THEORETICAL NOTE

Hebbian Learning of Cognitive Control: Dealing With Specific and Nonspecific Adaptation

Tom Verguts and Wim Notebaert
Ghent University

The conflict monitoring model of M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, and J. D. Cohen (2001) triggered several research programs investigating various aspects of cognitive control. One problematic aspect of the Botvinick et al. model is that there is no clear account of how the cognitive system knows where to intervene when conflict is detected. As a result, recent findings of task-specific and context-specific (e.g., item-specific) adaptation are difficult to interpret. The difficulty with item-specific adaptation was recently pointed out by C. Blais, S. Robidoux, E. F. Risko, and D. Besner (2007), who proposed an alternative model that could account for this. However, the same problem of where the cognitive system should intervene resurfaces in a different shape in this model, and it has difficulty in accounting for the Gratton effect, a hallmark item-nonspecific effect. The authors of the current article show how these problems can be solved when cognitive control is implemented as a conflict-modulated Hebbian learning rule.

Keywords: cognitive control, Hebbian learning

Cognitive control refers to the requirement to repress one's instantaneous urges in favor of less obvious but more appropriate responses. Despite its importance in daily life, cognitive control has long remained an intractable concept. To confront this problem, Botvinick, Braver, Barch, Carter, and Cohen (e.g., Botvinick et al., 2001) proposed a quantified measure of conflict that can be used to exert top-down control. This was a significant step forward because it proposed an account of how the control system "knows" when to intervene on bottom-up processing routes.

The success of the Botvinick et al. (2001) model stems partly from its conceptualization of conflict but also from its ability to account for extant data patterns obtained in congruency tasks. In a Stroop task, for instance, participants respond to the color of the stimulus and ignore the color-word information. Whereas congruency effects, such as the Stroop effect, are typically taken to be failures of selective attention (failing to ignore the irrelevant word information), variations in the size of congruency effects are often considered cognitive control effects. The most studied cognitive control effect is the Gratton effect. Gratton, Coles, and Donchin (1992) observed that the flanker effect was smaller after incongruent trials than after congruent trials, suggesting that control is exerted after an incongruent trial. This effect was later replicated for the Simon effect (e.g., Stürmer, Leuthold, Soetens, Schröter, & Sommer, 2002), the Stroop effect (e.g., Kerns et al., 2004), and

prime–target congruency effects (Kunde, 2003). The standard explanation for the Gratton effect is that relatively more attention is directed to the relevant information after an incongruent trial than after a congruent trial, reducing the congruency effect. It is important to note that several researchers demonstrated that the Gratton effect persists when stimulus and response repetitions are excluded from the analysis, indicating that the Gratton effect is not merely caused by confounded repetition effects (e.g., Akçay & Hazeltine, 2007; Notebaert, Gevers, Verbruggen, & Liefooghe, 2006; Notebaert & Verguts, 2006, 2007, in press; Ullsperger, Bylsma, & Botvinick, 2005; Verbruggen, Notebaert, Liefooghe, & Vandierendonck, 2006). These studies were originally motivated by the suggestion of Mayr, Awh, and Laurey (2003) that the Gratton effect was caused merely by an item-specific priming effect. Besides refuting Mayr et al.'s claim and more important for the present purpose, these findings show that the Gratton effect, and therefore cognitive control, is not entirely item specific: Control triggered by one item also influences processing of other items.

Botvinick et al.'s (2001) conflict monitoring theory is able to account for the Gratton effect by assuming that the cognitive system continuously monitors the level of conflict (in the anterior cingulate cortex; ACC), where conflict is defined as the energy in the response layer; put simply, the more response units are simultaneously active, the more conflict there is in the system. This conflict signal is transferred into a control signal that is used to adjust attention to the relevant and the irrelevant processing route (in dorsolateral prefrontal cortex). When conflict is detected (typically on incongruent trials), more attention is dedicated to the relevant processing route and less attention to the irrelevant processing route.

Although the Gratton effect is not entirely item specific and in this sense consistent with conflict monitoring theory, the effect was recently found to be task specific in a number of studies. In a task-switching design, Kiesel, Kunde, and Hoffmann (2006) mixed

---

a magnitude task (determining whether digits were smaller or larger than 5) with a parity task (determining whether digits were odd or even). When large numbers require left responses in the magnitude task and even numbers a left response in the parity task, 6 would be a congruent stimulus and 7 an incongruent one. In line with a task-specific control mechanism, an adaptation effect was observed for task repetitions (e.g., parity–parity) but not for task alternations (e.g., parity–magnitude). In a study by Brown, Reynolds, and Braver (2007), a number and a letter were presented on each trial and participants had to switch between letter and number processing. Although not emphasized or formally tested by Brown et al., the data (their Figure 4b) also suggest smaller congruency effects after incongruent trials than after congruent trials (Gratton effect) for task repetitions; in addition, their data suggest larger congruency effects after incongruent trials than after congruent trials (reversed Gratton effect) for task alternations. Finally, we (Notebaert & Verguts, in press) mixed two congruency tasks in which both relevant and irrelevant dimensions were different for the two tasks. On task repetitions, a regular Gratton effect was observed, whereas on task switches, a reverse Gratton effect was found.

