

# A Mechanism for Error Detection in Speeded Response Time Tasks

Clay B. Holroyd and Nick Yeung  
Princeton University

Michael G. H. Coles  
University of Illinois at Urbana–Champaign and F. C. Donders  
Centre for Cognitive Neuroimaging

Jonathan D. Cohen  
Princeton University and University of Pittsburgh

The concept of error detection plays a central role in theories of executive control. In this article, the authors present a mechanism that can rapidly detect errors in speeded response time tasks. This error monitor assigns values to the output of cognitive processes involved in stimulus categorization and response generation and detects errors by identifying states of the system associated with negative value. The mechanism is formalized in a computational model based on a recent theoretical framework for understanding error processing in humans (C. B. Holroyd & M. G. H. Coles, 2002). The model is used to simulate behavioral and event-related brain potential data in a speeded response time task, and the results of the simulation are compared with empirical data.

Frontal parts of the brain, including the prefrontal cortex (Luria, 1973; Stuss & Knight, 2002), the anterior cingulate cortex (Devinsky, Morrell, & Vogt, 1995; Posner & DiGirolamo, 1998), and their connections with the basal ganglia (L. L. Brown, Schneider, & Lidsky, 1997; Cummings, 1993), are thought to compose an executive system for cognitive control. The functions of this system are thought to include setting high-level goals, directing other cognitive systems to execute behaviors in accordance with those goals, monitoring the progress of these systems as they carry out their tasks, and intervening when they fail (Logan & Gordon, 2001; E. K. Miller & Cohen, 2001). Of particular concern to the executive system are behavioral errors, because, by definition, these events indicate when the system falls short of a goal. Thus, theories of executive control assume the existence of an error detection system capable of identifying failures in performance. Once an error is detected, the executive system can act to improve

performance on the task, both in the short term (by initiating remedial actions in conjunction with the error) and in the long term (by updating the response production system such that the errors are not repeated; Ohlsson, 1996; Schall, Stuphorn, & Brown, 2002).

In this article, we present a biologically plausible mechanism for the detection of speeded response errors. The mechanism is based on a recent theory that holds that the system operates according to principles of reinforcement learning (Holroyd & Coles, 2002). In a previous report, we showed that this mechanism can learn from feedback in a trial-and-error learning task. We further showed that the mechanism can, in principle, detect errors in speeded response time (RT) tasks by identifying particular combinations of internally generated responses and externally presented stimuli that are associated with negative outcomes. An attractive feature of this account is that it makes use of information about stimuli and responses that could be readily available to the monitoring system. In this way, the model represents one of the first computationally plausible mechanisms for online detection of errors. However, in our previous work, the details of the speeded RT model were oversimplified: The previous model did not simulate behavior and thus did not simulate the underlying cognitive processes on which the error detection mechanism depends. This raises the possibility that the error detection mechanism might fail were it implemented in a more realistic system.

In the present study, we address this issue by first building a neurally and cognitively plausible model of perceptual categorization and response selection that simulates human behavior in a speeded RT task. We then demonstrate that the error monitor can rapidly detect errors produced by the task module and illustrate how the error signal produced by the error monitor can be used by the task module for control. In so doing, we further show that the model can account for electrophysiological evidence of a neural

---

Clay B. Holroyd and Nick Yeung, Department of Psychology, Princeton University; Michael G. H. Coles, Department of Psychology, University of Illinois at Urbana–Champaign, and F. C. Donders Centre for Cognitive Neuroimaging, Nijmegen, the Netherlands; Jonathan D. Cohen, Department of Psychology and Center for the Study of Brain, Mind and Behavior, Princeton University, and Department of Psychiatry, University of Pittsburgh.

This research was supported in part by National Institute of Mental Health Grants MH41445 and MH62196, National Institute of Mental Health Predoctoral Fellowship MH11530, and National Institute of Mental Health Postdoctoral Fellowship MH63550. We are grateful to Emanuel Donchin and Sander Nieuwenhuis for their commentaries on various versions of this article, to Gordon Logan for helpful discussions, and to Kevin Spencer for kindly providing his computer code for simulating the Eriksen flanker task.

Correspondence concerning this article should be addressed to Clay B. Holroyd, who is now at the Department of Psychology, University of Victoria, PO Box 3050 STN CSC, Victoria, British Columbia V8W 3P5, Canada. E-mail: holroyd@uvic.ca

system for error detection.<sup>1</sup> When taken together with previous modeling work (Holroyd & Coles, 2002), these simulations demonstrate, first, that the monitor can evaluate performance on the basis of both internal information related to the response and external information related to feedback and, second, that the system can use those evaluations to internalize representations of appropriate behavior and to improve performance on the task at hand.

### The Reinforcement Learning Theory

The error monitor detects errors by identifying states of the system and the environment associated with negative valence. In speeded response tasks, these states consist of combinations of internally generated responses and externally presented stimuli. Figure 1 shows a schematic of the model, which consists of two components: a task module (called an *actor* or *motor controller* in the language of reinforcement learning), and a monitoring module (called a *critic*; for reviews of the literature on reinforcement learning, see Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). The task module produces overt behaviors in response to external input. The monitor evaluates the output of the task module in the given task context, reinforcing the task module for good performance and punishing the task module for bad performance. More specifically, the monitor receives stimulus-related information associated with external events and response-related information associated with the behaviors produced by the task module. From this information, the monitor assigns a degree of *value* (goodness or badness) to ongoing events. In addition, the monitor detects instantaneous changes in the value, called *temporal difference* (TD) *errors* (Sutton, 1988). These TD signals are sent to the task module, where they reinforce task-appropriate behaviors and extinguish inappropriate behaviors. The TD signals are also used by the monitor itself to improve its estimates of the value.<sup>2</sup>

