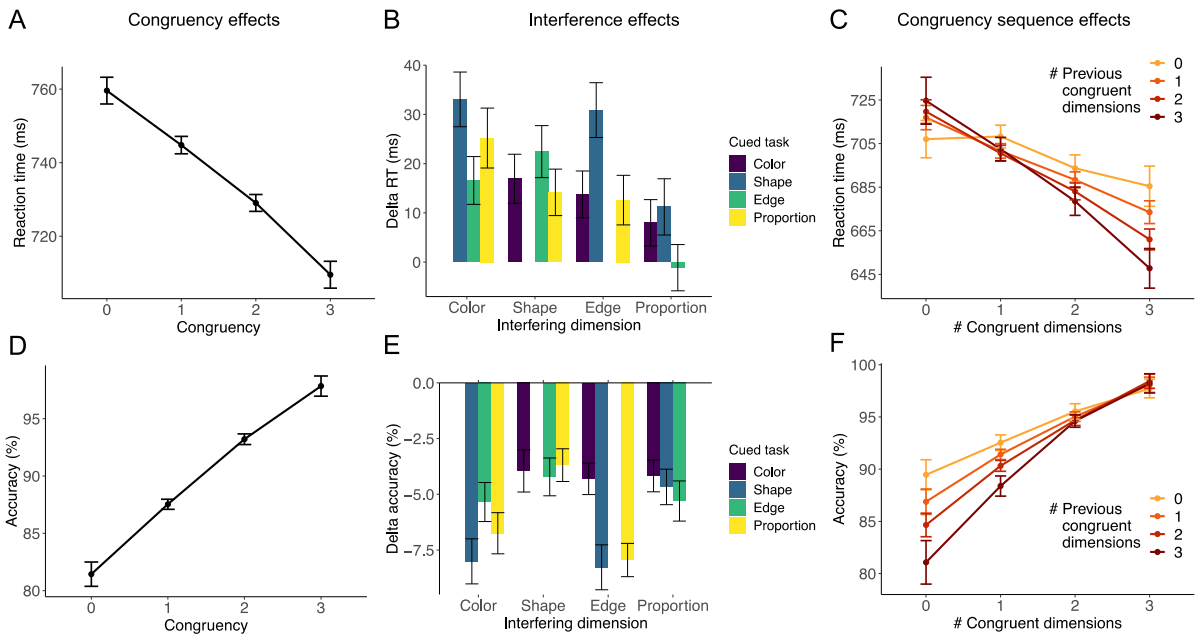

Fig. 7. Experiment 3 (*N* = 103). RT results in top row, accuracy results in bottom row. Compare to Figures 2 and 5. **a,d,** Participants showed a parametric congruency effect with faster and more accurate responses when the congruency level was higher. **b,e,** Each non-cued dimension (x axis) generated a reliable interference effect on any cued task (color coded), except for the interfering dimension dot proportion on the edge task (BF = 0.08). **c**,**f**, Parametric congruency sequence effects: Participants are less sensitive to the current congruency (x-axis) after experiencing a low level of congruency in the previous trial, as indicated by flatter lines (lighter colors). Error bars represent within-subject 95% confidence intervals.

As before, adaptation effects were resolved selectively within each dimension (Figure 6, panels E-H) for accuracy. For RT, we observed larger evidence for within compared to across dimension adaptation except for the new dot proportion dimension. Note that this asymmetry in adaptation for dot proportion between measures is consistent with the weak interference effect on RT but not accuracy for this dimension. Moreover, we found moderate evidence for a lack of across-dimension adaptation for all dimensions for both RT and accuracy.

These results indicate that adaptation effects in the MULTI are not driven by associative priming or other stimulus identity confounds. They instead suggest that the task induces multi-dimensional attentional control demands, and that these are resolved in a dimension-specific way.

A neural network model of dimension-specific adaptation

Classic computational models of cognitive control suggest that conflict is monitored after evidence from all dimensions have been merged into stimulus or response representations. However, such a mechanism is incompatible with our results, because conflict adaptation is sensitive to the source of conflict. This suggests that conflict monitoring occurs at a lower level in the information processing stream.

To demonstrate this, we developed a novel neural network model of our task based on prior models of classic interference paradigms (Cohen et al., 1990; Cohen & Huston, 1994; Ritz & Shenhav, 2022). An outline of the model is shown in Figure 8. The model processes stimulus information through four feed-forward pathways, corresponding to the four dimensions of the MULTI paradigm. Each pathway starts with two units that represent the location of the target feature of the corresponding dimension on the current trial. Their activation propagates to intermediate units, which also receive activation from four 'task' units that represent attention allocated to each dimension. Each task unit is connected to the two intermediate units in the pathway of the corresponding dimension and is activated when the corresponding dimension is cued. Their activation sensitizes processing in the intermediate units by offsetting a constant negative bias (Cohen et al., 1990). Each intermediate unit then propagates its activity to two response units. In order to simulate RTs and accuracy, each trial's evidence (the difference in activation between the two response units) is sampled through a diffusion process (Ratcliff et al., 2016). We hand tuned the parameters of the model, aiming to reduce complexity by equating them where possible.