The finding that cognitive control is task specific fits with conflict monitoring theory because it assumes that attention to a specific task is enhanced after the detection of incongruency. However, a conceptual problem is attached to the model because, although the model specifies when the control system should intervene, there is no clear account in the model of how the control system should know where to intervene. In other words, it has no information available to choose which task to pay more attention to in cases where an experiment has more than one relevant task (i.e., a task-switching design).

Besides the Gratton effect, there is another instance of cognitive control, the proportion congruency effect. Tzelgov, Henik, and Berger (1992) demonstrated that the Stroop effect was larger in a condition where congruent trials were more frequent than incongruent trials relative to a condition where incongruent trials were more frequent. In other words, the size of the congruency effect increases as the proportion of incongruent trials decreases. Botvinick et al.'s (2001) conflict monitoring theory is able to account for the proportion congruency effect because the control signal integrates conflict levels from all previous trials (i.e., not just the previous trial); hence, in a condition with a high proportion of incongruent trials, a strong control signal will be generated that diminishes the congruency effect overall. However, in a recent article, Blais et al. (2007) pointed out that the Botvinick et al. model cannot account for one version of the proportion congruency effect, which is the item-specific proportion congruency effect (ISPC; Jacoby, Lindsay, & Hessels, 2003). Jacoby et al. administered a Stroop task in which some stimuli or stimulus features (e.g., the color red) are presented with a high-congruency proportion and other stimuli or stimulus features are presented with a low-congruency proportion. They demonstrated that the congruency effect was larger for stimuli with a higher congruency proportion. This shows that cognitive control is not only task specific but also (in this particular sense) item specific. However, the view that the ISPC is a manifestation of cognitive control has recently been challenged (Schmidt & Besner, in press). These authors put forward an alternative contingency account of the ISPC and argued that the larger congruency effect for stimuli with a high-congruency proportion is due to word-response contingencies in the experimental design: For example, if the word *red* is often presented in red (high-congruency

proportion), the incongruent combinations (e.g., the word *red* in blue) will be presented infrequently, slowing response times (RTs) on these items. To confront this possible confound, we note that the ISPC is a special case of the more general context-specific proportion congruency effect (CSPC), according to which contextual cues can lead to systematic fluctuations in cognitive control (Crump, Gong, & Milliken, 2006). The ISPC is a special case because the color and/or word information function as contextual cues that lead to such fluctuations. In a Stroop task, Crump et al. (2006) showed that other contextual cues such as location can lead to different sizes of the congruency effect. In this case, each location is equally associated with the different colors, words, and responses. This CSPC is not compatible with the contingency account of Schmidt and Besner (in press), unless one argues that it is the word-location-response contingencies that matter. However, Crump, Jamieson, and Milliken (2007) have shown that if different locations are associated with different proportions of congruency, the CSPC transfers to other items presented at these locations; it is important to note that these transfer items can be presented with equal frequencies at these locations and yet generate a CPSC. Hence, although we do not question that the contingency account captures some proportion of the variance, we doubt that it is the whole story and are confident that a control account is also needed to account for the CSPC.

Following up on Blais et al. (2007), we focus on the ISPC effect, although the principles also hold for the more general CSPC effect. Blais et al. introduced a variation of the conflict monitoring model where control is exerted on specific items rather than on an entire route. Although this seems to be a valuable solution for the ISPC, this alternative model is not able to account for the item-nonspecific Gratton effect, one of the key empirical findings favoring the original conflict monitoring model. Blais et al. argued that a combination of both models might be required to account for both the ISPC effect and the item-nonspecific Gratton effect. Besides being less parsimonious than the original model, the solution proposed by Blais et al. retains the conceptual problem inherent in the Botvinick et al. (2001) model discussed above: If cognitive control is item specific, how does the system know which item caused the conflict (i.e., where to intervene)? To tackle this problem, we found inspiration in the reinforcement learning literature, which is briefly described next.

