Psychological Review
2007, Vol. 114, No. 4, 1076–1086

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.4.1076

THEORETICAL NOTE

Item-Specific Adaptation and the Conflict-Monitoring Hypothesis: A Computational Model

Chris Blais, Serje Robidoux, Evan F. Risko, and Derek Besner
University of Waterloo, Canada

M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, and J. D. Cohen (2001) implemented their conflict-monitoring hypothesis of cognitive control in a series of computational models. The authors of the current article first demonstrate that M. M. Botvinick et al.'s (2001) conflict-monitoring Stroop model fails to simulate L. L. Jacoby, D. S. Lindsay, and S. Hessels's (2003) report of an item-specific proportion-congruent (ISPC) effect in the Stroop task. The authors then implement a variant of M. M. Botvinick et al.'s model based on the assumption that control must be able to operate at the item level. This model successfully simulates the ISPC effect. In addition, the model provides an alternative to M. M. Botvinick et al.'s explanation of the list-level proportion-congruent effect in terms of an ISPC effect. Implications of the present modeling effort are discussed.

Keywords: conflict, conflict monitoring, control, Stroop

Botvinick, Braver, Barch, Carter, and Cohen's (2001) conflict-monitoring hypothesis proposes a neurally plausible mechanism for cognitive control. Evidence consistent with this hypothesis has accumulated across a number of behavioral, imaging (Carter et al., 2000; Kerns et al., 2004), and neurophysiological studies (Yeung, Botvinick, & Cohen, 2004). Botvinick et al. (2001) implemented this idea in a series of computational models that successfully simulated results from both behavioral and functional magnetic resonance imaging studies of control-related phenomena. Here, we extend Botvinick et al.'s conflict-monitoring model of Stroop performance to another control-related phenomenon heretofore unexamined from the conflict-monitoring perspective. Although the initial modeling effort was successful, it requires a critical change to the model's architecture. This change has important implications for a more general understanding of cognitive control.

Cognitive Control

Cognitive control is required to adapt behavior to situational demands. The most common example of cognitive control comes from studies of the Stroop task (see MacLeod, 1991). In the Stroop task, participants are presented with color words printed in various colors (i.e., the word *red* in BLUE) and asked to respond to the print color of the color word. Here, participants must suppress their habitual tendency to read the word in order to perform the less practiced task of naming the print color. Lack of cognitive control would lead to behavior dominated by habitual responses (i.e., reading the color word).

The Conflict-Monitoring Hypothesis

The conflict-monitoring hypothesis posits a neural mechanism that supports the adaptive control of behavior in situations like the Stroop task. According to the conflict-monitoring hypothesis, a specific subsystem of the human brain, the anterior cingulate cortex (ACC), responds to conflict in information processing. The detection of conflict triggers adjustments in cognitive control (e.g., by reinforcing goal representations) via the prefrontal cortex in order to reduce conflict in subsequent performance (Botvinick, Cohen, & Carter, 2004). For example, on an incongruent Stroop trial, the color word and the display color produce response conflict. This response conflict is detected by the ACC, which signals the prefrontal cortex that is responsible for enhancing top-down control (e.g., by enhancing the strength of goal representation), thus reducing conflict (Botvinick et al., 2001, 2004).

The Conflict-Monitoring Model of Stroop Performance

Botvinick et al. (2001) implemented their conflict-monitoring hypothesis in a series of neural network simulations. Their approach consisted of adding a conflict-monitoring module, representing the ACC, to extant computational models of various tasks (e.g., Stroop). In these simulations, activation within the conflict-monitoring module increased or decreased as a function of the amount of conflict between response units in the neural network. Using this basic architecture, Botvinick et al. simulated both behavioral manifestations of cognitive control and brain activation in the ACC.

Chris Blais, Serje Robidoux, Evan F. Risko, and Derek Besner, Department of Psychology, University of Waterloo, Canada.

This work was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC) Graduate Scholarship to Evan F. Risko and by NSERC Grant A0998 to Derek Besner. We thank Ilona Medvedovsky for programming assistance.

Correspondence concerning this article should be addressed to Chris Blais, who is now at the University of California, Berkeley, Department of Psychology, Helen Wills Neuroscience Institute, 132 Barker Hall, Berkeley, CA 94720. E-mail: cblais@berkeley.edu

For the present purposes, the most relevant of these simulations is the one relating to the Stroop effect. To simulate the Stroop effect, Botvinick et al. (2001) used a version (Cohen & Huston, 1994) of Cohen, Dunbar, and McClelland's (1990) popular model of performance in the Stroop task. A brief overview of Cohen et al.'s (1990) model appears here, but we refer the reader to the original article for a more complete treatment.

Cohen et al.'s (1990) model (see Figure 1) consists of a color-processing pathway and a word-processing pathway. These pathways feed activation through to a common set of response units. It is the competition between these response units that Botvinick et al. (2001) used as an index of conflict in information processing. Given that we are more practiced at reading words than naming colors, the connection weights between the word units and the response units are stronger than those between the color units and the response units. This difference in the connection weights gives rise to the asymmetrical interference pattern where words interfere with color naming more than colors interfere with word naming. This asymmetry in the strength of connections along the two processing pathways creates the need for a selective attention mechanism; otherwise the model would simply read the word on every trial. In Cohen et al.'s model, selective attention is instantiated through a set of task demand units, subsequently argued to reside in the prefrontal cortex (e.g., Botvinick et al., 2004), that serve to strengthen activation along the relevant pathway. In Botvinick et al.'s (2001) model, the conflict monitor sends a signal to the task demand units altering its input to the color pathway thus strengthening the connection between input and response node.

Botvinick et al. (2001) first demonstrated that their conflict-monitoring module, adjoined to the response units in the Cohen et al. (1990) framework, behaved in a fashion similar to that found in brain activation studies of the ACC. For example, they demonstrated that activity in the conflict module was higher in the incongruent condition (e.g., the word *red* in BLUE) than in the neutral (e.g., the word *house* in BLUE) and congruent conditions (e.g., the word *blue* in BLUE). Numerous imaging studies have reported a similar pattern of brain activation in the ACC while participants performed a Stroop task (e.g., Carter et al., 2000; Kerns et al., 2004).