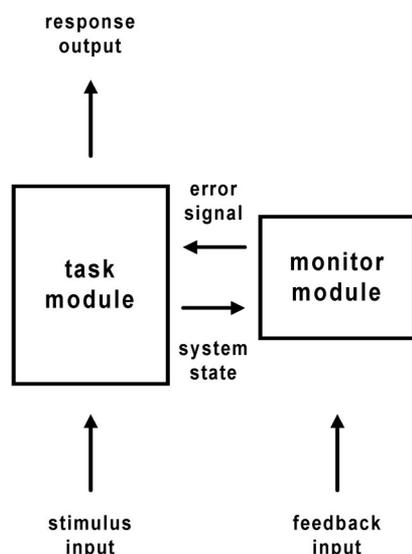


Figure 1. An architecture for solving reinforcement learning problems. See text for details.

As specified by the algorithm that computes them, these values are predictions: Positive values indicate that future events will be favorable to the organism, whereas negative values indicate that future events will be unfavorable to the organism. Thus, according to this definition, TD signals indicate changes in prediction: Positive TD signals (associated with positive changes in value) indicate that ongoing events are better than expected, and negative TD signals (associated with negative changes in value) indicate that ongoing events are worse than expected. These TD signals are used by the monitoring system to improve its predictions of future reward; if the TD signal is zero, then the monitoring system has learned to predict the value of ongoing events perfectly, and no more learning need occur. The method is in fact a generalization of the Rescorla–Wagner learning rule to the continuous time domain (Rescorla & Wagner, 1972; Sutton & Barto, 1990). For our purposes, the crucial point is that a negative TD signal indicates that an unfavorable event has occurred that the monitor did not foresee and that adjustments in behavior and estimates of value are therefore required.

The architecture of our model was motivated by two separate literatures on the neural basis of error detection and learning. First, several decades of behavioral research have provided indirect evidence for the existence of an error detection system (e.g., Angel, 1976; Diggle, 1987; Higgins & Angel, 1970; Laming, 1979; Rabbitt, 1966a, 1966b). More recently, the discovery of a neural signal associated with error commission has provided more direct evidence of its existence. When participants make errors in choice RT tasks (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1990; Gehring, Goss, Coles, Meyer, & Donchin, 1993) or receive feedback indicating that they made an incorrect response (Miltner, Braun, & Coles, 1997), a negative-going deflection occurs in the event-related brain potential (ERP) called the error-related negativity (ERN; for a review of ERPs, see Coles & Rugg, 1995). Both the response ERN (e.g., Dehaene, Posner, & Tucker, 1994; Holroyd, Dien, & Coles, 1998) and the feedback ERN (e.g., Gehring & Willoughby, 2002; Miltner et al., 1997) are distributed over a frontal–central region of the scalp, appear to be generated in the anterior cingulate cortex, and are associated with an error detection mechanism (for reviews of the ERN, see Falkenstein, Hoormann, Christ, & Hohnsbein, 2000; Holroyd, Nieuwenhuis, Mars, & Coles, 2004; Nieuwenhuis, Holroyd, Mol, & Coles, 2004).

The second source inspiring the architecture of our model was a large collection of computational studies that assumed that the monitor is located in the basal ganglia, that the output of the monitor is carried by the mesencephalic dopamine system, and that this output consists of TD signals indicating when events are better

<sup>1</sup> The conflict monitoring theory (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Yeung, Botvinick, & Cohen, 2004) and the mismatch theory (Coles, Scheffers, & Holroyd, 2001) provide alternative accounts of this mechanism. Although these theories and the reinforcement learning theory promise to inform one another in important ways, we defer a discussion about their relation until the end of this article. Note that the ERN data presented here have also been simulated by Yeung et al. (2004) according to the principle of response conflict.

<sup>2</sup> We use *TD signal* instead of *TD error* to avoid confusion with the concept of a performance error.

or worse than expected (e.g., Barto, 1995; Braver et al., 2001; Contreras-Vidal & Schultz, 1999; Daw, Kakade, & Dayan, 2002; Daw & Touretzky, 2002; Doya, 1999, 2002; Egelman, Person, & Montague, 1998; Houk, Adams, & Barto, 1995; McClure, Daw, & Montague, 2003; Montague, Dayan, & Sejnowski, 1996; Nakahara, Doya, & Hikosaka, 2001; Rougier & O'Reilly, 2002; Schultz, Dayan, & Montague, 1997; Sporns & Alexander, 2002; Suri, 2002; Suri & Schultz, 2001; cf. J. Brown, Bullock, & Grossberg, 1999). Expanding on this work, we illustrated how the monitor in the basal ganglia can detect behavioral errors, producing dopaminergic error signals that indicate that events are worse than expected. We argued that the ERN is produced when the error signals carried by the mesencephalic dopamine system disinhibit motor neurons in the anterior cingulate cortex (see also Holroyd, 2004). Thus, when the monitor determines that an unexpected error response has occurred, a response ERN is produced, and when the monitor determines that an unexpected error feedback stimulus has occurred, a feedback ERN is produced. Furthermore, the theory holds that motor areas in the anterior cingulate cortex use the error information to improve performance on the task at hand.