Learning and Cognitive Control

The recent interest in the ACC and its role in cognitive control is accompanied by a parallel literature interested in reinforcement learning (e.g., Montague, Hyman, & Cohen, 2004; Roelfsema & Van Ooyen, 2005) and the role of the ACC in reinforcement learning (e.g., Cohen, McClure, & Yu, 2007; Holroyd & Coles, 2002; Seymour et al., 2004). Reinforcement learning is typically implemented as a modulated Hebbian learning rule. A typical modern reinforcement learning rule specifies weight updates of the form

$$\Delta w_{ij} = \lambda \times \delta \times x_i \times x_j, \tag{1}$$

where λ is the learning rate, δ is a reinforcement factor, and x_i and x_j are the activation of the sending and receiving neurons, respectively (Roelfsema & Van Ooyen, 2005). The reinforcement factor δ is typically contrastive in the sense that current reward (or predicted reward) is compared with some expected value (e.g., temporal difference learning; Montague et al., 2004; Sutton &
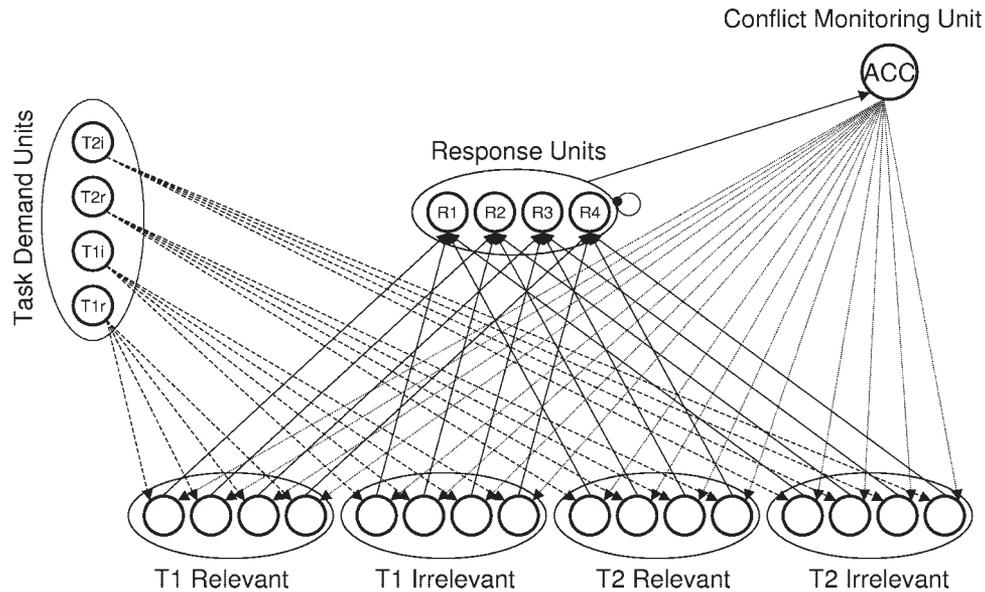
Figure 1. Outline of the model in a task-switching situation as implemented here. There are two task-relevant input layers for the first (T1) and second (T2) tasks, respectively, and two task-irrelevant input layers for T1 and T2, respectively. ACC = anterior cingulate cortex.

Barto, 1998). In the context of experimental tasks such as the Stroop task, the contrastive reinforcement signal is typically taken to be derived from two estimates of probability of success (e.g., Holroyd & Coles, 2002). However, given the very low number of errors that are usually made in such tasks, it is unlikely that reliable error probability estimates can be constructed and adaptively used by the cognitive system. For instance, in Notebaert and Verguts (in press), an error rate of only 4% (3.38% on congruent trials and 4.59% on incongruent trials) was observed. More generally, not only actual errors and error estimates are informative as to how well one is doing but also other, more online measures such as experienced conflict or response time are useful. For this reason, we here consider a *conflict-modulated Hebbian learning rule* instead of a reinforcement-modulated Hebbian learning rule. Our learning rule is similar to Equation 1 (see the Appendix, Equations A3 and A4), but the factor δ is based on conflict rather than error (likelihood). The learning rule is also contrastive in the sense that it compares the current level of conflict with the average level of conflict experienced up to that trial.

An outline of our model is shown in Figure 1. Because we deal with the conceptual problem of how the system knows where to intervene, we implemented a task-switching study to simulate task- and item-specific control effects; for this reason, the figure also represents a task-switching context. The input layers for two tasks are shown (denoted T1 and T2). Each of the two tasks has both a relevant and an irrelevant input layer. For example, in a Stroop task, the relevant input layer of one task might code for the color of the stimulus and the irrelevant input layer might code for the word. Task demand units bias individual items in their corresponding input layers. In addition, a conflict monitoring unit calculates a running (weighted) average of conflict across trials; hence, this unit integrates priming (influence of the previous trial) and more long-lasting context effects (influence of earlier trials). The learn-

ing rule uses the conflict monitoring signal to adapt the top-down influence from the task demand units to the units from the input layers. In the learning rule, a small constant is subtracted from the activation of the task demand unit so that this term becomes negative if the activation is small (see the Appendix, Equations A3 and A4). The model inherits the idea of conflict configuring the cognitive system from the Botvinick et al. (2001) model and the item specificity of weight adaptations from the Blais et al. (2007) model. Note also that the learning rule is completely local: All information necessary for weight adaptation (input activation, output activation, conflict term) is locally available at the synapse.