In a second set of simulations, Botvinick et al. (2001) demonstrated that this conflict-monitoring module could serve as a signaling device for control-related processes to adapt to the presence

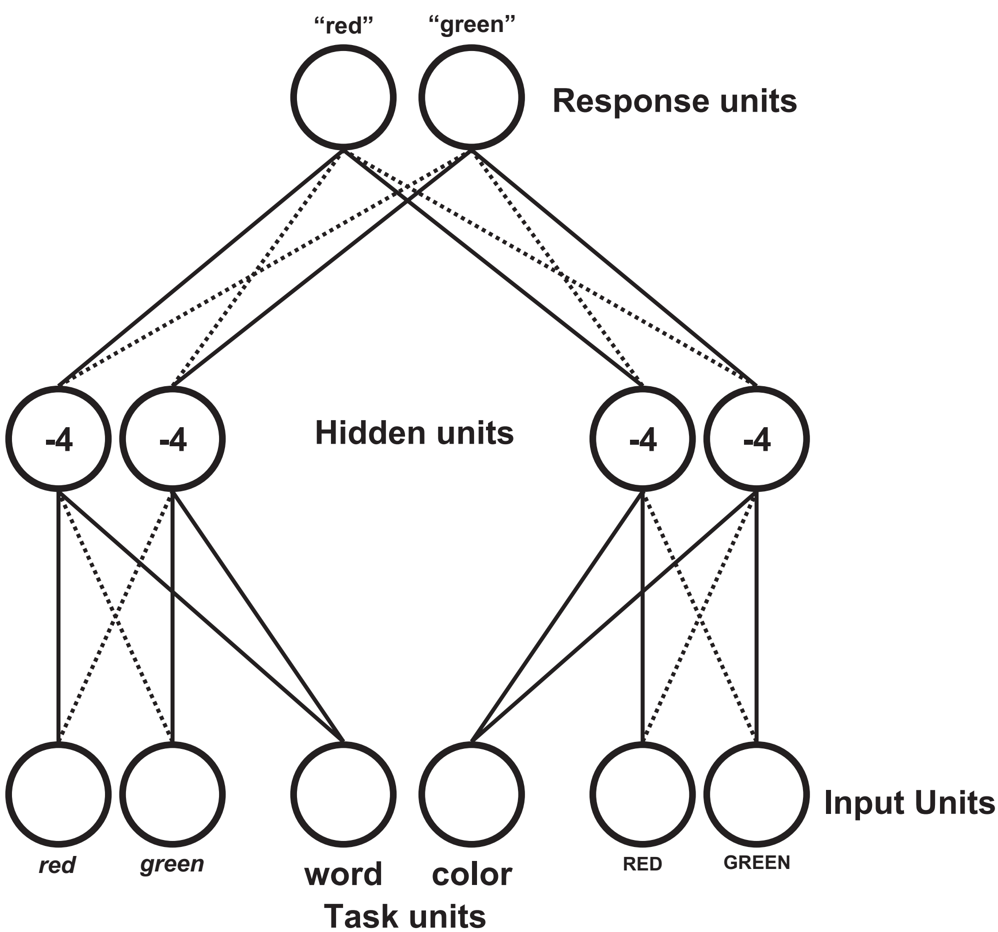

Figure 1. Representation of the Cohen et al. (1990) model of Stroop performance. Adapted from "On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect," by J. D. Cohen, K. Dunbar, and J. L. McClelland, 1990, *Psychological Review, 97,* p. 336. Copyright 1990 by the American Psychological Association.

of conflict in information processing. On each trial, the conflict-monitoring module provided a control signal to the task demand units. High levels of conflict increased the input from the task demand units to the color pathway, whereas low levels of conflict reduced the input from the task demand units.[1]

Botvinick et al. (2001) simulated the proportion-congruent effect, a prominent control phenomenon in the Stroop literature. This effect refers to the observation that the relative proportion of congruent trials in an experimental list affects the magnitude of the Stroop effect (e.g., Cheesman & Merikle, 1986; Lindsay & Jacoby, 1994; Logan & Zbrodoff, 1979; Tzelgov, Henik, & Berger, 1992). That is, the magnitude of the Stroop effect *increases* as the proportion of congruent trials *increases*. Botvinick et al. (2001) demonstrated that their model simulates this pattern of data. As the proportion of congruent trials increases, the amount of conflict in the system decreases, leading to a decrease in top-down control (i.e., input from the task demand unit to the color pathway). This decrease in top-down control results in a larger Stroop effect because the word pathway exerts a stronger influence on the response nodes. Thus, the implementation of the conflict-monitoring hypothesis in the form of a computational neural network model accounts for cognitive control in both brain activation and behavioral studies.

In the present investigation, we attempted to model a more recent instance of cognitive control, the item-specific proportion-congruent (ISPC) effect (Jacoby, Lindsay, & Hessels, 2003), in the context of the conflict-monitoring hypothesis and more specifically within Botvinick et al.'s (2001) conflict-monitoring model of Stroop performance.

The ISPC Effect

Jacoby and colleagues (Jacoby et al., 2003; Jacoby, McElree, & Trainham, 1999; Trainham, Lindsay, & Jacoby, 1997) have demonstrated that the magnitude of the Stroop effect is affected by the proportion of congruent trials at the *item level*. For example, Jacoby et al. (2003) varied the proportion of congruent trials at the level of specific items (i.e., a specific pairing of a color word and color) while keeping the list-level congruency proportion at 50% and found that items that were "mostly congruent" showed a larger Stroop effect than items that were "mostly incongruent." They referred to this item-level effect as the ISPC effect. Jacoby et al. (2003) claimed that a control mechanism like Cohen et al.'s (1990) and by association Botvinick et al.'s (2001), where a single task demand unit controls attention to all items within a processing pathway, could not account for this effect.

The reason Cohen et al.'s (1990) and Botvinick et al.'s (2001) models encounter difficulty in explaining the ISPC effect is because control operates at the level of the color pathway, rather than at the level of individual color representations. As previously discussed, Botvinick et al.'s (2001) model simulated the list-level proportion-congruent (LLPC) effect with a conflict module that signaled the task demand unit to increase or decrease input to the color pathway. Critically, increasing or decreasing the task demand unit's input to the color pathway affects the processing of all colors. Implementing control in this manner is unlikely to be able to explain the ISPC effect because the ISPC effect suggests that control can occur at the level of individual items.

Rather than discarding Botvinick et al.'s (2001) model of control in Stroop performance, we suggest that an amendment to this model will allow it to account for the ISPC effect. Specifically, we suggest that control has to be implemented at the item level.

Present Investigation

The present investigation tests the hypothesis that a conflict-monitoring model based on Cohen et al.'s (1990) model of Stroop performance can account for the ISPC effect when control is implemented at the item level but not when control is implemented at the pathway level.

In the first set of simulations, we added a conflict module to Cohen et al.'s (1990) model of Stroop performance and attempted to simulate the ISPC effect with control acting at the pathway level. As suggested, this approach failed to simulate the ISPC effect. Next, we implemented control at the item level and demonstrated that the model was then able to simulate the ISPC effect. Thus, by implementing control at the item level rather than the pathway level we were able to simulate the ISPC effect using the principles outlined in the conflict-monitoring hypothesis.