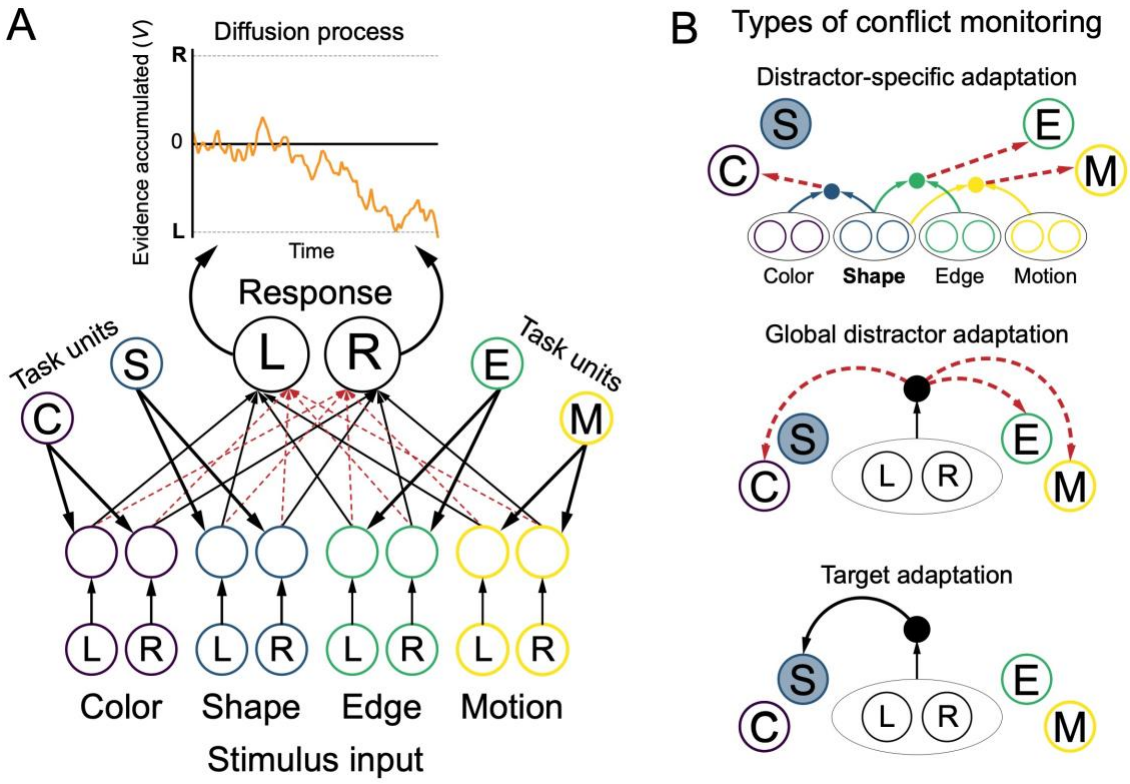

Fig. 8. Computational model. a, Schematic description of the neural network model developed for the MULTI paradigm. (Bottom) The model processes stimulus information through four feed-forward pathways (color coded), corresponding to the four dimensions of the MULTI paradigm. Each pathway projects to a set of intermediate neurons, where stimulus input activity converges with activity from task units and a constant negative bias. The intermediate layer projects to two response units using balanced excitatory (black lines) and inhibitory connections (red dashed lines). (Top) In order to simulate RTs and accuracy, each trial's evidence (the difference in activation between the two response units) is sampled through a diffusion process. **b**, Three variants of conflict-control loops. In all, conflict (solid circles) is transformed into a measure of control, which is used to update the activity of the task-demand units on the next trial (red dashed arrows: subtraction, black solid arrows: addition). (Top) Dimension-specific adaptation assumes that multiple conflict measures are computed within each processing pathway. Three conflict-monitoring units compute conflict by comparing 'representational evidence' for each non-cued dimension against evidence in the cued dimension. Each dimension's conflict is then transformed in a control value which is used to change activity in the corresponding task unit on the next trial. Because each dimension has a separate conflict-control loop, this mechanism allows for distractor-specific adaptation. (Middle) To implement global distractor adaptation, conflict is detected at the response level, where the processing pathways have fully converged. Specifically, the single conflict value is transformed into a control value, and then subtracted from the input to all non-cued dimensions' task demand units. (Bottom). For target adaptation, conflict is detected at the response level. This single conflict value is transformed into a control value, which is then added to the input to the cued dimension's task demand unit.

The model captures single-trial level performance well. Average RT decreased and accuracy increased when the number of congruent non-cued dimensions increased (Figure 9AB). This happens because on trials with more congruent dimensions, evidence in the response layer is more favorable towards the correct response, leading to a faster and more accurate resolution of the diffusion process.

Next, we aimed to simulate the distractor-specific conflict adaptation effects from Experiments 1-3. To do this, we considered three variants of models that implement some form of conflict monitoring. In all, "monitoring" units measure the degree of conflict, or simultaneous activation of

incompatible pieces of information (e.g., responses). This measure of conflict is then transformed to a measure of *control*, which is used to configure the activity of the task units on the next trial. Attentional control (target enhancement or distractor suppression) is increased after trials with a high degree of conflict (Egner, 2017; Gratton et al., 1992).

We simulated three models to demonstrate that distractor-specific congruency sequence effects require individual conflict monitors for each distractor. The key difference between these models is how they monitor conflict. The first detects conflict separately for each input dimension, whereas the last two models detect conflict at the response level.

Our first model assumed multiple conflict-monitoring units at the intermediate layer of the network (Figure 8B top). Here, we define conflict as the simultaneous activation of incompatible representations within each processing pathway (i.e., at the stimulus representation level). To quantify dimension-specific conflict, the model first computes 'representational evidence' within each non-cued pathway, indicating whether the dimension provides evidence for a left or right response. Second, this evidence gets translated into conflict by comparing it to the direction of the cued dimension's evidence: when evidence in a non-cued dimension is consistent with the cued dimension's evidence, conflict is low, and when they are inconsistent, conflict is high. The non-cued dimension's evidence determines the intensity of conflict. The resulting control value is then applied to the corresponding task unit on the next trial.

This model reproduced the parametric congruency sequence effects observed across all three experiments (see Figure 9, panels A & B). More importantly, this model successfully captured dimension-specific congruency sequence effects. As can be seen in Figure 9 panel C, adaptation effects were determined exclusively by congruency sequences within the same dimension, and not across dimensions – in full agreement with the empirical evidence.