If, on a given trial, the level of conflict is high relative to earlier trials, this provides a signal that task-relevant connections should be strengthened. This is achieved by the conflict-modulated Hebbian learning rule that strengthens connections to the extent that the corresponding input and output units are both active and conflict is high. Because task-relevant connections are connected to an active task demand unit, these connections are strengthened, especially for the weights attached to the current item, because the activation of this item is elevated. However, task-irrelevant connections are connected to an inactive task demand unit and are therefore weakened. In this way, the model solves the conceptual problem outlined above because the model does not need to be told which connections should be strengthened or weakened.

Simulations

A task-switching situation was implemented in which simulated participants switched randomly between one of two tasks.[1] The

---

[1] Documentation and source code files can be downloaded from http://users.ugent.be/~tverguts/.

two tasks used different relevant and irrelevant input dimensions (see Figure 1; cf. Notebaert & Verguts, in press). There were 200 simulated participants; each simulated participant completed 900 trials. After each trial, activation in the task demand, input, and response layers was reset at zero.

In the RT analyses, errors were discarded from the analyses. For the Gratton effect, task repetition trials and task-switching trials are analyzed separately. In the Gratton task repetition analyses, trial-to-trial stimulus repetitions (either on the relevant or on the irrelevant dimension) are discarded. This is the same procedure as in recent behavioral experiments on the Gratton effect to exclude low-level priming effects (e.g., Kerns et al., 2004). The ISPC effects were calculated on task repetitions only without removal of trial-to-trial stimulus repetitions to match it as closely as possibly to the usual method of ISPC calculation in single-task situations. Including task switches does not change the results. In Simulation 1, we used four task-relevant stimulus values for each task, each presented with 50% congruency. In Simulation 2, we again used four task-relevant stimulus values for each task, but Stimulus Value 1 was presented with 80% congruency, Stimulus Values 2 and 3 with 50% congruency, and Stimulus Value 4 with 20% congruency. On congruent trials, the corresponding task-relevant and task-irrelevant input units were activated (e.g., in a Stroop

context, the color red and word *red* at the input layer); on incongruent trials, if, for example, the first input unit (e.g., the color red) was task relevant, then a task-irrelevant unit was chosen among the other three units (e.g., the word *blue*, *green*, or *yellow*).

In each simulation, we calculated the percentage of models that exhibited a particular effect as $n_+/(n_+ + n_-) \times 100$, where n_+ denotes the number of models (out of 200) that showed the effect (e.g., Gratton) and n_- the number that showed the reverse effect. This was done to ensure that the no-effect null hypothesis (i.e., the baseline for comparison) always predicts a percentage of 50%. A model (a) was considered to exhibit a Gratton effect if the congruency effect was larger after an incongruent trial than after a congruent one and (b) was considered to exhibit an ISPC (in Simulation 2) if the congruency effect for mostly congruent stimuli was larger than for neutral stimuli (50% congruency) and the congruency effect for neutral stimuli was larger than for mostly incongruent stimuli.

Simulation 1: Equal Proportions

The results are plotted in the first row of Figure 2. In line with previous studies (e.g., Kerns et al., 2004; Notebaert et al., 2006), the model exhibits a Gratton effect on task repetitions (i.e., as in
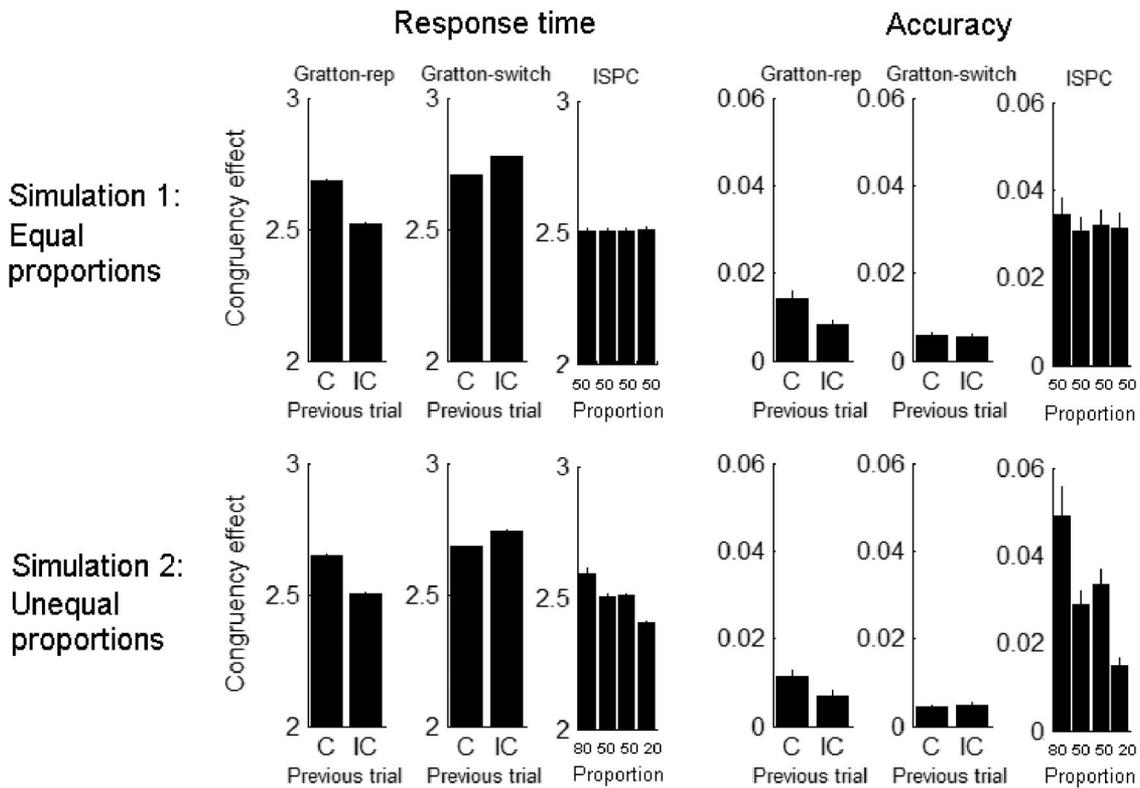