Computational Properties of the Model

Architecture

Cohen et al.'s (1990) model has already been briefly described above. Our implementation of this part of the model did not include a learning algorithm.[2] As such, the connection weights were hard-coded, rather than learned (cf. Cohen et al., 1990), to values that had been used successfully in other simulations with four response alternatives (Kanne, Balota, Spieler, & Faust, 1998). These are shown in Table 1. The remaining details of this model are not critical to the implementation of control but are presented in the Appendix. What follows is a discussion of the modifications introduced by Botvinick et al. (2001) and those introduced by us for the current simulations.

Botvinick et al.'s (2001) most significant change to Cohen et al.'s (1990) model was the introduction of a control-monitoring device. This device relies on a measure of the "energy" in the response units (Hopfield, 1982) to determine the degree of conflict resulting from a given trial. Hopfield's formula is provided in Equation 1:

$$\text{Hopfield Energy} = E_H = -\sum_i \sum_j a_i a_j w_{ij}. \qquad (1)$$

[1] The reader might wonder why the participant would not simply modulate attention to the relevant dimension in such a way as to completely eliminate conflict. The question, however, depends on the assumption that removing attention from the word dimension would stop processing of the word. This view is inconsistent with the claim that word processing can occur in the absence of attention (Brown, Gore, & Carr, 2002; Cohen, Dunbar, & McClelland, 1990). For example, in Cohen et al.'s (1990) Stroop model, removing attention entirely from the word dimension does not eliminate the Stroop effect. Furthermore, participants may perceive the irrelevant dimension as useful and thus choose to attend to it at the cost of increasing conflict (Dishon-Berkovits & Algom, 2000).

[2] The source code can be downloaded from http://www.arts.uwaterloo.ca/~dbesner2/blais2007/

Table 1
Connection Weights Used in the Simulations

Connection weight	Excitatory	Inhibitory
Task units	4.0	
Words		
Input to hidden units	3.1	−2.3
Hidden to response units	3.3	−2.6
Colors		
Input to hidden units	2.5	−2.1
Hidden to response units	1.8	−1.5

E_H then acted as the input to a formula that determined the extent of control to be exerted (Equation 2):

$$\text{Control} = C(t+1) = \lambda C(t) + (1-\lambda)(\alpha E(t) + \beta), \quad (2)$$

where $C(t)$ is control on trial t, $E(t)$ is the energy at the end of trial t, λ is a weighting parameter (restricted to values between 0 and 1) that indicates the degree to which control will change on each trial, and α and β are scaling parameters. Theoretically, parameters α and β are uninteresting as they simply scale the energy value to a range that will have a functional role at the control level. Conversely, the λ parameter is very important. It functionally operates as a memory "buffer." When λ is large, the memory buffer contains information about the amount of conflict on many previous trials and, as such, "history" dominates the control value. When this value is small, the memory buffer heavily weights the amount of conflict on the most recent trial and, as such, the "previous trial" dominates the control setting.

Consistent with this interpretation, when simulating LLPC effects, Botvinick et al. (2001) set λ to a large value (.95), presumably reflecting the dominant view in the field that proportion effects result from participants recognizing that the word is informative and thus actively reading it on every (or most) trial(s). Also, when simulating sequential effects, λ was set to a relatively smaller value (.50) consistent with the view that participants are affected by the previous trial's congruency. In all the present simulations, α was 21.01873, β was −5.41268, and λ was .50.

$C(t)$ provides the weight for the connections from the task demand unit to the relevant word/color units for the next trial, subject to the constraint that the connection weights must fall between 2.3 and 4.0. That is, at the end of each trial control is calculated according to Equation 3, and the weights on the relevant connections are then set to $C(t + 1)$. Details about which connections are adjusted and under what circumstances follow a brief discussion of the differences between Botvinick et al.'s (2001) approach and ours.

Due to the cascaded (as opposed to interactive activation) nature of our model (chosen for ease of implementation), certain adjustments to the equations specified above are required. First, as our model does not include within-level inhibitory connections, we adjusted Equation 1 to reflect the assumption that connections between response units are −1.0. Furthermore, the Hopfield formula as described in Equation 1 cycles over each pair of units twice. This makes sense only if connections between units are not symmetric (i.e., unit_i inhibits unit_j more than unit_j inhibits unit_i), however, when connections are symmetric between any two units, there is no need to compare each pair twice. The assumption of unit negative inhibition implies that all connections are symmetric so our formula calculates the energy for each pair only once. These changes represent minor adjustments to the original formula and give rise to (Equation 3):

$$\text{Energy} = E = \sum_{i} \sum_{j>i} a_i a_j. \quad (3)$$

As with Botvinick et al.'s (2001) approach, energy is calculated on the basis of activation in the response units. This energy provides the input to a formula that calculates control in the same way described by Botvinick et al. (2001; Equation 2).

Pathway-level control (Model P). Our implementation of pathway control (see Figure 2) differs slightly from that described in Botvinick et al. (2001). In lieu of adjusting the activation of the task demand unit, the connection weights between the task demand units and the relevant color or word units were modified. When implementing the item-specific mechanism, we found that changing the connection weights rather than activity in the task demand unit was more tractable. Implementing an item-specific mechanism via changing activity in the task demand unit would require a move from two task demand units representing word and color to eight separate task demand units representing the four color words and the four colors. The present option is less cumbersome and does not require a drastic change to the architecture of the underlying model. Mathematically these are equivalent operations as the effect on units in the subsequent layers is given by the weight of the connection multiplied by the activation strength (please refer to Paragraph 1 of the Appendix). Whether one doubles the activation level in the task demand unit or the weight between the task demand unit and the relevant color/word units, the net result will be a doubling of the input from the task demand unit.

Item-level control (Model I). Our implementation of item-level control (see Figure 3) involves modifying only the connection from the task demand unit to the hidden unit that was associated with the color that was relevant on that trial. For illustration, if the task is color naming, and the previous trial consisted of the word *red* presented in BLUE (an incongruent trial), $C(t)$ would be large, and thus the weight on the connection to the presented color (BLUE) is increased to reflect increased attention. Conversely if the previous trial was congruent (e.g., the word *red* presented in RED), $C(t)$ would be relatively small, and the weight on the connection to the presented color (RED) would be decreased. The operation of control is the only way in which Models P and I differ.

Simulation 1: The ISPC Effect

Here, we compare the ability of Model P and Model I to simulate the ISPC effect. In Model P, control is implemented at the pathway level following Botvinick et al. (2001), and in Model I control is implemented at the item level.

One thousand randomized lists of 192 (48 practice, 144 experimental) items were constructed such that the words *blue* and *red* appeared mostly incongruent and the words *yellow* and *green* appeared mostly congruent. The relationship between items and congruency was consistent within the practice and experimental trials. Table 2 shows the trial structure. These lists were presented to both versions of the model.