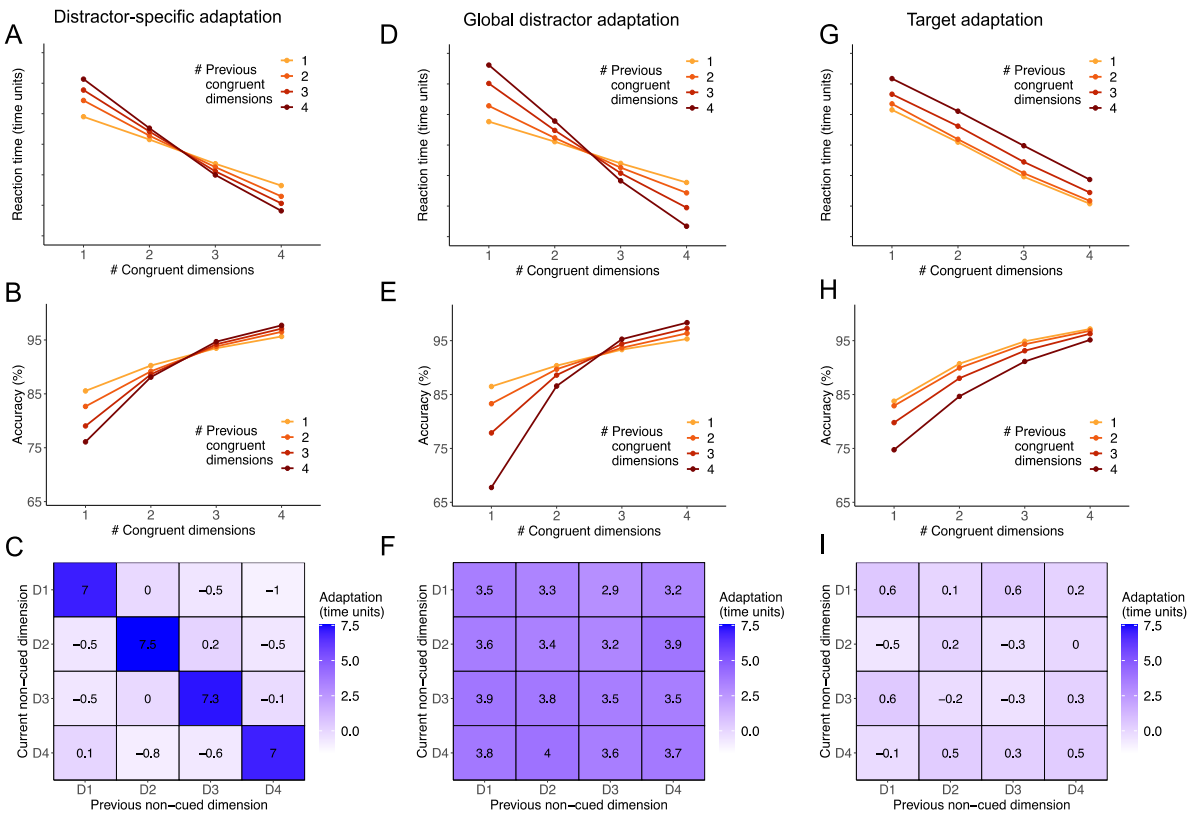

Fig. 9. Simulation results of the three conflict adaptation models. a,b,c Distractor-specific adaptation reproduces the parametric congruency sequence effects observed across all three experiments, and yields dimension-specific congruency sequence effects: Adaptation effects are determined exclusively by congruency sequences within the same dimension (on diagonal cells). **d,e,f,** Global distractor adaptation reproduces the parametric congruency sequence effects, but fails to capture dimension-specific congruency sequence effects: Adaptation effects are determined both within- and across-dimensions. **g,h,I**, Target adaptation fails at capturing the parametric congruency sequence effects observed across the three experiments, and also the within-dimension specificity of this adaptation. However, it does yield generally increased performance after high-conflict trials (lower RT and higher accuracy), but these effects are orthogonal to the current congruency level. For accuracy, it may appear we observe a parametric sequence effect, since the slope of the curves change with prior congruency, but this reflects a linear shift in accuracy combined with a ceiling effect at the higher congruency levels.

To demonstrate that the existence of multiple conflict monitors is crucial for capturing our results, the last two models implement a single conflict detector at the level of the response nodes (Botvinick et al., 2001; Verguts & Notebaert, 2008), Specifically, we quantify response conflict as the absolute difference in activation between the two response nodes. When the activation of one response unit is larger than the other, evidence is large and conflict is low. When the activity of the two units is similar, evidence is low and conflict is high[1]. This single conflict signal is then applied to the task demand units.

[1] Botvinick and colleagues (2001) quantified conflict as Hopfield energy, which corresponds to the multiplication of the response units' activations. This fit their purpose, because response selection is modeled through mutually inhibitory connections, that continue to multiply their activation until one response node hits a threshold. Because we resolve decisions using a diffusion process on the signed difference between response nodes, however, the absolute difference between response nodes provides a more natural measure of conflict.

For the second model, we moved conflict detection to the response level (Botvinick et al., 2001), but still used this to adapt distractor representations (Figure 8B middle). After conflict detection, the control value was subtracted from the input to all non-cued task units. This way, attention to all distractors reduces (increases) after particularly incongruent (congruent) trials, yielding global distractor adaptation. This model produced parametric congruency effects. For both accuracy and reaction time, the effect of congruency on a given trial was dependent on the congruency of the previous trial (Figure 9, panels D & E). However, the model did not produce dimension-specific adaptation. As shown in Figure 9F, adaptation effects were observed both for sequences within and across dimensions. This pattern arises because the system is agnostic to the source of the conflict. At the response layer, conflict reflects the cumulative contributions of all non-cued dimensions, making dimension-specific contributions unidentifiable.

The third model implemented conflict adaptation by enhancing the target representation (Figure 8 bottom). Specifically, after the conflict value is transformed into a control value, it is added to the input to the cued dimension's task unit. This model failed to capture parametric congruency sequence effects (Figure 9GH). High-conflict trials induced increases in performance (lower RT and higher accuracy) on the next trial, but they did not change susceptibility to distractor information. Naturally, this model was also unable to capture the distractor-specific adaptation we observed in our experiments (Figure 9I).