Figure 2. Simulation data: Congruency effects for response times and accuracies. Gratton-rep = Gratton effect on task repetition trials; Gratton-switch: Gratton effect on task-switching trials; ISPC = item-specific proportion congruency; C = congruent, IC = incongruent. Error bars denote 2 standard errors of measurement (calculated as the standard deviation of the effect across model replications divided by the square root of the number of replications), although in some instances they are very small and therefore not visible. The labels *C* and *IC* on the *x*-axes of the Gratton graphs refer to the congruency of the previous trial. The labels on the *x*-axes of the ISPC graphs refer to the proportions of congruency of each stimulus type.

single-task situations) in RT and accuracy data, even after removing (relevant or irrelevant dimension) stimulus repetitions. Eighty-nine percent of the models exhibited a Gratton effect (i.e., smaller congruency effect after incongruent trial) in the RT data and 64% exhibited a Gratton effect in the accuracy data.

The reason why the model generates an (item-nonspecific) Gratton effect is that there is always some baseline activation in the input layers (see the Appendix for equations). This baseline activity is not an isolated assumption to account for the Gratton effect: The assumption was already needed in the Botvinick et al. model to account for autocorrelation and posterror slowing (see Botvinick et al., 2001, p. 642). Because of this baseline activity, connection weights corresponding to task-relevant but nonpresented colors are strengthened slightly also.

On task-switching trials and also in line with earlier studies (Brown, Reynolds, & Braver, 2007; Notebaert & Verguts, in press), the Gratton effect is reversed. Eighty-one percent of the simulated models exhibited this trend in the RT data and 49% exhibited the trend in the accuracy data. The reason why there is a reversed Gratton effect on task-switching trials (in the RT data) is the following: Suppose that Task 1 was used on trial $n - 1$ and the trial was incongruent. This increases attention (through the conflict-modulated Hebbian learning rule) on the Task 1–relevant dimension and decreases attention to the other dimensions, including the Task 2–relevant dimension. It also decreases attention to the Task 2–irrelevant dimension; however, because this dimension is relatively inactive due to the fact that an irrelevant dimension is never strengthened during the task, the decrease in attention to the Task 2–irrelevant dimension is less pronounced. The net result is that attention to the Task 2–relevant dimension is decreased most strongly after an incongruent trial on Task 1, and therefore performance on Task 2 suffers from incongruency on Task 1. There is obviously no ISPC effect in Simulation 1, as each stimulus is presented with 50% congruency.

Simulation 2: Unequal Proportions

The results are plotted in the second row of Figure 2. There is again a Gratton effect for task repetitions (87% of all models in the RT data, 67% in the accuracy data) and a reversed Gratton effect for task switches (81% of all models in the RT data, 52% in the accuracy data). In this case, there is also an ISPC effect: The congruency effect is larger for stimuli with higher congruency proportions (88% and 85% in RT and accuracy data, respectively). There is an ISPC effect in this model because if an item is presented with a low proportion of congruency (i.e., high levels of conflict), the model increases its top-down influence for this item.

We have focused on two critical empirical phenomena here. However, because of its similarity with the Botvinick et al. (2001) and Blais et al. (2007) models, the current model also inherits many other properties of these models (e.g., listwise proportion congruency effect; see Blais et al., 2007).