Analysis. Following Botvinick et al. (2001), for each list, the 48 practice trials served to allow the task weights in the model an

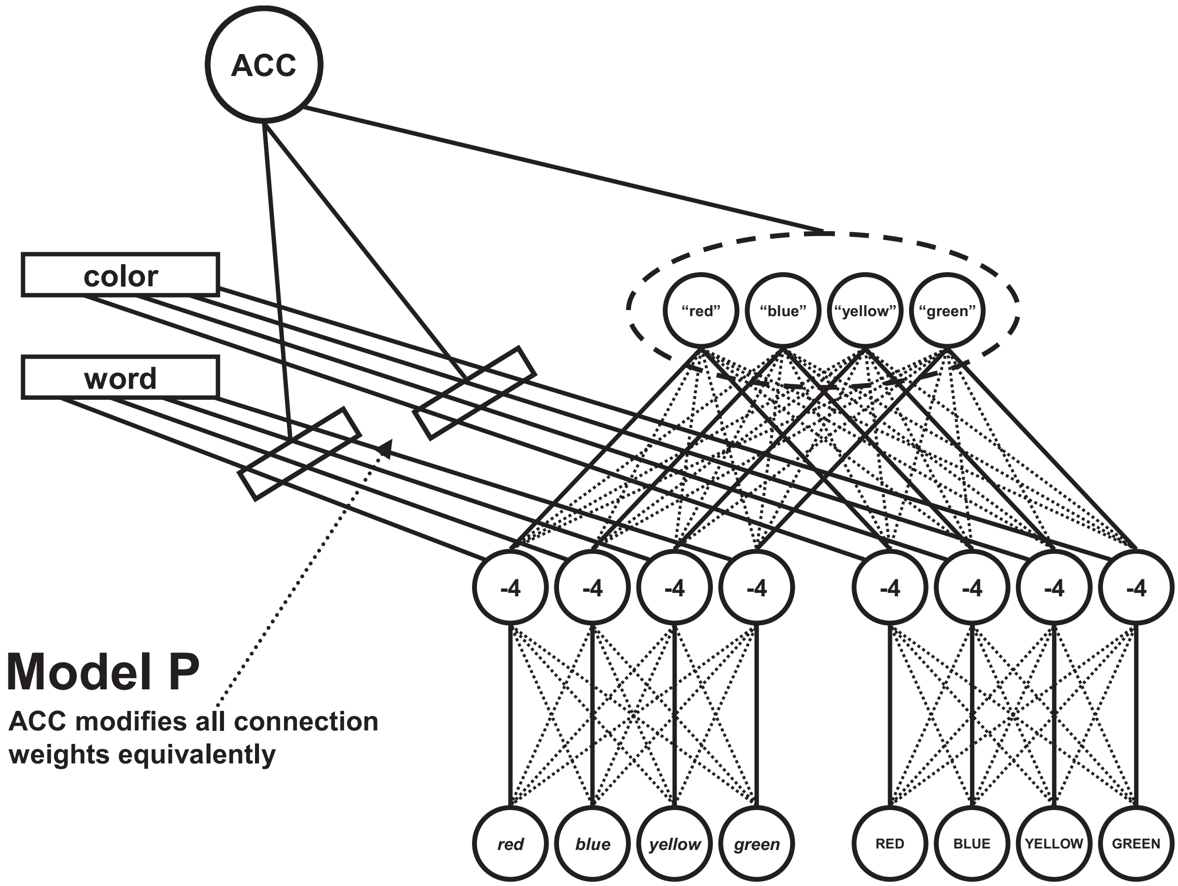

Figure 2. Representation of Model P, an implementation of pathway-level control. ACC = anterior cingulate cortex.

opportunity to settle. The remaining 144 experimental trials were analyzed for the presence of an ISPC effect. Table 3 shows the average in each condition across the 1,000 simulations for each version of the model.

As is evident in Figure 4, Model P shows no trace of an ISPC effect. Specifically, across the 1,000 simulations, the average size of the Stroop effect (incongruent–congruent) for the "mostly congruent" items is 164.2 cycles, and the average size of the Stroop effect for the "mostly incongruent" items is 163.5 cycles. Model I, however, shows an ISPC effect of 30.2 cycles. The size of the Stroop effect for the "mostly congruent" items is 191.5 cycles, and the size of the Stroop effect for the "mostly incongruent" items is 161.3 cycles.

Discussion. In Simulation 1, Model P failed and Model I succeeded in simulating the ISPC effect. The only difference between these two models was the level at which control was implemented. Not surprisingly, when control was implemented at the level of the pathway the model did not produce an ISPC effect. However, when control was implemented at the level of the item, the model readily produced an ISPC effect. Thus, with control implemented at the item level, the ISPC effect can be added to a growing list of control-related phenomena that are explicable in terms of a conflict-monitoring mechanism. In a second set of simulations, we tested the ability of the item-level mechanism just introduced to account for an LLPC effect.

Simulation 2: The LLPC Effect

Some researchers (Jacoby et al., 1999; Trainham et al., 1997) have suggested that the ISPC effect and the LLPC effect could emerge from a similar mechanism. Essentially, LLPC manipulations could be construed as item-level manipulations with all items sharing the same congruency proportion. Thus, it is possible that the list-level effect is simply an ISPC effect in disguise. We tested this possibility in Simulation 2 by simulating the LLPC effect using both Model P and Model I. The magnitude of the proportion-congruent effect in Model P relative to Model I should determine the extent to which the LLPC effect can be considered an ISPC effect, at least within the context of the present model.

One thousand randomized lists of 192 items each were constructed. In 500 of the lists, 25% of the items were congruent. In the remaining 500 lists, 75% of the items were congruent. Table 4 shows the trial structure for each of these lists. All of these lists were presented to each version of the model.

Analysis. As in Simulation 1, the first 48 trials in each list served to allow the task weights in the model an opportunity to settle. The remaining 144 trials in each list were analyzed for the presence of a proportion effect. Table 5 shows the average in each condition across the 1,000 simulations for each version of the model.