This last set of simulations demonstrate that, at least in the gold-standard neural network model of attentional control (Cohen et al., 1990), congruency sequence effects cannot be achieved by modulating attention to the cued task. This is because target adaptation primarily changes the availability of evidence related to the cued dimension, leading to similar changes in accuracy and RT (i.e., in the signal-to-noise ratio of the target feature) independently of the current level of congruency. Importantly, target adaptation leaves unaffected the processing of the distractors[2]. Therefore, no changes in facilitation nor interference from non-cued dimensions are to be expected.

In short, only our novel network was able to capture our behavioral results, due to a control mechanism that independently monitors conflict for each non-cued dimension.

Discussion

In everyday life, we constantly need to switch attention between multiple competing sources of information. However, attentional control is typically investigated assuming a clear dichotomy between one task-relevant and one task-irrelevant dimension. Here, we designed a multidimensional task-set interference paradigm that allowed us to observe attentional control dynamics in multidimensional environments. We show that attentional control selectively adjusts the processing of each task-irrelevant dimension based on its own congruency history. Computational modeling demonstrated that, in order to capture this pattern of behavior, conflict must be monitored for each distractor separately, and detected at the level of stimulus representation.

---

[2] Note that in the original model by Botvinick and colleagues (2001), attention deployed to targets and distractors was bound together. Their model implemented control adaptation by first allocating attention to the target, and then setting complementary values to distractors. Here we decoupled these two mechanisms by allowing independent contributions to the response units from cued and non-cued dimensions' pathways. Therefore, here a reduced processing of cued dimensions does not necessarily entail an increased processing on non-cued ones, and vice versa.

This finding invites important revisions to established models of conflict adaptation. As mentioned before, classic models assume a single conflict detector at the response level, where information about individual stimulus features is no longer accessible. Our simulations show that this produces adaptation effects that transfer between distractor dimensions, unlike our findings. Interestingly, the domain-specific model proposed by Egner (2008) assumes separate monitors for different conflict types along the information processing stream (e.g., semantic vs. response conflict), but at levels where individual dimension information is unavailable. This model explains conflict adaptation in tasks where distractors elicit different conflict types (e.g., a combined Stroop-Simon task; Braem et al., 2014). In our task, however, all conflict happens exclusively at the stimulus level (Kornblum et al., 1990), yet we see no transfer of adaptation between dimensions. This flattens any proposed hierarchy of conflict monitors, allowing them to operate independently at the level of feature-based attention (Algom & Chajut, 2019). Finally, another class of models applies control to the connections between task representations and stimulus features (Blais et al., 2007; Verguts & Notebaert, 2008). For example, Verguts and Notebaert (2008) proposed the 'adaptation by binding' by which conflict changes the connections between all task representations (cued and non-cued) and stimulus features, depending on the activity level of the connected units. Therefore, these models implements both target and distractor adaption. Because they measure conflict at the response level, they are agnostic to the original decomposition of the conflict. This means that they are unable to adapt individual stimulus dimensions based on dimension-specific information. Neither of these assumptions are met in our data. However, these models can simulate how people learn to attach attentional control settings to specific stimulus features (Bugg & Crump, 2012; Bugg & Hutchison, 2013; Jacoby et al., 2003). Our model does not account for such behavior. Future work may attempt to unify these perspectives.

Our computational model bears resemblance to the recently proposed Learned Attention for Control model (Alexander, 2022). This model learns to deploy attention to dimensions that are useful for responding correctly. Similar to our account, this model assumes independent learning processes for each dimension. However, the model has no access to task cues, since it is completely reliant on trial-and-error learning. Therefore, it would likely struggle to effectively capture behavior on tasks where relevant dimensions alternate often and unpredictably. Yet, this bottom-up solution, relying on simple associative learning, may complement the costly mechanisms of rule-based executive control (Kool et al., 2010).

Interestingly, most models resolve conflict by enhancing target information, in clear disagreement with our data. This leads to a simple question: Why did our participants not engage in target enhancement? One possibility is that target enhancement is not a feasible physiological mechanism. A growing body of work suggests that a relative amplification of task-relevant information relies on inhibition of regions processing distractors (Jensen, 2023; Jensen & Mazaheri, 2010). Alternatively, it may be more effective to suppress distractors in our task, especially considering the high number of them. Moreover, more granular control adaptation leads to more optimal control settings, because each dimension can be individually tuned to its response history. For example, it allows for facilitation of a congruent dimension, even when the majority of dimensions was incongruent. This predicts that target enhancement is more likely to happen when distractors are mostly incongruent (Bugg & Crump, 2012; Jacoby et al., 2003), since their potential utility would be smaller. The MULTI is perfectly suited to test this prediction. However, it should be noted that our simulations were unable to capture congruency sequence effects when implementing target enhancement, underscoring it may not be a viable algorithmic solution. Perhaps other computational frameworks can demonstrate adaptation through target

enhancement. However, our model closely followed the benchmark neural network developed by Cohen and colleagues (1990), suggesting that multiple innovations would be necessary for this to become feasible. At any rate, such adaptation would produce generalized transfer of control, which is inconsistent with what we observe.

Our results seem at odds with a small body of literature that finds general control adaptation in two-dimensional task designs. Transfer between distractor dimensions has been reported between two spatially-related dimensions (Kunde & Wühr, 2006), and cross-task transfer has been reported when both tasks shared identical response requirements (Notebaert & Verguts, 2008). These findings suggest that transfer of control may be possible if dimensions share some representational similarity (Braem et al., 2014). Such similarity could arise from perceptual factors (Kunde & Wühr, 2006), statistical associations between co-occurring features over time (Schapiro et al., 2012), or alignment of behavioral outcomes (Ileri-Tayar et al., 2022). In future studies, we will use the general framework of the MULTI to test this hypothesis.