This hypothesis is consistent with a range of biological evidence that suggests that the midbrain dopamine system plays a critical role in reinforcement learning (e.g., Bao, Chan, & Merzenich, 2001; Gurden, Takita, & Jay, 2000; Kimura & Matsumoto, 1997; Reynolds, Hyland, & Wickens, 2001; Reynolds & Wickens, 2002) by indicating errors in reward prediction (Schultz, 2002; Waelti, Dickinson, & Schultz, 2001). This dopamine activity modulates processes in the frontal cortex (Dreher & Burnod, 2002; Dreher, Guigon, & Burnod, 2002; Ferron, Thierry, Le Douarin, & Glowinski, 1984; Gao, Krimer, & Goldman-Rakic, 2001; Lewis & O'Donnell, 2000; Yang & Seamans, 1996), including the anterior cingulate cortex (e.g., Crino, Morrison, & Hof, 1993; Vogt, Vogt, Nimchinsky, & Hof, 1997), where the dopamine system appears to convey evaluative information (e.g., Porrino, 1993; Richardson & Gratton, 1998). Furthermore, the caudal part of the anterior cingulate cortex, which is thought to contribute to high-level cognitive control of motor behavior (reviewed in Devinsky et al., 1995; Dum & Strick, 1993; Paus, 2001; Picard & Strick, 1996), is sensitive to reward (e.g., Shidara & Richmond, 2002; Tzschentke, 2000) and error information (e.g., Gamba, Sasaki, & Brooks, 1986; Ito, Stuphorn, Brown, & Schall, 2003; Niki & Watanabe, 1979; Ullsperger & Von Cramon, 2003; see also Stuphorn, Taylor, & Schall, 2000) and appears to use this evaluative information to guide action selection (e.g., Bush et al., 2002; Elliott & Dolan, 1998; Gabriel, 1993; Procyk, Tanaka, & Joseph, 2000; Shima & Tanji, 1998). It is important to note that the anterior cingulate cortex is also believed to be the source of the ERN (e.g., Dehaene et al., 1994; Holroyd et al., 1998; Miltner et al., 1997; reviewed in Holroyd, Nieuwenhuis, Mars, & Coles, 2004).

Recently, several studies of the feedback ERN have provided empirical support for the theory (for reviews, see Holroyd, Nieuwenhuis, Mars, & Coles, 2004; Nieuwenhuis, Holroyd, Mol, & Coles, 2004). These experiments all involved pseudo-trial-and-error learning tasks in which, on each trial, participants were presented with an imperative stimulus, were required to make a response, and were then presented with a feedback stimulus indicating a reward or penalty. In these experiments, the participants

were not informed of the appropriate stimulus-response mappings and were required to infer them by trial and error. Mostly focusing on the feedback ERN, these studies have provided evidence suggesting that the ERN indeed reflects a prediction error in reward and have explored how the evaluative system that produces the ERN determines whether an event is good or bad (Holroyd & Coles, 2002; Holroyd, Larsen, & Cohen, 2004; Holroyd, Nieuwenhuis, Yeung, & Cohen, 2003; Holroyd, Nieuwenhuis, Yeung, et al., 2004; Mars, de Bruijn, Hulstijn, Miltner, & Coles, 2004; Nieuwenhuis et al., 2002; Nieuwenhuis, Yeung, Holroyd, Schurger, & Cohen, 2004; Yeung & Sanfey, 2004).

Other recent results have also provided some preliminary support for the position that the ERN depends on the mesencephalic dopamine system. For example, ERN amplitude is increased by administration of d-amphetamine, which releases dopamine and inhibits its reuptake; this suggests that dopamine is involved in ERN production (De Bruijn, Hulstijn, Verkes, Ruigt, & Sabbe, 2004). Conversely, alcohol consumption reduces ERN amplitude (Ridderinkhof et al., 2002), possibly because alcohol affects dopamine receptors (Holroyd & Yeung, 2003). Although Parkinson's disease damages the mesencephalic dopamine system, evidence of abnormal ERNs in people with mild to moderate Parkinson's disease has been mixed (Falkenstein et al., 2001; Holroyd, Praamstra, Plat, & Coles, 2002). However, ERN amplitude decreases with age, an observation that has been attributed to age-related changes in dopamine function (Nieuwenhuis et al., 2002). The ERN is also abnormal in schizophrenia (Bates, Kiehl, Laurens, & Liddle, 2002; Kopp & Rist, 1999; Mathalon et al., 2002) and in Gilles de la Tourette syndrome (Johannes et al., 2002), both of which are associated with disruption of the midbrain dopamine system (Cohen & Servan-Schreiber, 1992; Davis, Kahn, Ko, & Davidson, 1991; Devinsky, 1983; Harrison, 2000; Singer, Butler, Tune, Seifert, & Coyle, 1982).

## Research Overview

As we have reviewed, much evidence is consistent with the reinforcement learning theory of the ERN (RL-ERN theory). In particular, the evidence suggests that the ERN reflects an error in reward prediction and that the error signal is associated with activity of the midbrain dopamine system. In the present study, we adopt the reinforcement learning framework to implement an error detection mechanism in a neurally plausible cognitive system. Our theory makes the critical claim that the system can accomplish online error detection by monitoring for relatively simple sources of information about stimulus categorization and response selection, such that it can detect errors as stimulus-response conjunctions associated with negative value. In this study, we investigate the plausibility of this theory as an account of human error detection and its reflection in the response ERN.

Our principal objective is to demonstrate that the monitor in the RL-ERN model can rapidly detect behavioral errors in a speeded RT task, on the basis of state information continuously provided to it from the task module. This question is of paramount concern given that the response ERN is typically observed in speeded response tasks, in which it reaches maximum amplitude within 100 ms following the error. This places tight constraints on the dynamics of any mechanism proposed to explain the ERN in terms of

error monitoring. Our second objective is to illustrate that the error signal produced by the monitor can be used to modify behaviors carried out by the task module. Specifically, we show that the signal can induce response slowing that is considered to be evidence of an error-sensitive control process (Laming, 1979; Rabbitt, 1966b). Further, the model predicts that the amount of slowing is proportional to the size of the error signal, a prediction that is confirmed by the empirical data. When taken together with previous modeling work (Holroyd & Coles, 2002), these simulations illustrate how the monitoring system can evaluate both internal and external sources of performance-related information and how the system can apply that information for behavioral control.