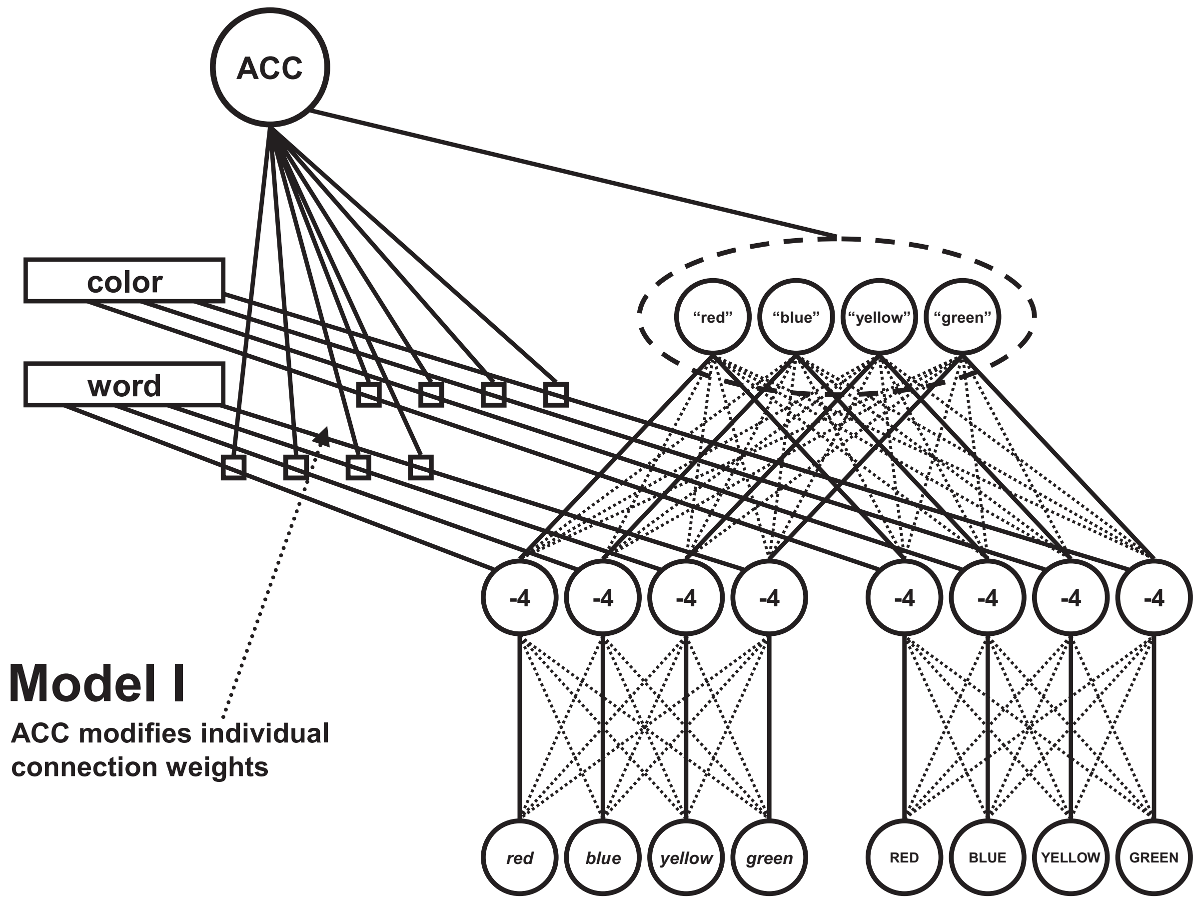

Figure 3. Representation of Model I, an implementation of item-level control. ACC = anterior cingulate cortex.

As is evident in Figure 5, both Model P and Model I show an LLPC effect. For Model P, the size of the Stroop effect for the 25% lists is 143.9 cycles, and the size of the Stroop effect for the 75% lists is 185.9 cycles. For Model I, the size of the Stroop effect for the 25% lists is 158.2 cycles, and the size of the Stroop effect for the 75% lists is 201.3 cycles. The magnitude of the interaction between congruency proportion (25% vs. 75%) and the magnitude of the Stroop effect (incongruent vs. congruent) is nearly identical for both Model P (42.0 cycles) and Model I (43.1 cycles).

Discussion. Both Model P and Model I produced an LLPC effect. In addition, the size of this interaction was the same for both versions of the model. Thus, the LLPC effect can be accounted for by an ISPC effect.

General Discussion

The current investigation presented a computational account of the ISPC effect in the context of the conflict-monitoring hypothesis. The control mechanism hypothesized to account for the ISPC effect was also demonstrated to be able to account for the LLPC effect. In Botvinick et al. (2001), this latter effect was accounted

Table 2
Trial Structure for Simulation 1

	Word			
Color	*yellow*	*green*	*blue*	*red*
YELLOW	27	9		
GREEN	9	27		
BLUE			9	27
RED			27	9

Table 3
Cycles for Models P and I as a Function of Congruency and Congruent Proportion Manipulated at the Item Level

Model type	Incongruent	Congruent	Stroop effect
Model P			
25% congruent	326.6	163.1	163.5
75% congruent	327.4	163.2	164.2
Model I			
25% congruent	323.8	162.4	161.3
75% congruent	358.4	166.8	191.5

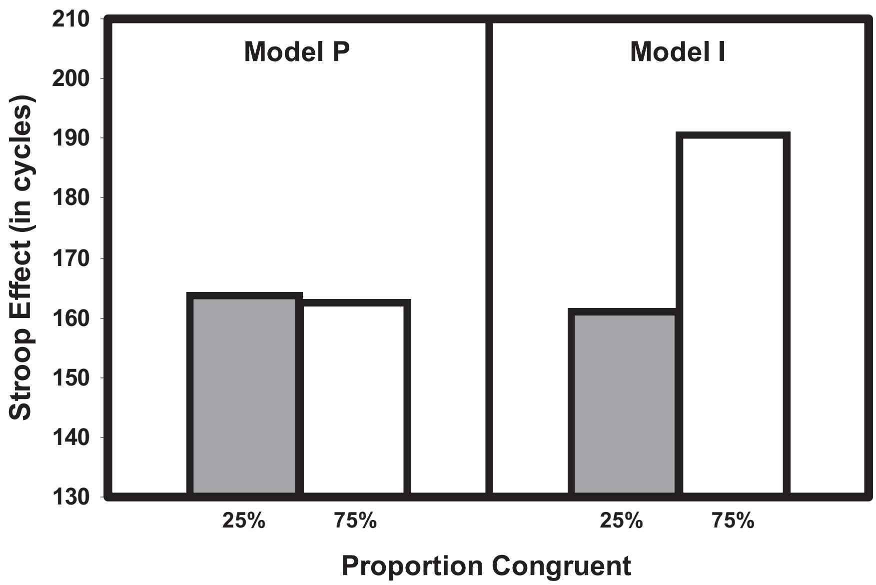

Figure 4. The interaction between the proportion of congruent trials and congruency when the proportion of congruent items is manipulated at the item level as a function of the level at which control operates. When control operates equally across the entire pathway as in Model P (left panel), this interaction is absent. When control operates at the level of the individual items as in Model I (right panel), this interaction is present.

for via control that operated at the level of the entire pathway. Critically, we showed that this locus of control, although sufficient to account for the LLPC effect, is insufficient to account for the ISPC effect. However, the presently proposed item-level control is both necessary to account for the ISPC effect and sufficient to account for the list-level effect.