The distractor-specificity of conflict adaptation has implications for neuroscientific models of cognitive control. These models posit the anterior cingulate cortex (ACC) as the central hub for conflict detection (Botvinick et al., 2001; Holroyd & McClure, 2015), transmitting need-for-control signals to lateral prefrontal regions (Cavanagh & Frank, 2014; Shenhav et al., 2013; Verguts, 2017). Our findings challenge the notion that the ACC listens to compressed estimates of conflict in high-level representations. Instead, they suggest that the brain is equipped to monitor a wide variety of stimulus dimensions, and at least three of them in parallel. This calls for alternative neural models of conflict monitoring. One possibility is that the ACC relies on mixed selectivity in neuronal coding. In lateral prefrontal cortex, such mixed selectivity may enable complex behavior and a high capacity for control by expanding the dimensionality of neural representations (Fusi et al., 2016; Rigotti et al., 2013). Similar coding properties may allow the ACC to simultaneously represent multiple sources of conflict from multiple novel stimulus dimensions. For example, the ACC could combine lower-dimensional subspaces, known as manifolds, to flexibly allow for monitoring of incongruency from multiple sensory regions of the cortex. In such a framework, the role of the ACC, integrating task-relevant information from across the brain, would be consistent with Expected Value of Control Theory (Shenhav et al., 2013). Conflict could then be resolved by suppressing interfering information (Danielmeier et al., 2011), represented within distal neural populations, possibly through oscillatory neural dynamics (Buschman & Miller, 2022; Jensen, 2023; Verguts, 2017). Alternatively, and more speculatively, ACC could use these manifolds to represent stimulus information locally. Conflict could then be measured as the overlap between neural representations of task-relevant aspects merged together (i.e., conjunctive representations; Kikumoto & Mayr, 2020), such as the instructed stimulus features and their position on screen. In this framework, conflict would emerge when such representations in the ACC are partially overlapping, thus preventing a linear readout of a single response. Conflict could then be resolved locally by orthogonalizing the representations of mutually incompatible dimensions, on a trial-by-trial basis. This explanation alleviates the need for a high degree of flexibility in connectivity between ACC and posterior regions, but it comes at the expense of an (inefficient) duplication of information from more sensory areas to medial frontal cortex. Our task and computational model may inspire a new set of experiments to contrast these hypotheses.

Conclusion

By expanding the dimensionality of traditional interference paradigms, we have gained valuable insights for our understanding of cognitive control. Our findings suggest that attentional control

selectively adapts task-irrelevant stimulus dimensions based on their congruency history. Control adaptation occurs at the level of feature-based attention and involves multiple independent conflict-control loops. These results challenge previous models of control adaptation and highlight the need for alternative neurocomputational models of conflict representation. We hope that these findings can serve as inspiration for future investigations, leveraging increased complexity while maintaining experimental rigor.

Method

Participants

For Experiment 1, we recruited 151 participants from the student population of Washington University in St Louis. The experiment was administered online. To take part to the study participants were screened for color blindness, and their participation was compensated with curriculum credits. The experimental procedure was approved by the ethical committee. Digital informed consent was obtained before the experiment, and participants were debriefed after the experiment. Inclusion criteria were set to 80% accuracy on each experimental task. Accordingly, data from 47 participants were excluded from further analysis, leading to a final sample size of N=104.

For experiment 2, we recruited 172 participants from the same student population. The procedure and exclusion criteria were the same as in Experiment 1. Accordingly, data from 68 participants were excluded from further analysis, leading to a final sample size of N=104.

For experiment 3, we recruited 115 participants from the same student population. The procedure was the same as in Experiment 1 and 2. The exclusion criteria was lowered from 80% to 70% accuracy on each experimental task, to account for higher difficulty emerged for this design. Accordingly, data from 12 participants were excluded from further analysis, leading to a final sample size of N=103.

Multidimensional Task-set Interference Paradigm

Participants had to attend to the cued dimension and chose one of the two stimuli. Specifically, for any dimension, participants had to find a target feature, and to select the corresponding side (pressing the ‘f’ or ‘j’ key for the left or right stimulus, respectively). At the beginning of the experiment, for each of the four dimensions, participants were instructed which feature was the target. For instance, a participants could receive instructions to ‘find the blue shape’ when the letter cue between the stimuli was C, ‘find the oval’ if it was S, ‘find the dashed-edged object’ if it was B, or ‘find the upward-moving dots’ if it was M. The target features were randomly assigned to each participant at the start and would remain the same for the entire experiment. Participants were trained on five blocks of practice trials to consolidate the dimension-target feature association, first on each cued dimension separately (until criterion of 5 correct trials in a row), and then on the four dimensions interleaved (10 correct trials in a row).

Participants were instructed to respond as quickly and as accurately as possible, and a response deadline of 1.5 s was imposed. An incorrect or late response was signaled by a red X at the center of the screen lasting 0.75 s. The inter trial interval (ITI) was set to 1 s, during which a white fixation cross was shown at the center of the screen. During the practice phase, the fixation cross would turn green for 0.75 s after a correct response.

Cued dimensions could repeat or switch across trials. Repeat trials were preceded by one to four trials of the same dimension, while switch trials were preceded by trials of a different dimension (Liston et al., 2006). During the experiment, the sequence of cued dimensions was pseudo-randomized, so that the same dimension could be cued a minimum of three and a maximum of

five times in a row. All four dimensions was cued within four switch trials. Therefore, participant constantly cycled through the four tasks, organized in series of three to five repeat trials. Participant completed a total of 1600 trials split in two blocks. The average experiment duration was ~60 min, including one self-paced break in the middle.

From this design it follows that task-set interference is parametrically manipulated on a discrete scale ranging 0-3. On each trial, any of the three non-cued dimensions could be congruent if it primed the response side where the cued dimension's target feature was located (i.e., the target features of the cued and non-cued dimensions were on the same side), or incongruent if it primed the opposite response side. Henceforth, congruency is defined as the number of non-cued dimensions that primed the correct response.

All versions of the task were developed with the "jsPsych" library (de Leeuw, 2015), using a customized version of the random-dot-kinematogram plugin (Rajananda et al., 2018).

Experiment 2 – single stimulus task

In this experiment, participants completed a version of the MULTI used in Experiment 1, but with a few important modifications. First, a single stimulus was presented at the center of the screen. Across trials, the stimuli differed along the same four dimensions (color, shape, edge, and dot motion). The cue still indicated which dimension was relevant but was now presented at the center of the single stimulus object. In other words, only four features (one per dimension) were presented on any given trial. Second, for any cued dimension, participants were now instructed to report whether the target feature was present (by pressing the 'f' key) or absent (by pressing the 'j' key).