We begin by presenting the empirical data associated with a modified version (Holroyd & Coles, 2002; Holroyd, Praamstra, et al., 2002) of the Eriksen flanker task (B. A. Eriksen & Eriksen, 1974).<sup>3</sup> The Eriksen flanker task is commonly used to study the ERN (e.g., Gehring et al., 1993), so it provides an appropriate means to test the predictions of the model. Moreover, the dynamics of the cognitive processes that underlie performance of the task are well understood. Studies involving this task have demonstrated that the cognitive systems involved in stimulus categorization and response selection can operate simultaneously and that these systems continuously make available their output to other systems—both before the computations reach completion and afterward (C. W. Eriksen, Coles, Morris, & O'Hara, 1985; C. W. Eriksen & Schultz, 1979; for an opposing view, see J. Miller, 1988). Moreover, the timing of these neural processes is critical, as the order of the neural events contributes to the outcome (correct or incorrect) of each response (Coles, De Jong, Gehring, & Gratton, 1991; Coles, Gratton, Bashore, Eriksen, & Donchin, 1985; Coles, Scheffers, & Fournier, 1995; Gratton, Coles, Sirevaag, Eriksen, & Donchin, 1988). As we show, the empirical data in the modified version of this task provide further insight into the timing of these cognitive processes.

Our second step is to describe a neurally plausible and cognitively accurate module for this task based on principles that capture the continuous, dynamic, and parallel nature of the human information processing system. Accordingly, we implement the model with a parallel distributed processing approach, a neural network method that describes cognitive phenomena in terms of the flow of activity among pools of interconnected processing units (Rumelhart & McClelland, 1986). The units in such models have continuous and graded activation functions (McClelland & Rumelhart, 1988) that capture the continuous and parallel nature of the human information processing system (McClelland, 1992). A number of previous studies have already used these computational principles to simulate performance on this task (Botvinick et al., 2001; Cohen, Servan-Schreiber, & McClelland, 1992; McClelland, 1992; Servan-Schreiber, 1990; Servan-Schreiber, Bruno, Carter, & Cohen, 1998; Spencer & Coles, 1999; Yeung et al., 2004; see also Zhang, Zhang, & Kornblum, 1999), and the task module in our simulation reflects a straightforward extension of this research. We demonstrate the appropriateness of the task module by using it to account for the behavioral data (average RT, average accuracy, and RT distributions for correct and incorrect trials) in the task and for the latency of the P300, a component of the ERP sensitive to the duration of the stimulus evaluation process.

Our third step is to add a monitor module based on the RL-ERN theory that detects errors on the basis of information provided to it by the task module. It is important to note that these errors are detected in real time, that is, as the information to detect the error becomes available from the task module. We also show that the size and latency of the error signals produced by the monitor are consistent with empirical observations of the ERN. This step demonstrates the critical result of this study: that the error detection mechanism described by the RL-ERN theory can, in fact, rapidly detect errors and shows dynamic properties consistent with the empirical data concerning the ERN.

Finally, we show that the error signal produced by the monitor can be used to modulate an attentional bias directed at the response units: Error signals that indicate that an error has occurred reduce the amount of bias, slowing response production on upcoming trials. We further show that the slowing is greater on trials associated with large error signals than on trials associated with small error signals. The effect of the slowing is to increase the probability that subsequent trials will be correct. When taken together with previous modeling work (Holroyd & Coles, 2002), these simulations indicate how the system can use the error signals to improve performance, both to internalize representations of appropriate behavior, such that the system can operate in the absence of external feedback, and to adjust the state of the motor control system to improve performance on the task at hand.

## Empirical Data

### Method

#### Participants

Fifteen undergraduate students at the University of Illinois at Urbana-Champaign (9 male and 6 female) were paid \$5 per hour for participating in the experiment. The experiment consisted of two sessions conducted on different days; participants were paid a \$5 bonus for completing both sessions.

#### Task

Participants sat in front of a computer monitor in a dimly lit room and performed a modified version (cf. Holroyd, Praamstra, et al., 2002) of the Eriksen flanker task (B. A. Eriksen & Eriksen, 1974). Each session consisted of a practice block followed by 12 blocks of 200 trials each, with 5–10-min breaks between blocks. For each participant, 4,800 trials of data were collected. The stimulus-onset asynchrony was 1.5 s, and the duration of each stimulus was 50 ms. Stimuli consisted of four 5-letter stimulus arrays composed of *H*s and *S*s. The central letter of each array was designated the *target*, and the flanking distractor letters were either *compatible* (i.e., *HHHHH*, *SSSSS*) or *incompatible* (i.e., *SSHSS*, *HSSH*) with the target. Additionally, two of the stimuli with the same target (e.g., *HHHHH* and *SSHSS*) were *frequent*, each appearing on 40% of the trials. The remaining two stimuli were *infrequent*, each appearing on 10% of the trials. Together, target frequency and flanker compatibility defined four stimulus conditions: *infrequent compatible* (III), *infrequent incompatible*

<sup>3</sup> This presentation consists of a more detailed analysis of the empirical data reported in Holroyd and Coles (2002).

(FIF), frequent incompatible (IFI), and frequent compatible (FFF).<sup>4</sup> Participants sat about 1 m away from the computer display, such that each letter in the stimulus array subtended about 0.5° of visual angle. They were instructed to respond with the left hand to one target and with the right hand to the other target. The stimulus–response mappings and target probabilities were systematically varied across participants.