General Discussion

We have shown how cognitive control may be conceptualized as an instance of a modulated Hebbian learning based system. This provides for conceptual links between reinforcement learning accounts of ACC activity and cognitive control on the one hand (e.g.,

Holroyd & Coles, 2002) and the conflict monitoring theory of Botvinick et al. (2001) on the other. Each approach can explain aspects of the empirical data that the other approach is incapable of; for example, the reinforcement learning approach can account for the shift of the ERN to the earliest predictor of an error and conflict monitoring theory can account for ACC activity in the absence of error. Hence, each model appears to provide only an incomplete account of cognitive control, and a unified framework integrating both accounts is in order. An initial step toward such integration was already taken by Botvinick (2007) who, consistent with our proposal, argued that conflict can be used as a learning signal. Botvinick formulated a skeletal model of how this can be achieved, but his model did not include a Hebbian component and was not able (or intended) to account for different specific and nonspecific aspects of cognitive control that we have captured with our learning rule. In this sense, our proposal can be seen as a next step building on his earlier conceptualization.

Besides integrating two views on ACC activity, our model also has the potential of integrating two theoretical frameworks proposed to account for effects reported in the cognitive control literature, such as the Gratton effect or the ISPC. According to one view (e.g., Botvinick et al., 2001), these effects are reflections of a cognitive control system attempting to optimize task behavior. According to a different view, these effects are caused by (low-level) stimulus–stimulus or stimulus–response bindings that are confounded with control-related experimental manipulations (e.g., Mayr et al., 2003; Schmidt, Crump, Cheesman, & Besner, 2007). In the present framework, all of these effects are conceived of as instantiations of Hebbian learning: If the Hebbian learning rule is modulated by conflict, it reflects cognitive control, whereas if the learning rule is unmodulated (and operates between stimulus and/or response representations), it reflects binding. It has been experimentally validated that each account is at least partially valid (e.g., Akçay & Hazeltine, 2007); future modeling in the unified framework is needed to understand their interactions and relative contributions to experimental data.

In attempting to integrate Hebbian learning and cognitive control, our model relates to Brown and Braver's (2005) account of ACC activity in congruency tasks. These authors proposed that stimulus elements (e.g., colors) become connected to ACC cells depending on their correlation with error. Brown and Braver's account, like ours, is a modulated Hebbian approach of how connections between stimulus elements and control units change over time to produce cognitive control phenomena. However, their model also differs in important respects; for example, their learning rule is based on error (likelihood) only, does not contrast obtained with expected outcomes, and modifies bottom-up rather than top-down changes. Another model that is related to the present one is the task-switching model of Gilbert and Shallice (2002), which also applied Hebbian learning to connections between input and task demand units to explain selected aspects of cognitive control phenomena. Their model, however, did not use modulated Hebbian learning and assumed that Hebbian changes were remembered in the network for one trial only. The reason for the latter assumption was that connection weights would grow without bound if Hebbian learning accrued over many trials; note that this is automatically prevented in our modulated Hebbian learning framework because weights can either increase or decrease due to the δ term. Also, their model was concerned with

explaining task switch costs rather than incongruency costs (as is the present model). A model that accounts for both switch costs and incongruency costs was recently formulated by Brown et al. (2007). Their model combines an (unmodulated) Hebbian learning rule in conjunction with a top-down incongruency detector and a (task and response) switch detector. It remains to be seen whether a general principle (such as Hebbian learning) can be formulated that accounts for switch costs and incongruency costs equally as well as the Brown et al. model.

How Is Control Exerted?

We have argued that control is implemented by changing connection weights from task demand units to the input layers. An alternative suggestion, also raised by Botvinick et al. (2001), is that control operates by decreasing the baseline level of response units, thus slowing down responding and improving accuracy. This is an option that we did not pursue in the present article. Yet, as for the exact form that δ takes in the learning rule, we are not committed to the idea that increasing attention to input layers is the only way to exhibit cognitive control. It is feasible that also updating the response unit baseline levels is governed by a modulated Hebbian learning rule; in that case, one predicts item-specific adaptations that nevertheless have an influence on the baseline of other items, similar to what happens in the Gratton effect. In particular, because posterror slowing has been associated with decreasing baseline levels by Botvinick et al., it could be that posterror slowing is partly item specific and partly not. This remains to be tested empirically.

Reinforcement Learning and Conflict Learning

What the conflict-modulated Hebbian learning rule does in the present task setting is to place more emphasis on the task-relevant route when needed. When this is not needed, the model allows for using the easier, task-irrelevant route. Consider the following analogy: You are using your new mobile phone and realize that it shares key combinations with your old phone. It also has some new key combinations. For shared (congruent) combinations, there is no use in learning new key combinations, and the old connections can profitably be generalized between tasks (mobile phones). So, if no conflict is experienced in using old combinations, you continue using these combinations. If, however, conflict is detected between an old and a new combination, more attention is dedicated to this specific new combination.