In the following sections, we discuss alternative accounts of the ISPC effect, explore the implications of recent conflict-adaptation findings for conflict-monitoring models, present predictions derived from our account of the ISPC effect, and offer some future directions for the conflict-monitoring framework.

Alternative Accounts of the ISPC Effect

An alternative explanation of the ISPC effect attributes it to stimulus-specific priming. Specifically, manipulations of ISPC effects involve increasing the frequency of some stimuli and decreasing the frequency of others and this difference in frequency could produce an ISPC effect (see Table 2). For example, in the mostly congruent condition, the congruent stimuli (i.e., yellow–YELLOW and green–GREEN) occurred more often than the incongruent stimuli (i.e., yellow–GREEN and green–YELLOW). In the mostly incongruent condition, the incongruent stimuli (i.e., red–BLUE and blue–RED) occurred more often than the congruent stimuli (i.e., red–RED and blue–BLUE). If increases in frequency of occurrence speeds processing, for example through a priming mechanism, then the Stroop effect should be larger in the mostly congruent condition than in the mostly incongruent condition because congruent trials were more frequent than incongruent trials in the former condition, and incongruent trials were more frequent than congruent trials in the latter condition. Critically, this account would require only that stimulus processing be sensitive to relative frequency. A control mechanism, such as the one proposed here, would not be required. It is important to note that the same stimulus-specific priming mechanism could also account for the global proportion effect.

Table 4
Trial Structure for Simulation 2

Proportion congruent	Color	Word			
		yellow	*green*	*blue*	*red*
25%	YELLOW	9	9	9	9
	GREEN	9	9	9	9
	BLUE	9	9	9	9
	RED	9	9	9	9
75%	YELLOW	27	3	3	3
	GREEN	3	27	3	3
	BLUE	3	3	27	3
	RED	3	3	3	27

Table 5
Cycles for Models P and I as a Function of Congruency and Congruent Proportion Manipulated at the List Level

Model type	Incongruent	Congruent	Stroop effect
Model P			
25% congruent	305.5	161.6	143.9
75% congruent	350.9	165.0	185.9
Model I			
25% congruent	321.1	162.9	158.2
75% congruent	367.5	166.2	201.3

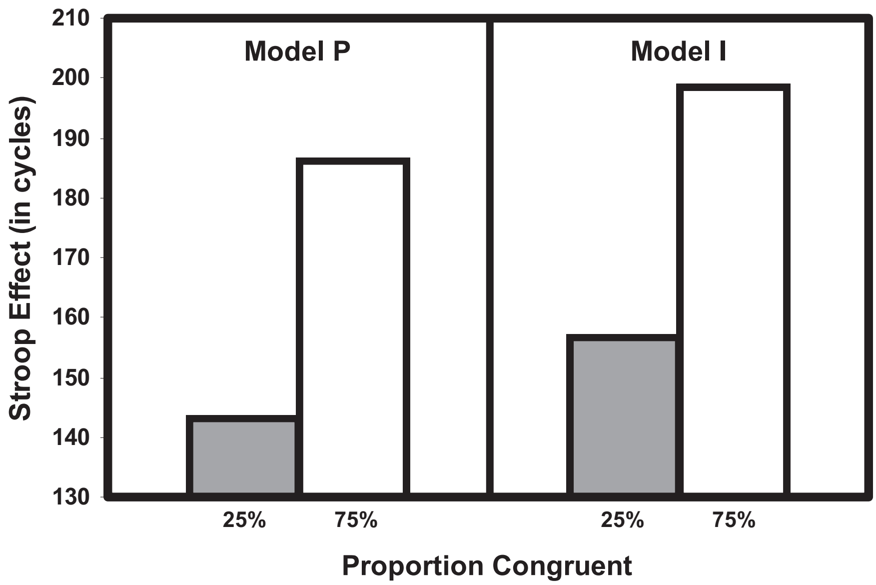

Figure 5. The interaction between the proportion of congruent trials and congruency when the proportion of congruent items is manipulated at the list level as a function of the level at which control operates. Note that the size of the interaction is nearly identical in Model P (left panel) as compared to Model I (right panel).

Evidence against a priming account of both the ISPC effect and the global proportion-congruent effect has been provided by numerous researchers (Jacoby et al., 1999, 2003; Lindsay & Jacoby, 1994; Schmidt, Crump, Cheesman & Besner, 2007). With respect to the ISPC effect, Jacoby et al. (1999, 2003) used the process-dissociation procedure and demonstrated that the ISPC manipulation affects the word process and not the color process. If stimulus-specific priming had been responsible for the ISPC effect, the process dissociation procedure should have revealed influences on both the word and color processes. Thus, it appears that the ISPC effect cannot be entirely accounted for by a priming mechanism. Lindsay and Jacoby (1994), also using the process-dissociation procedure, provided a similar demonstration using a global manipulation of proportion congruence.

Jacoby et al.'s (1999, 2003) results support the need for some form of control mechanism to account for the ISPC effect but suggest an alternative means through which this control is exerted in the context of the Stroop task. Using the process-dissociation procedure, Jacoby et al. (1999, 2003) demonstrated that the ISPC manipulations affect the influence of *word* processes, not the influence of *color* processes, and they demonstrated the viability of this account in the context of the Stroop Counter model (i.e., by modulating the influence of the word process as a function of proportion congruent). In Botvinick et al.'s (2001) model and the present model, the task demand parameters along the color pathway are modulated to account for proportion-congruent effects. Thus, the two accounts differ with respect to what is being controlled.

At this point it is useful to draw attention to the distinction between the conflict-monitoring hypothesis (i.e., the detection of conflict in the system signals the need for control) and the mechanism responsible for control. Specifically, the conflict-monitoring hypothesis addresses the question of how the need for control is determined (i.e., conflict detection); how control is implemented following the detection of conflict is a separate and independent issue. The models differ only with respect to the latter issue. Jacoby et al.'s (1999) model does not specify how the need for control is determined. We have demonstrated here that a conflict-monitoring module in concert with a control mechanism similar to Botvinick et al.'s (2001) can be used to model the ISPC effect. However, we see no reason why a conflict-monitoring mechanism could not be added to a model that implements control in a different way (such as Jacoby et al.'s, 1999, Stroop Counter model). The conflict-monitoring hypothesis is not tied to a specific model (e.g., Cohen et al.'s, 1990, Stroop model), and its application to other existing models (e.g., the Stroop Counter model) represents an interesting direction for future computational work.

Conflict-Adaptation Effects

Botvinick et al.'s (2001) conflict-monitoring model also provides an account of conflict-adaptation effects (Gratton, Coles, & Donchin, 1992). The conflict-adaptation effect in the Stroop task refers to the observation that the size of the Stroop effect is smaller following an incongruent trial than following a congruent trial (Kerns et al., 2004). According to the conflict-monitoring theory, the conflict-adaptation effect is due to variation in the amount of top-down control on a trial-by-trial basis. Following the detection of conflict on an incongruent trial, the amount of top-down control will be higher, thus reducing the word's influence on the subsequent trial. This would serve to reduce the magnitude of the Stroop effect following incongruent trials.