Experiment 3 – continuous stimuli task

For this experiment, we presented again two stimuli flanking the task cue. Importantly, on each trial and for each dimension, the two objects varied in their appearance along a continuum. For the color dimension, the objects varied from more blue to more green (drawn from a 130° section of a circle in L*a*b color space, L= 70, center at [a,b] = [0, 0], radius = 55, and from 111° to 241°). For the shape dimension, the objects varied from more square-shaped to more circle-shaped, with the degree of 'circleness' being implemented as the size of the radius of rounded corners (0px to 50% of the square's width). We ensured that both objects were drawn with the same surface size. For the edge dimension, the objects varied in the degree of thickness of their borders (1px to 10px). Finally, to accommodate a more perceptible continuous change on the fourth dimension, we manipulated the relative proportion of static black and white dots (100 dots per object; range 0% to 100% black dots). To ensure that within each dimension, differences between objects were perceptible, we drew random features with a minimum distance of 20% of the entire range.

Data analysis and inferential statistics

Switch trials were excluded from the current analyses unless specified. Additionally, we rejected trials with responses faster than 200ms or slower than 5 standard deviations of individual participants' reaction time (RT) distribution. Incorrect trials were further excluded for RT analyses.

Accuracy and reaction times were modeled with a Bayesian multilevel generalized linear approach implemented in R (R Core Team, 2017) with the “brms” package (Bürkner, 2017), which interface R with the Stan language (Carpenter et al., 2017). The analysis pipeline included defining a probability model, computing the posterior distributions for each model’s parameter, and testing hypothesis based on the posterior probability distribution of the parameters and/or marginal means of interest. All models predicting reaction times assumed a shifted lognormal response distribution, with identity and log link functions for the distributional parameters. All models predicting accuracy assumed a Bernoulli distribution with a logit link function. Median and 95% credible intervals of the posterior parameters’ distributions are reported in the log scale (reaction times) or logit scale (accuracy), while estimated marginal means are reported in the response scales (with units ms and percentage, respectively). All models assumed weakly informative priors centered on zero for the parameters of interest, and on the raw mean for the intercept. For hypothesis testing we report evidence ratios (ER). For directional hypotheses, ER corresponded to the proportion of posterior probability samples supporting a hypothesis divided by the proportion supporting its alternative. For point hypothesis (two-sided), the ER was the Bayes Factor obtained with Savage-Dickey density ratio method, i.e. the posterior density at the point of interest divided by the prior density at that point.

Neural network

We simulated behavior on our task using three different neural networks, modeled after influential neural networks of response interference (Botvinick et al., 2001; Cohen et al., 1990). Full code for these models will be provided on GitHub (https://osf.io/j9fxc/).

All these networks shared the same basic structure.

Basic structure

This structure consists of four independent pathways, one per dimension, each starting with a pair of input units that project to a set of hidden units in a hidden layer. Stimulus input reflected the location of the corresponding dimension’s target feature (left or right). Each input unit only projects to a single corresponding hidden unit with the weight of this connection set to 1. Therefore, representation of each individual dimension is maintained in the hidden layer, with eight units reflecting the location of each dimension’s target feature. Moreover, each pair of hidden units receives input from one corresponding ‘task’ unit (four units in total). Critically, the weight of each task unit’s connection to the corresponding pathway was set to 3, which was meant to partially offset a constant negative bias term of -4 for each hidden unit. In the basic model, task units are activated when the corresponding dimension is cued. Total net input to each hidden unit was transformed to an activity value between 0 and 1 using the logistic activation function.

As described in detail by Cohen and colleagues (1990), activation of a task unit facilitates processing in the hidden units of the corresponding dimension. In short, differences in stimulus input activity will be larger for the cued dimension than for the non-cued dimensions, because the bias reduces sensitivity to input. Only for the cued dimension, the task units provide input that relieve this bias, allowing those units to reflect stimulus evidence with more sensitivity.

The hidden layer units then project to two response units (corresponding to a left and a right response respectively). Each hidden unit is connected to both response units, but with opposite signs. Specifically, hidden units that represent the presence of target features on the left project to the left response with a weight of 1, and to the right response with a weight of -1. In contrast, hidden units that represent the presence of target features on the right project to the right response with a weight of 1, and to the left response with a weight of -1.

As before, net input for each response unit was transformed to an activity value between 0 and 1 using the logistic activation function.

Diffusion process

In order to simulate response times and accuracy, we feed the difference in activity between the left and the right response unit (the 'evidence') to a diffusion process (Ratcliff et al., 2016). The model updates a decision variable V (starting at $V_0 = 0$) by accumulating noisy information about the evidence over time. New evidence is accumulated until it exceeds either a negative threshold $(A = -1)$ representing a left response, or a positive threshold $(A = 1)$ representing a right response. Specifically, the model updates the decision variable as.

$$V_{t+1} = V_t + d \cdot e + \epsilon_t.$$

Here, V_t is the decision variable at time t, d is a drift rate parameter indicating the amount of evidence e that is sampled at each point in time, and ϵ_t the amount of noise added to the decision variable. Critically, noise is distributed normally as $\epsilon_t \sim \mathcal{N}(0, \sigma)$, with σ the degree of noise. In all our simulations, we set $d = 0.1$ and $\sigma = 0.1$.

Conflict monitoring

In order to simulate the parametric congruency sequence effects in our data, we took the core structure of this model, and implemented three different conflict monitoring mechanisms. Each of these mechanisms detects conflict (either between conflict response tendencies, or between feature representations), transforms this conflict into a measure of control, and then uses this to change the input to the task units on the following trial. Because congruency sequence effects were not present on switch trials, we did not implement this mechanism on those trials for any of the models.