Participants responded on each trial by squeezing one of two zero-displacement dynamometers (Model 152A, Daytronic Linear Velocity Force Transducers, Dayton, OH) connected to an amplifier system (Model 830A, Conditioner Amplifiers, Dayton, OH). During the experiment, overt responses were registered when the participant's squeeze force exceeded 25% of his or her maximum squeeze force, which was determined for each participant at the start of the session. During a practice block of trials, participants received auditory feedback when their squeeze force exceeded this criterion, which enabled them to learn the amount of force necessary for a response to register.

Following each block, feedback informing the participant of his or her accuracy (percentage correct) and average speed (in milliseconds) was presented on the video display. For the purposes of the feedback, responses that were generated during the first 50 ms following stimulus onset were considered errors. Participants were asked to respond as quickly as possible while maintaining an accuracy of about 85%. Participants were told that if their accuracy fell below 80%, then on the following block they should respond more slowly to improve their performance. Conversely, they were told that if their accuracy rose to 90%, then on the following block they should exploit the opportunity to improve their speed. Verbal feedback was provided as well—for example, as encouragement to the participants to break their personal records in speed and performance.

Electrophysiological recording and data analysis methods are given in Appendix A. For the purpose of data analysis, RT was determined from the onset of the electromyogram (see Appendix A).

## Results and Discussion

### Overt Behavior

Figure 2 presents the accuracies (proportion correct, top) and the RTs (for correct and incorrect responses, bottom) associated with each condition. Two 2-factor analyses of variance (ANOVAs) with repeated measures on both factors (frequency, compatibility) were performed separately on the accuracy and correct RT data. Responses to frequent stimuli were faster,  $F(1, 14) = 223.8, p < .001, \eta_p^2 = .94$ , and more accurate,  $F(1, 14) = 512.0, p < .001, \eta_p^2 = .97$ , than responses to infrequent stimuli, which indicates that the participants developed a bias to respond with the hand that was mapped to the frequently appearing target. Furthermore, responses to incompatible arrays were slower,  $F(1, 14) = 149.3, p < .001, \eta_p^2 = .91$  (B. A. Eriksen & Eriksen, 1974; C. W. Eriksen et al., 1985; C. W. Eriksen & Schultz, 1979), and less accurate,  $F(1, 14) = 282.6, p < .001, \eta_p^2 = .95$  (Coles et al., 1985; Gratton et al., 1988), than responses to compatible arrays, which indicates an effect of the compatibility of the stimulus array on the response generation process. An interaction between frequency and compatibility revealed that the decrease in accuracy on incompatible trials compared with compatible trials was larger for infrequent trials than for frequent trials,  $F(1, 14) = 34.0, p < .001, \eta_p^2 = .71$ . In contrast, there was no interaction between frequency and compatibility with respect to RT ( $p = .50$ ).

We also examined the relationship between RT on correct and on incorrect trials. For both of the infrequent conditions (III, FIF), errors were faster than correct responses: III,  $t(14) = 16.3, p <$

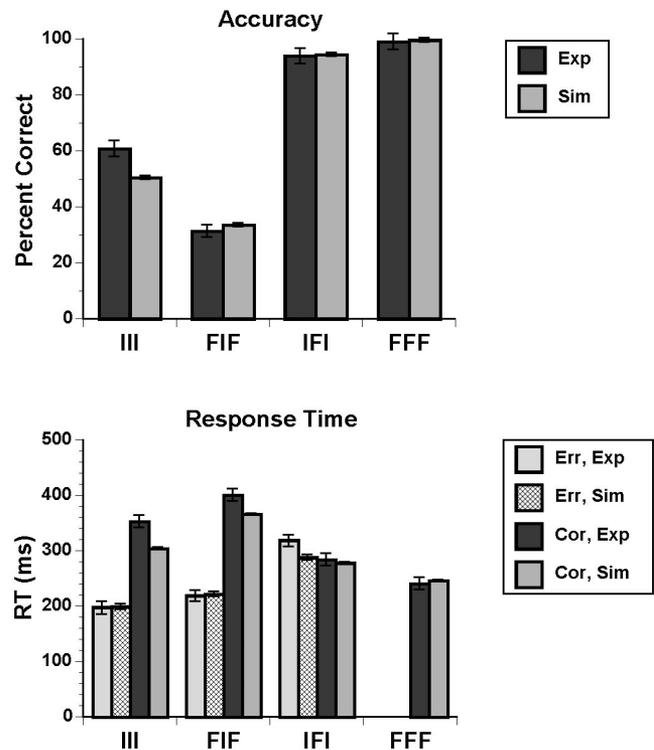


Figure 2. Accuracies and response times (RTs) in the biased Eriksen flanker task. The top panel shows average accuracies (percentage correct) for the empirical (Exp) and simulated (Sim) data. The bottom panel shows average RTs (in milliseconds) for the empirical and simulated data. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition; FFF = frequent compatible condition; Err = errors; Cor = correct responses.

.001; FIF,  $t(14) = 16.3, p < .001$ . This result confirms that the participants developed a bias to respond quickly with the hand that was mapped to the frequently appearing target. In contrast, for the IFI condition, errors were slower than correct responses,  $t(14) = 3.8, p < .005$  (cf. Holroyd, Praamstra, et al., 2002).

These results suggest that the response generation process was influenced by three primary factors: a bias associated with the frequency of appearance of the target stimuli, information pertaining to the flanking letters, and information pertaining to the target letter. Examination of the RT distributions (see Figure 3) suggests that these influences occurred in sequence, such that the data can be explained by three periods of information processing. (The left column of Figure 3 shows RT histograms for each stimulus, for correct and incorrect responses, pooled across participants; the middle column shows RT histograms for the same conditions for

<sup>4</sup> The format of the abbreviations is analogous to the flanker stimuli themselves. Thus, if one takes the center of the stimulus array to be the target and the flanking letters to be noise, then III and FFF are both compatible stimuli and FIF and IFI are both incompatible stimuli. Furthermore, if the central target letter specifies the frequency of appearance of that stimulus, then III and FIF are both infrequently occurring stimuli and IFI and FFF are both frequently occurring stimuli.