Recent work has suggested that the conflict-adaptation effect may be modulated by stimulus repetition (see Mayr, Awh, & Laurey, 2003; Nieuwenhuis et al., 2006; Ullsperger, Bylsma, & Botvinick, 2005). The exact relation between stimulus repetition and conflict adaptation is currently unclear, however, it appears that although the conflict-adaptation effect may be reduced after removal of stimulus repetitions, it is not eliminated. This pattern of data is difficult to account for in terms of a conflict-monitoring device. With a pathway-level mechanism, an individual item produces conflict on trial $n - 1$, and attention is increased to all items equally, which would reduce the word's impact on trial n for all items. Thus, the conflict-adaptation effect would be unaffected by the inclusion/exclusion of stimulus repetitions. With an item-level mechanism, the item that produces the conflict on trial $n - 1$ has attention to it increased. Thus, the word's impact on the following trial would only be reduced when the stimulus repeats.[3]

Thus, both Model P and Model I have difficulty explaining the interaction between stimulus repetition and conflict adaptation. However, the results produced by Model P and Model I hint at the possibility that a hybrid model that combines an item-level control mechanism with a pathway-level control mechanism might provide a better account of the conflict-adaptation effect. In this hybrid model, the conflict-adaptation effect should be larger when stimulus repetitions are included rather than excluded because in the latter case the item-specific contribution would be removed and only the item-independent contribution would remain.[4] Such a model is discussed in the *Future Directions* section.

New Predictions

The utility of a theory and by extension any computational model instantiating that theory is found not just in what it explains but in what new predictions it makes. Critically, the present account of the ISPC effect in terms of conflict monitoring makes a strong prediction that has yet to be tested.

If the ISPC effect is to be explained in the context of the conflict-monitoring hypothesis, as we have proposed, then the ACC, claimed to house the conflict monitor, should be sensitive to conflict at the level of individual items. Specifically, if an ISPC manipulation were employed, ACC activation on incongruent trials should be higher for mostly congruent items than for mostly incongruent items. Critically, Botvinick et al.'s (2001) model predicts no difference between mostly congruent and mostly incongruent items in terms of ACC activation on incongruent trials in the context of an ISPC manipulation. In the only similar study that we are aware of, Carter et al. (2000) demonstrated that, in the context of an LLPC manipulation, ACC activation was highest on incongruent trials in a mostly congruent list. This result, as we demonstrated in the above simulations, is explicable in terms of either item-level control or pathway-level control. It is only with an ISPC manipulation, wherein individual items differ in proportion congruence but the overall list is 50% congruent, that these two theories can be differentiated. This experiment has yet to be conducted, but it is clear that if the ISPC effect is to be explained in the context of the conflict-monitoring hypothesis, the ACC itself would need to be sensitive to conflict at the item level.

A second, more specific, prediction is that if the ISPC effect is the result of ACC modulation, then persons with impaired (e.g., schizophrenics or stroke patients with frontal lobe lesions) or otherwise dysfunctional (e.g., children with attention-deficit/hyperactivity disorder, see van Meel, Heslenfeld, Oosterlaan, & Sergeant, 2007) ACC function should show a decreased, or no, ISPC effect. Temporary lesions to the ACC made via rapid transcranial magnetic stimulation (see also Gobell, Rushworth, & Walsh, 2006) may also reduce the ISPC effect.

Future Directions—A Challenge for Formal Models of Conflict Monitoring?

In Botvinick et al. (2001), a number of different control mechanisms, all subserved by a conflict-detection device, were introduced to account for various control phenomena.[5] To account for conflict-adaptation effects in the flanker task, they had a conflict monitor signal a unit responsible for the focusing of spatial attention. To account for the error-related slowdown in a two-choice speeded response task, they had a conflict monitor signal a response priming unit that served to shift the system to a new point on the speed-accuracy tradeoff curve. Finally, as discussed here, to account for the LLPC effect in Stroop, they had a conflict monitor signal a task demand unit that increased/decreased selective attention to the color dimension. Here, we have introduced an additional control mechanism (i.e., item-level control) that is able to account for the ISPC effect. In addition, this item-level control mechanism is able to account for the LLPC effect. Thus, the item-specific mechanism is the only mechanism that appears able to account for more than a single phenomenon. Given all of these different control mechanisms that have been postulated, an important future direction for the conflict-monitoring approach is to integrate all of these mechanisms within a single model that produces all relevant phenomena, preferably with a single parameter set. A rough sketch of such a model is presented next.

As briefly noted in the section on conflict-adaptation effects, a hybrid Stroop model that combines an item-level mechanism with a pathway-level mechanism would produce the ISPC effect, the LLPC effect, and a conflict-adaptation effect that is reduced but still present after removal of stimulus repetitions. To account for the error-related slowdown, this model would require a response priming mechanism. Thus, the model would have three separate loci for control: (a) item-level control, (b) pathway-level control, and (c) response priming, all of which would be subserved by a conflict-detection module. In addition to addressing behavioral phenomena, the described model would also inherit the ability to account for data from a large number of brain activation (e.g., Botvinick et al., 2001; Kerns et al., 2004) and electrophysiological (e.g., Yeung et al., 2004; Yeung & Cohen, 2006) studies. The integration of these mechanisms is a computationally daunting task but would mark an important step forward in the formal development of the conflict-monitoring account.

[3] Simulations with Model P and Model I support this claim. Both produce conflict-adaptation effects. In Model P, this effect is unaffected by the inclusion/exclusion of stimulus repetitions, whereas in Model I, the conflict-adaptation effect is eliminated following removal of stimulus repetitions.

[4] In a preliminary set of simulations, such a model produced an ISPC effect, an LLPC effect, and a conflict-adaptation effect that was reduced but not eliminated following the removal of stimulus repetitions.

[5] Botvinick et al. (2001) used a different parameter set for each different control phenomenon. Here, the same parameter set was used to simulate both the ISPC effect and the LLPC effect.

Conclusions

The present investigation introduced the idea of item-level control in the context of the conflict-monitoring hypothesis. We demonstrated that an existing model of the Stroop task with the addition of a conflict-monitoring mechanism based on this principle could account for both the ISPC effect and the LLPC effect with the same parameter set.

References

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.

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences, 8,* 539–546.

Brown, T. L., Gore, C. L., & Carr, T. H. (2002). Visual attention and word recognition in Stroop color naming: Is word recognition "automatic"? *Journal of Experimental Psychology: General, 131,* 220–240.