This example illustrates that a conflict-modulated Hebbian learning rule can help to optimize task performance but not to learn new tasks. For learning new tasks, reinforcement learning algorithms or other learning rules are necessary. As already mentioned, reinforcement learning rules can be formulated with exactly the same modulated Hebbian rule as used here (Equation 1), with different δ signals. Reinforcement-related δ signals are also thought to involve ACC: The recent interest in ACC and surrounding cortical areas has delivered a wide range of evaluative signals in ACC and surrounding mediofrontal cortex (e.g., conflict, pain, reward, error likelihood; for a review, see Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). We speculate that ACC delivers conflict and other evaluative δ signals to subcortical structures (e.g., locus coeruleus and dopaminergic ventral tegmen-

tal area or substantia nigra; Schultz, Dayan, & Montague, 1997). In this scenario, ACC would then function as a collector and propagator of different evaluative signals. Because of its proximity to motor, premotor, and orbitofrontal cortices, ACC is ideally located for this task. Because some of these subcortical structures project strongly to cortical areas (e.g., dopaminergic ventral tegmental area and substantia nigra to prefrontal areas; locus coeruleus also to posterior cortical areas), they may propagate the δ signal to cortical areas for implementation of the modulated Hebbian learning rule.

References

Akçay, C., & Hazeltine, E. (2007). Feature overlap and conflict monitoring: Two sources of sequential modulations. *Psychonomic Bulletin & Review, 14,* 742–748.

Blais, C., Robidoux, S., Risko, E. F., & Besner, D. (2007). Item-specific adaptation and the conflict monitoring hypothesis: A computational model. *Psychological Review, 114,* 1076–1086.

Botvinick, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulated function. *Cognitive, Affective, and Behavioral Neuroscience, 7,* 356–366.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Brown, J. W., & Braver, T. S. (2005, February 18). Learned predictions of error likelihood in the anterior cingulate cortex. *Science, 307,* 1118–1121.

Brown, J. W., Reynolds, J. R., & Braver, T. S. (2007). A computational model of fractionated conflict-control mechanisms in task-switching. *Cognitive Psychology, 55,* 37–85.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploration and exploitation. *Philosophical Transactions of the Royal Society of London, Series B, 362,* 933–942.

Crump, M. J. C., Gong, Z., & Milliken, B. (2006). The context-specific proportion congruent Stroop effect: Location as a contextual cue. *Psychonomic Bulletin & Review, 13,* 316–321.

Crump, M. J. C., Jamieson, R. K., & Milliken, B. (2007). *Fluid improvisation and memory-driven cognitive control: Testing an instance-based account of context-specific learning and control.* Manuscript submitted for publication.

Gilbert, S. J., & Shallice, T. (2002). Task switching: A PDP model. *Cognitive Psychology, 44,* 297–337.

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.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109,* 679–709.

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

Kerns, J. G., Cohen, J. D., MacDonald, A. W., 3rd, Cho, R. Y., Stenger, V. A., & Carter, C. S. (2004, February 13). Anterior cingulate conflict monitoring and adjustments in control. *Science, 303,* 1023–1026.

Kiesel, A., Kunde, W., & Hoffmann, J. (2006). Evidence for task-specific resolution of response conflict. *Psychonomic Bulletin & Review, 13,* 800–806.

Kunde, W. (2003). Sequential modulations of stimulus–response correspondence effects depend on awareness of response conflict. *Psychonomic Bulletin & Review, 10,* 198–205.

Mayr, U., Awh, E., & Laurey, P. (2003). Conflict adaptation effects in the absence of executive control. *Nature Neuroscience, 6,* 450–452.

Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004, October 14). Computational roles for dopamine in behavioural control. *Nature, 431,* 760–767.

Notebaert, W., Gevers, W., Verbruggen, F., & Liefooghe, B. (2006). Top-down and bottom-up sequential modulations of congruency effects. *Psychonomic Bulletin & Review, 13,* 112–117.

Notebaert, W., & Verguts, T. (2006). Stimulus conflict predicts conflict adaptation in a numerical flanker task. *Psychonomic Bulletin & Review, 13,* 1078–1084.

Notebaert, W., & Verguts, T. (2007). Dissociating conflict adaptation from feature integration: A multiple regression approach. *Journal of Experimental Psychology: Human Perception and Performance, 33,* 1256–1260.

Notebaert, W., & Verguts, T. (in press). Cognitive control acts locally. *Cognition.*

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004, October 15). The role of the medial frontal cortex in cognitive control. *Science, 306,* 443–447.

Roelfsema, P. R., & Van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Computation, 17,* 2176–2214.

Schmidt, J. R., & Besner, D. (in press). The Stroop effect: Why proportion congruent has nothing to do with congruency and everything to do with contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Schmidt, J. R., Crump, M. J. C., Cheesman, J., & Besner, D. (2007). Contingency learning without awareness: Evidence for implicit control. *Consciousness & Cognition, 16,* 421–435.