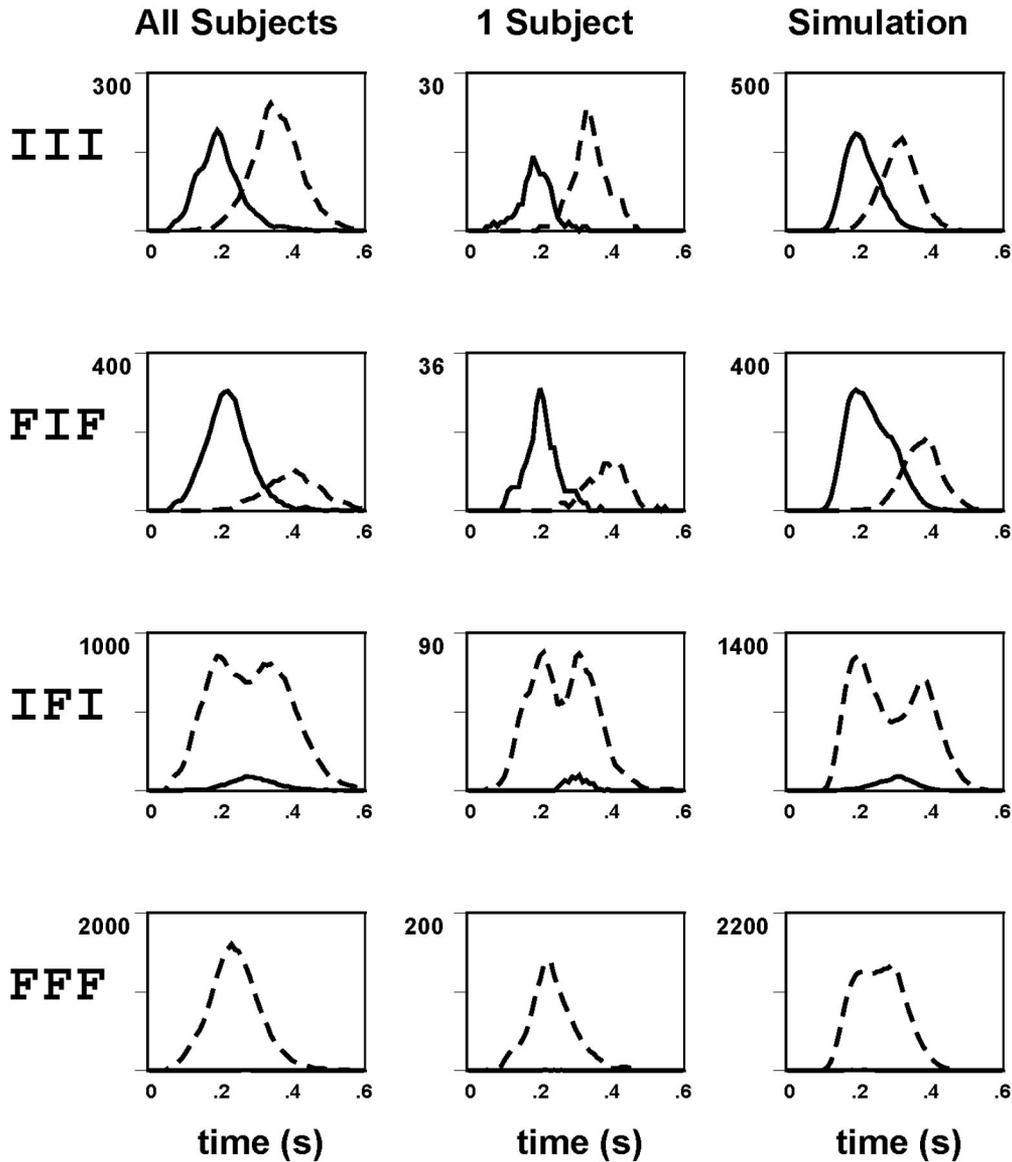


Figure 3. Response time (RT) histograms (10-ms bins) for correct (dashed lines) and incorrect (solid lines) trials. The left column shows empirical data pooled across participants. The middle column shows empirical data for a representative participant. The right column shows simulated data pooled across simulated participants. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition; FFF = frequent compatible condition.

a representative participant.) First, during an early period of information processing, errors tended to occur earlier than correct responses in the infrequent conditions (III, FIF). In contrast, responses in the frequent conditions (FFF, IFI) were nearly all correct during this early period. This observation is consistent with the inference that participants developed a bias to respond with the hand mapped to the frequent target, and it suggests that this bias was expressed during this early period. Second, during an intermediate period of information processing, a greater number of errors occurred at intermediate RTs in the incompatible conditions

(FIF, IFI) than in the compatible conditions (III, FFF). In particular, in the IFI condition, the distribution of error RTs was located in the middle of the distribution of correct RTs (for all but 1 participant, who did not make many errors in this condition), but virtually no errors occurred during this period in the FFF condition. Thus, the errors in the IFI condition are most likely attributable to the presence of the flankers, which favored the incorrect response. We infer that information related to the flankers impacted the motor system during this intermediate period. Third, during a late period of information processing, the responses

tended to be correct in all four of the conditions. We conclude that during this late period, information pertaining to the target stimuli impacted the motor system. Note that very few errors occurred in the FFF condition. In this case, all factors that influenced the choice of response (bias, flanker processing, and target processing) favored the correct response. Because few errors were generated in this condition, we do not consider these errors further.