Carter, C. S., MacDonald, A. M., Ross, L. L., Stenger, V. A., Noll, D., & Cohen, J. D. (2000). Parsing executive processes: Strategic versus evaluative functions of the anterior cingulate cortex. *Proceedings of the National Academy of Sciences, USA, 97,* 1944–1948.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology, 40,* 343–367.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97,* 332–361.

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

Dishon-Berkovits, M., & Algom, D. (2000). The Stroop effect: It is not the robust phenomenon that you have thought it to be. *Memory & Cognition, 28,* 1437–1449.

Gobell, S. M., Rushworth, M. F., & Walsh, V. (2006). Inferior parietal rTMS affects performance in an addition task. *Cortex, 42,* 774–781.

Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation and responses. *Journal of Experimental Psychology: General, 121,* 480–506.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, USA, 79,* 2554–2558.

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

Jacoby, L. L., McElree, B., & Trainham, T. N. (1999). Automatic influences as accessibility bias in memory and Stroop tasks: Toward a formal model. In D. Gopher & A. Koriat (Eds.), *Attention and performance XVII: Cognitive regulation of performance. Interaction of theory and application* (pp. 461–486). Cambridge, MA: The MIT Press.

Kanne, S. M., Balota, D. A., Spieler, D. H., & Faust, M. E. (1998). Explorations of Cohen, Dunbar, and McClelland's (1990) connectionist model of Stroop performance. *Psychological Review, 105,* 174–187.

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

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,* 219–234.

Link, S. W. (1975). The relative judgment theory of two-choice response time. *Journal of Mathematical Psychology, 12,* 114–135.

Logan, G. D., & Zbrodoff, N. J. (1979). When it helps to be misled: Facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. *Memory & Cognition, 7,* 166–174.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109,* 163–203.

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

Nieuwenhuis, S. T., Stins, J. F., Posthuma, D., Polderman, J. C., Boomsma, D. I., & de Geus, E. J. (2006). Accounting for sequential trial effects in the Flanker task: Conflict adaptation or associative priming? *Memory & Cognition, 34,* 1260–1272.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 45–76). Cambridge, MA: MIT Press.

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

Trainham, T. N., Lindsay, D. S., & Jacoby, L. L. (1997). Stroop process dissociations: Reply to Hillstrom and Logan (1997). *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1579–1587.

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, & Behavioral Neuroscience, 5,* 467–472.

van Meel, C. S., Heslenfeld, D. J., Oosterlaan, J., & Sergeant, J. A. (2007). Adaptive control deficits in attention-deficit/hyperactivity disorder (ADHD): The role of error processing. *Psychiatry Research, 151,* 211–220.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review, 111,* 931–959.

Yeung, N., & Cohen, J. D. (2006). The impact of cognitive deficits on conflict monitoring: Predictable dissociations between the ERN and N2. *Psychological Science, 17,* 164–171.

(*Appendix follows*)

Appendix

Model Details

Below is a brief description of how activation in the model operates; more details are available in Cohen et al. (1990). Within the network, units update their activations by taking a weighted sum of the input they receive from other levels in the network. As the present network is cascaded (rather than based on interactive activation), input only arises from units in the preceding level. Mathematically, the net input at time t for unit_j (at level_n) is given by:

$$\text{net}_j(t) = \sum_i a_i(t)w_{ij}, \quad \text{(A1)}$$

where $a_i(t)$ is the activation of each unit_i (at level_{n-1}) from which unit_j receives input, and w_{ij} is the strength of the connection weight between unit_i at level_n and unit_j at level_{n-1}. The activation of a unit is a weighted average of its current net input and its previous net input:

$$a_j(t) = \overline{\text{net}}_j(t) = \tau\text{net}_j(t) - (\tau - 1)\overline{\text{net}}_j(t-1), \quad \text{(A2)}$$

where $\overline{\text{net}}_j(t)$ is the average of the net input to unit_j over time, $\text{net}_j(t)$ is the net input to unit_j at time t, and τ is the cascade rate ($0 < \tau \leq 1$). In a sense, τ can be thought of as a resistance parameter. If one thinks of activation in a unit as resistant to change, then τ provides the strength required to overcome the resistance. As τ increases, its strength increases: If τ were set to 0, there would be no change at all; if it were set to 1, change would be very rapid (indeed, it would completely overcome the resistance). The value for τ was .005 in all simulations.

One problem with having a network cascade in this fashion is that Equation A2 provides a linear activation function. Such networks suffer from a number of concerns related to computational limitations (see Rumelhart, Hinton, & McClelland, 1986, for discussion). To overcome these limitations, researchers introduce nonlinearity into the activation function by using a sigmoid function to calculate the activation of a unit based on its instantaneous net input:

$$a_j(t) = \text{logistic}[\text{net}_j(t)] = \frac{1}{1 + e^{-\text{net}_j(t)}}, \quad \text{(A3)}$$

where $\text{net}_j(t)$ is given by Equation A1. The logistic function constrains values between 0 and 1. This nonlinearity provides all the important behaviors of cascaded networks without suffering from an unlimited buildup of activation over time. Furthermore, the dynamic properties of the cascade model can be introduced by incorporating the concept of τ and using $\overline{\text{net}}_j(t)$ as it is defined in Equation A2 as the input to the logistic function. This gives us the following activation rule:

$$a_j(t) = \text{logistic}[\overline{\text{net}}_j(t)]. \quad \text{(A4)}$$

With this function, activation builds up slowly over time (as controlled by the cascade rate, τ) and is constrained to a value between 0 and 1. All that remains is to specify the mechanism for response selection.

Response Selection

The model uses principles from a random walk (Link, 1975) and a diffusion process (Ratcliff, 1978) to select its ultimate response. Each potential response is paired with an evidence accumulator that takes its input from the output units of the network. At the beginning of each trial, the evidence accumulators are set to 0, and at each time-step of processing (a cycle), evidence accumulates as a function of the activation in the relevant output unit. The amount of evidence accumulated for response i is given by the following equation:

$$\text{evidence}_i = \text{N}(\alpha[\text{act}_i - \max(\text{act}_{j \neq i})], \sigma). \quad \text{(A5)}$$

$\text{N}(\mu, \sigma)$ is a random value sampled from a normal (Gaussian) distribution with mean μ and standard deviation, σ; α determines the rate of evidence accumulation; act_i is the activation in output unit i; $\max(\text{act}_{j \neq i})$ is the maximum activation of the other output units; and σ is a noise parameter. Taking the difference between activation in the output unit of interest, and the other output unit with the strongest activation allows evidence in the response units to differentiate more quickly as the activation in the output units differentiates between potential responses. A response is generated when one of the accumulators reaches a fixed threshold. For all of the present simulations, the value of α was .08, σ was 0.015, and the response threshold was 1.0.

Received September 27, 2006
Revision received April 16, 2007
Accepted April 17, 2007 ■