Our first distractor-specific conflict monitoring model detects the degree of conflict for each non-cued dimension separately. Specifically, the model first calculates a measure of "representational evidence" for each dimension individually, by subtracting activity of the hidden unit representing a target feature on the right from the activity of the hidden unit representing a target feature on the left. Next, this representational evidence is multiplied with the sign of the cued dimension's representational evidence (reflecting the correct response), and this measure is multiplied by -1 to compute dimension-specific conflict. This way, dimension-specific conflict is positive if a non-cued dimension provides evidence inconsistent with the correct response, and negative if it provides evidence consistent with it.

Then, for each dimension, conflict is transformed into a measure of control using the following formula, adapted from Botvinick and colleagues (2001):

$$control_{d,i+1} = \lambda \cdot control_{d,i} + (1 - \lambda)(\alpha \cdot conflict_{d,i} + \beta)$$

The parameters d and i index the dimension and trial, respectively. The parameter λ allows control to reflect a time-discounted measure of prior conflict. However, for simplicity and parsimony we set $\lambda = 0$ (but qualitatively similar effects were observed for non-zero values of this parameter). The parameters α and β are scaling parameters. Here, we set these parameters to $\alpha = 4$ and $\beta = 0$. Unlike in the simulations described below, where β was used to set the sign of the control update, it was not necessary to assign β a non-zero value here, because the sign of each distractor-specific conflict value perfectly corresponded to whether the corresponding dimension was consistent or inconsistent with the correct response.

At the start of the next trial, each non-cued dimension's control value was then subtracted from the (zero) input of that dimension. As described in the main text, this algorithm predicts in distractor-specific adaptation effects, with the model becoming more susceptible to a non-cued dimension after a congruent trial on that dimension, and less susceptible after an incongruent trial.

The remaining two models measured a single form of conflict at the response level, by simply taking the negative of the evidence. Here, more positive values indicate smaller evidence and therefore increased conflict, and vice versa. For both, this measure of conflict is then transformed into control as

$$control_{i+1} = \lambda \cdot control_i + (1 - \lambda)(\alpha \cdot conflict_i + \beta).$$

Not that this formula is almost identical to the one above, except that we do not index conflict and control by dimension anymore, since this information is unavailable at the response level. For the simulations of the two remaining models, we set $\alpha = 4$ and $\beta = 0.44$. We set β so that it would offset the conflict measure, yielding positive values for trials with more conflict, and negative values with less conflict. To do this, we first estimated the average amount of conflict in the model ($\approx$ 0.11), and we then multiplied this by α.

For the second model, implementing global distractor adaptation, the control measure was applied negatively to the task units of all non-cued dimension on the next trial. As described in the main text, this results in a parametric congruency sequence effect, but this adaptation is not distractor-specific. Instead, conflict arising from any distractor leads to an adaptation of susceptibility to all non-cued dimensions on the subsequent trial.

For the third model, implementing target adaptation, the control measure was applied to the task unit of the cued dimension. As described in the main text, this does not result in congruency sequence effects. Instead, the model becomes faster and more accurate if previous trials were more incongruent, but sensitivity to distractor information remains identical.

Simulations

We simulated 1000 synthetic participants on each model. Each of these ran through the same number of trials as in our experiments (1600 trials). The process generating the sequences of cued dimensions was identical to the experiments as well.

References

Alexander, W. H. (2022). *Cognitive control strategies derive from dimension reliability*. PsyArXiv. https://doi.org/10.31234/osf.io/a9bmh

Algom, D., & Chajut, E. (2019). Reclaiming the Stroop Effect Back From Control to Input-Driven Attention and Perception. *Frontiers in Psychology*, *10*. https://www.frontiersin.org/articles/10.3389/fpsyg.2019.01683

Blais, C., Robidoux, S., Risko, E., & Besner, D. (2007). Item-Specific Adaptation and the Conflict-Monitoring Hypothesis: A Computational Model. *Psychological Review*, *114*, 1076–1086. https://doi.org/10.1037/0033-295X.114.4.1076

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624–652. https://doi.org/10.1037/0033-295X.108.3.624

Braem, S., Abrahamse, E. L., Duthoo, W., & Notebaert, W. (2014). What determines the specificity of conflict adaptation? A review, critical analysis, and proposed synthesis. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.01134

Bugg, J. M., & Crump, M. (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*. https://www.frontiersin.org/articles/10.3389/fpsyg.2012.00367

Bugg, J. M., & Hutchison, K. A. (2013). Converging evidence for control of color–word Stroop interference at the item level. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(2), 433–449. https://doi.org/10.1037/a0029145

Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, *80*(1), 1–28. https://doi.org/10.18637/jss.v080.i01

Buschman, T. J., & Miller, E. K. (2022). Working Memory Is Complex and Dynamic, Like Your Thoughts. *Journal of Cognitive Neuroscience*, *35*(1), 17–23. https://doi.org/10.1162/jocn_a_01940

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1). https://doi.org/10.18637/jss.v076.i01

Cavanagh, J. F., & Frank, M. J. (2014). Frontal Theta as a Mechanism for cognitive control. *Psychophysiology*, *51*(8), S8–S8. https://doi.org/10.1016/j.tics.2014.04.012.Frontal

Cohen, J. D., & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In *Attention and performance 15: Conscious and nonconscious information processing* (pp. 453–476). The MIT Press.

Cohen, J. D., McClelland, J. L., & Dunbar, K. (1990). On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect. *Psychological Review*, *97*((3)), 332–361.

Danielmeier, C., Eichele, T., Forstmann, B. U., Tittgemeyer, M., & Ullsperger, M. (2011). Posterior Medial Frontal Cortex Activity Predicts Post-Error Adaptations in Task-Related Visual and Motor Areas. *The Journal of Neuroscience*, *31*(5), 1780–1789. https://doi.org/10.1523/JNEUROSCI.4299-10.2011

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, *47*(1), 1–12. https://doi.org/10.3758/s13428-014-0458-y

Egner, T. (2008). Multiple conflict-driven control mechanisms in the human brain. *Trends in Cognitive Sciences*, *12*(10), 374–380. https://doi.org/10.1016/j.tics.2008.07.001

Egner, T. (Ed.). (2017). *The Wiley handbook of cognitive control*. Wiley Blackwell.