As Figure 3 suggests, when pooled across correct and incorrect responses, the RT distributions associated with the III and FIF conditions were bimodal. For infrequent trials, the trough of the bimodal distribution corresponded to the RT at which the slow tail of the error distribution intersected with the fast tail of the correct distribution (III, 244 ms  $\pm$  35 ms; FIF, 308 ms  $\pm$  45 ms). To test this interpretation, for each infrequent condition we compared the magnitude of the combined distribution at this cross-over point with the magnitude of the combined distribution at the RTs corresponding to the modes of each of the individual distributions. For III stimuli, fewer trials occurred at the cross-over point than occurred at the times corresponding to the modes of the incorrect,  $t(14) = 5.2, p < .001$ , one-tailed, and correct,  $t(14) = 5.7, p < .001$ , one-tailed, distributions. Similarly, for FIF stimuli, fewer trials occurred at the cross-over point than occurred at the times corresponding to the modes of the incorrect,  $t(14) = 10.2, p < .001$ , one-tailed, and correct,  $t(14) = 2.8, p < .01$ , one-tailed, distributions. Thus, for both infrequent conditions, fewer responses were generated at the time when the correct and incorrect RT distributions were equally valued than occurred at the modes of the histograms.

These results indicate that the infrequent conditions were both characterized by a time period during which response generation was relatively unlikely. The beginning of the period was defined by the modes of the error distributions (compatible trials, 171 ms  $\pm$  26 ms; incompatible trials, 208 ms  $\pm$  30 ms), and the end of the period was defined by the modes of the correct distributions (compatible trials, 335 ms  $\pm$  45 ms; incompatible trials, 386 ms  $\pm$  56 ms). We suggest that, during this period, stimulus information began to impact the motor system, counteracting the response bias. As the system determined the identity of the infrequent target, the response channel associated with the frequent response was inhibited (producing fewer errors), and the response channel associated with the infrequent response was excited (producing more correct responses). This inference is supported by data related to latency of the P300, presented below.

Error commission in speeded RT tasks is typically followed on subsequent trials by a slowing of the response generation process, which increases the probability that those responses will be correct (Laming, 1979; Rabbitt, 1966b). This posterror slowing appears to be induced by the activation of an error-related control process. To evaluate the impact of such a process in this task, we averaged RTs for error trials and for correct trials immediately preceding and following the error trials (see Figure 4). It is not surprising that responses on error trials were faster than responses on correct trials both preceding,  $t(14) = 6.6, p < .001$ , and following,  $t(14) = -6.7, p < .001$ , the error trials, which indicates that the errors tended to occur when participants responded impulsively. More important, correct responses immediately following errors were slower than correct responses immediately preceding errors,  $t(14) = -5.1, p < .001$ . This posterror slowing suggests that the system adopted a more conserva-

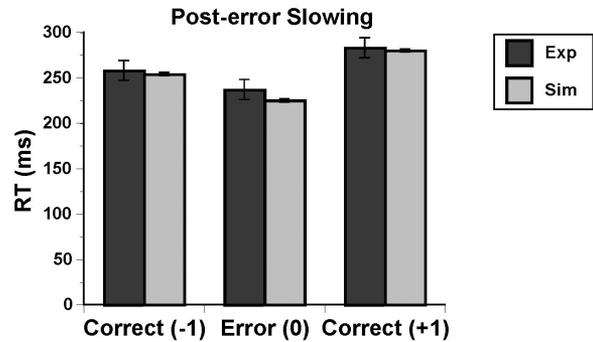


Figure 4. Posterror slowing. Average response times (RTs) on error (error 0) trials and on correct trials immediately preceding (correct -1) and following (correct +1) the error trials. Exp = empirical; Sim = simulated.

tive response bias following error commission, producing slower responses that were more likely to be correct.

### ERPs

**P300.** As is commonly found in ERP studies (for reviews, see Donchin & Coles, 1988; Johnson, 1988), the P300 was larger when elicited by infrequent compared with frequent stimuli,  $t(14) = 3.5, p < .005$ . A widely held view is that P300 latency indexes the duration of the stimulus evaluation process (Donchin, 1984; Duncan-Johnson & Donchin, 1982; Kutas, McCarthy, & Donchin, 1977; McCarthy & Donchin, 1981).<sup>5</sup> Accordingly, we used P300 latency to provide insight into the speed of stimulus evaluation in our study. Because the frequent targets elicited small P300s, we determined P300 latency from trials with infrequent targets only. For the infrequent stimuli, incompatible correct trials (628 ms  $\pm$  21 ms) were associated with longer P300 latencies than compatible correct trials (572 ms  $\pm$  18 ms),  $t(14) = 6.1, p < .001$  (Coles et al., 1985; Fournier, Scheffers, Coles, & Adamson, 1997; Masaki, Takasawa, & Yamazaki, 2000; Scheffers & Coles, 2000). Also for the infrequent stimuli, incorrect trials were associated with longer P300 latencies than correct trials: P300 latency to III stimuli (errors = 602 ms  $\pm$  17 ms, corrects = 572 ms  $\pm$  18 ms),  $t(14) = 3.2, p < .01$ ; P300 latency to FIF stimuli (errors = 649 ms, corrects = 628 ms),  $t(14) = 2.2, p < .05$ , as shown in Figure 5 (left panel). These results suggest that errors were more likely to occur when stimulus evaluation was slow compared with when stimulus evaluation was fast (Coles et al., 1985; Kutas et al., 1977; McCarthy, 1984; Scheffers & Coles, 2000; see also Donchin, Gratton, Dupree, & Coles, 1988).