Schultz, W., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science, 275,* 1593–1599.

Seymour, B., O'Doherty, J. P., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., et al. (2004, June 10). Temporal difference models describe higher-order learning in humans. *Nature, 429,* 664–667.

Stürmer, B., Leuthold, H., Soetens, E., Schröter, H., & Sommer, W. (2002). Control over location-based response activation in the Simon task: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 1345–1363.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning.* Cambridge, MA: MIT Press.

Tzelgov, J., Henik, A., & Berger, J. (1992). Controlling Stroop effects by manipulating expectations for color words. *Memory & Cognition, 20,* 727–735.

Ullsperger, M., Bylsma, L. M., & Botvinick, M. M. (2005). The conflict-adaptation effect: It's not just priming. *Cognitive, Affective, and Behavioral Neuroscience, 5,* 467–472.

Verbruggen, F., Notebaert, W., Liefooghe, B., & Vandierendonck, A. (2006). Stimulus and response conflict-induced cognitive control in the flanker task. *Psychonomic Bulletin & Review, 13,* 328–333.

Appendix

Model Equations

Time in a trial is indexed by t; the cascade rate of activation in a trial is denoted by τ. The activation equation for an arbitrary input unit i (in either the task-relevant or the task-irrelevant layer) is then as follows:

$$x_i^{in}(t + 1) = (1 - \tau)x_i^{in}(t) + \tau[I_i(t) + \beta_{in}]. \qquad (A1)$$

The indicator function $I_i(t)$ is 1 if the stimulus corresponding to that unit i is presented at time t, and 0 otherwise.

The activation equation for a response unit j is

$$x_j^{res}(t + 1) = (1 - \tau)x_j^{res}(t)$$

$$+ \tau\left\{\sum_i w_i^{ir}x_i^{in}(t)\left[C + \sum_{k=1}^{4} w_{ki}^{ti}(n)x_k^{td}(n)\right] + w^{inh}\sum_{k \neq j}x_k^{res}(t)\right\}. \qquad (A2)$$

The matrix $\boldsymbol{w}^{ir}$ contains bottom-up weights from the input layers to the response layer. The term $\left[C + \sum_{k=1}^{4} w_{ki}^{ti}(n)x_k^{td}(n)\right]$ incorporates the top-down attentional weighting from the task demand units to the input layers by weight matrix $\boldsymbol{w}^{ti}$, which is adaptively changed over trials. The term $w^{inh}\sum_{k \neq j}x_k^{res}(t)$ reflects response competition, and the summation $-w^{inh}\sum_j \sum_{k \neq j}x_j^{res}(t)x_k^{res}(t)$ over response units represents the total amount of conflict.

The activation equation for the control unit equals

$$x^{con}(n + 1) = \lambda_{con}x^{con}(n) + (1 - \lambda_{con})$$

$$\times \left(- w^{inh}\sum_j \sum_{k \neq j}x_j^{res}x_k^{res} + \beta_{con}\right).$$

This equation is applied at the end of each trial n. Finally, weights are adapted according to a conflict-modulated Hebbian learning rule:

$$w_{ki}^{ti}(n + 1) = \lambda_w w_{ki}^{ti}(n) + (1 - \lambda_w)(\alpha_w f + \beta_w). \qquad (A3)$$

The term f reflects the conflict-modulated Hebbian term

$$f = [x^{con}(n) - \overline{x^{con}}]x_i^{in}(t)[x_k^{td}(n) - 1/2], \qquad (A4)$$

where $\overline{x^{con}}$ denotes the mean activity of the control unit up to trial n. Weights w^{ti} are only adapted between task representations and their corresponding input layer units and are restricted to be nonnegative.

Parameters were chosen as follows: $\tau = 0.2$, $\beta_{in} = 0.2$, $w^{inh} = -0.5$, $C = 0.7$, $\beta_{con} = 1$, $\lambda_{con} = 0.8$, $\lambda_w = 0.7$, $\alpha_w = 20$, $\beta_w = 0.5$. The activation of the relevant task demand unit was set at 1 and that of the other task demand units (three units on each trial in the implemented simulations) at 0.3. The initial strength of each connection between a task demand unit and its corresponding input units (i.e., entries in matrix $\boldsymbol{w}^{ti}$) was 0.5. The strength of input–response connections for each of the two task-relevant input layers equals 1 (e.g., from color red to response "red" in a Stroop task; matrix $\boldsymbol{w}^{ir}$); the strength of input–response connections for each of the two task-irrelevant input layers equals 1.1 (e.g., from verbal word *red* to response "red" in a Stroop task). In each trial, activation of the input and response units was updated according to Equations A1 and A2 until one of the response units reached a threshold value of 0.8. The corresponding response was taken to be the model's response choice and the time needed to reach that unit was taken to be the model's RT. The qualitative pattern of results was robust to changes in these parameters.