Egner, T., Delano, M., & Hirsch, J. (2007). Separate conflict-specific cognitive control mechanisms in the human brain. *NeuroImage*, *35*(2), 940–948. https://doi.org/10.1016/j.neuroimage.2006.11.061

Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience*, *8*(12), Article 12. https://doi.org/10.1038/nn1594

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143–149. https://doi.org/10.3758/BF03203267

Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, *37*, 66–74. https://doi.org/10.1016/j.conb.2016.01.010

Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*, 480–506. https://doi.org/10.1037/0096-3445.121.4.480

Hazeltine, E., Lightman, E., Schwarb, H., & Schumacher, E. H. (2011). The boundaries of sequential modulations: Evidence for set-level control. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(6), 1898–1914. https://doi.org/10.1037/a0024662

Holroyd, C. B., & McClure, S. M. (2015). Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. *Psychological Review*, *122*(1), 54–83. https://doi.org/10.1037/a0038339

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

Ileri-Tayar, M., Moss, C., & Bugg, J. M. (2022). Transfer of learned cognitive control settings within and between tasks. *Neurobiology of Learning and Memory*, *196*, 107689. https://doi.org/10.1016/j.nlm.2022.107689

Jacoby, L. L., Lindsay, D. S., & Hessels, S. (2003). Item-specific control of automatic processes: Stroop process dissociations. *Psychonomic Bulletin & Review*, *10*(3), 638–644. https://doi.org/10.3758/BF03196526

Jensen, O. (2023). *Gating by alpha band inhibition revised: A case for a secondary control mechanism*. PsyArXiv. https://doi.org/10.31234/osf.io/7bk32

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*(November), 186. https://doi.org/10.3389/fnhum.2010.00186

Kerns, J. G., Cohen, J. D., MacDonald, A. W., Cho, R. Y., Stenger, V. A., & Carter, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science (New York, N.Y.)*, *303*(5660), 1023–1026. https://doi.org/10.1126/science.1089910

Kikumoto, A., & Mayr, U. (2020). Conjunctive representations that integrate stimuli, responses, and rules are critical for action selection. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(19), 10603–10608. https://doi.org/10.1073/pnas.1922166117

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision Making and the Avoidance of Cogntive Demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198.Decision

Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis for stimulus-response compatibility--A model and taxonomy. *Psychological Review*, *97*(2), 253–270. https://doi.org/10.1037/0033-295X.97.2.253

Kunde, W., & Wühr, P. (2006). Sequential modulations of correspondence effects across spatial dimensions and tasks. *Memory & Cognition*, *34*(2), 356–367. https://doi.org/10.3758/BF03193413

Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(2), 219–234. https://doi.org/10.1037/0096-1523.20.2.219

Liston, C., Matalon, S., Hare, T. A., Davidson, M. C., & Casey, B. J. (2006). Anterior Cingulate and Posterior Parietal Cortices Are Sensitive to Dissociable Forms of Conflict in a Task-Switching Paradigm. *Neuron*, *50*(4), 643–653. https://doi.org/10.1016/j.neuron.2006.04.015

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*(2), 163–203. https://doi.org/10.1037/0033-2909.109.2.163

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

Monsell, S., Sumner, P., & Waters, H. (2003). Task-set reconfiguration with predictable and unpredictable task switches. *Memory & Cognition*, *31*(3), 327–342. https://doi.org/10.3758/BF03194391

Notebaert, W., & Verguts, T. (2008). Cognitive control acts locally. *Cognition*, *106*(2), 1071–1080. https://doi.org/10.1016/j.cognition.2007.04.011

R Core Team. (2017). R. In *R Core Team* [Computer software]. https://doi.org/3-900051-14-3

Rajananda, S., Lau, H., & Odegaard, B. (2018). *A Random-Dot Kinematogram for Web-Based Vision Research* (1). *6*(1), Article 1. https://doi.org/10.5334/jors.194

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. *Trends in Cognitive Sciences*, *20*(4), 260–281. https://doi.org/10.1016/j.tics.2016.01.007

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), 585–590. https://doi.org/10.1038/nature12160

Ritz, H., & Shenhav, A. (2022). Humans reconfigure target and distractor processing to address distinct task demands. *BioRxiv*, 2021.09.08.459546. https://doi.org/10.1101/2021.09.08.459546

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of Object Representations in the Human Medial Temporal Lobe Based on Temporal Regularities. *Current Biology*, *22*(17), 1622–1627. https://doi.org/10.1016/j.cub.2012.06.056

Schmidt, J. R., & De Houwer, J. (2011). Now you see it, now you don't: Controlling for contingencies and stimulus repetitions eliminates the Gratton effect. *Acta Psychologica*, *138*(1), 176–186. https://doi.org/10.1016/j.actpsy.2011.06.002

Shenhav, A., Botvinick, M., & Cohen, J. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*(2), 217–240. https://doi.org/10.1016/j.neuron.2013.07.007

Simon, J. R., & Rudell, A. P. (1967). Auditory S-R compatibility: The effect of an irrelevant cue on information processing. *Journal of Applied Psychology*, *51*(3), 300–304. https://doi.org/10.1037/h0020586

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662. https://doi.org/10.1037/h0054651

Verguts, T. (2017). Binding by random bursts: A computational model of cognitive control. *Journal of Cognitive Neuroscience*. https://doi.org/10.1162/jocn_a_01117

Verguts, T., & Notebaert, W. (2008). Hebbian learning of cognitive control: Dealing with specific and nonspecific adaptation. *Psychological Review*, *115*(2), 518–525. https://doi.org/10.1037/0033-295X.115.2.518

Wühr, P., Duthoo, W., & Notebaert, W. (2015). Generalizing attentional control across dimensions and tasks: Evidence from transfer of proportion-congruent effects. *Quarterly Journal of Experimental Psychology*, *68*(4), 779–801. https://doi.org/10.1080/17470218.2014.966729

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