We explored this observation further by selecting, for each participant, trials in the infrequent conditions in which the RTs occurred between the modes of the error distributions and the

<sup>5</sup> More recent theories have suggested that P300 latency may also be sensitive to processes involved in response selection and execution (Donchin & Coles, 1988). This possibility has engendered some controversy in the literature (e.g., Leuthold & Sommer, 1998; Verleger, 1997). Nevertheless, the preponderance of experimental evidence indicates that the latency of the P300 is primarily sensitive to the duration of the stimulus evaluation process (for a review, see Donchin & Coles, 1988).

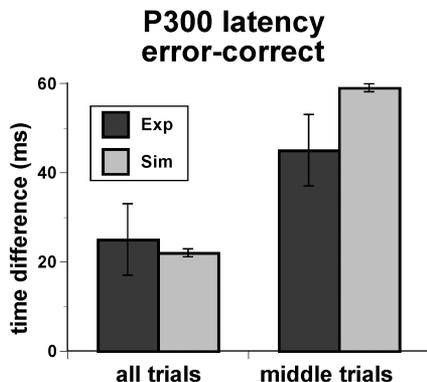


Figure 5. P300 latencies for the empirical (Exp) and simulated (Sim) data in the biased Eriksen flanker task. Shown are the differences in P300 latencies between error trials and correct trials, averaged across all trials in the infrequent conditions (left panel) and across a subset of those trials associated with intermediate response times (right panel). Positive values indicate that P300 latency was longer on error trials than on correct trials. Note that P300 latency was especially slow on error trials relative to correct trials for the subset of trials associated with intermediate response times.

modes of the correct distributions. When P300 latencies associated with these intermediate RTs were averaged according to response type, a two-way repeated measures ANOVA on compatibility and response type indicated that P300 latency was slower on incorrect trials (630 ms  $\pm$  69 ms) than on correct trials (585 ms  $\pm$  76 ms),  $F(1, 42) = 25.0, p < .001$ . Figure 5 (right panel) shows the difference in P300 latency between correct trials and error trials in this intermediate RT region, indicating that P300 latency occurred about 45 ms later on these error trials compared with these correct trials. It is important to note that the difference in latency was significantly larger for trials with RTs in the middle region (Figure 5, right panel) than for all trials (Figure 5, left panel),  $t(14) = -4.2, p < .001$ . Thus, not only were errors more likely to occur when stimulus evaluation was slow, the speed of the stimulus evaluation process had the greatest impact on the response selection process during this intermediate time period.

As we have discussed, this range between the modes of the RT distributions was associated with a period when response generation was relatively unlikely. We suggest that this range corresponds to the period during which accruing stimulus information began to impact the response selection process. During this period, the speed with which stimulus information was processed (as indicated by P300 latency) determined the course of response selection: When stimuli were evaluated quickly, the response was correct, but when stimuli were evaluated slowly, the response was incorrect.

**ERN.** Figure 6 (left column) shows the response-locked ERPs (for electrode Cz) for correct and incorrect trials for each stimulus condition; Figure 7 shows the amplitude of the ERN for the III, FIF, and IFI stimulus conditions. It is important to note that the amplitude of the ERN was larger in the IFI condition than in either of the infrequent conditions. These data were described in detail in Holroyd and Coles (2002; see also Coles et al., 2001; Holroyd, Praamstra, et al., 2002; Yeung et al., 2004). The positivity on

correct trials appears due, at least in part, to the P300, because the P300 was larger on the infrequent trials than on the frequent trials.

### Computational Model

To test our theory of error detection, we developed a computational model consisting of two components (see Figure 8): a task module that implements the stimulus–response mappings, and a monitor module that evaluates the appropriateness of the task module’s behavior. The specifics of these modules are described below. In brief, the task module consists of four layers: (a) A *perception* layer encodes the external input to the system, namely, the stimulus arrays appearing on each trial; (b) a *category* layer determines from the activity of the perception layer the identity of the target stimulus; (c) a *response* layer generates a response command by applying the stimulus–response mapping appropriate to the activity in the category layer; and (d) an *attention* layer facilitates the process of target categorization by increasing activity in the perception layer associated with the target stimulus and by inhibiting activity in the perception layer associated with the flanking stimuli. The attention layer also implements a response bias by differentially exciting the response options in the response layer. The computational principles that motivated the design of the task module are described below.

The monitor module implements a basic TD architecture, as described by the RL-ERN theory. Specifically, the monitor consists of three layers. First, a state layer determines the state of the system at any time. The state layer forms representations of the identity of the target stimulus and the executed response on each trial as well as the particular conjunctions of those stimuli and responses. For tasks in which external feedback is provided, the state layer also receives information about feedback stimuli presented on each trial. Second, a value layer attributes to the current state a probability that the trial will end in success or failure. Third, a TD unit determines the TD signal, that is, the instantaneous change in value associated with transitions between different states of the system. As we have described, a positive TD signal indicates that ongoing events are better than expected, whereas a negative TD signal indicates that ongoing events are worse than expected. The TD signal is sent to the value layer, where it is used to improve the system’s predictions of future success. The TD signal is also sent to the task module, where it is used to improve the task module’s performance of the task. The computational principles that motivated the design of the monitor module are described below.

In what follows, we used the model to simulate the performance and ERP data associated with the biased Eriksen flanker task.<sup>6</sup> For heuristic purposes, the simulation details associated with the task module and with the monitor module are presented separately.

<sup>6</sup> Model parameters were chosen so that the performance of the model matched the behavioral and electrophysiological data. Qualitatively similar patterns of results were found with a range of parameter values, which demonstrates that the simulation results follow from the processing principles incorporated into the model rather than the particular parameters used.

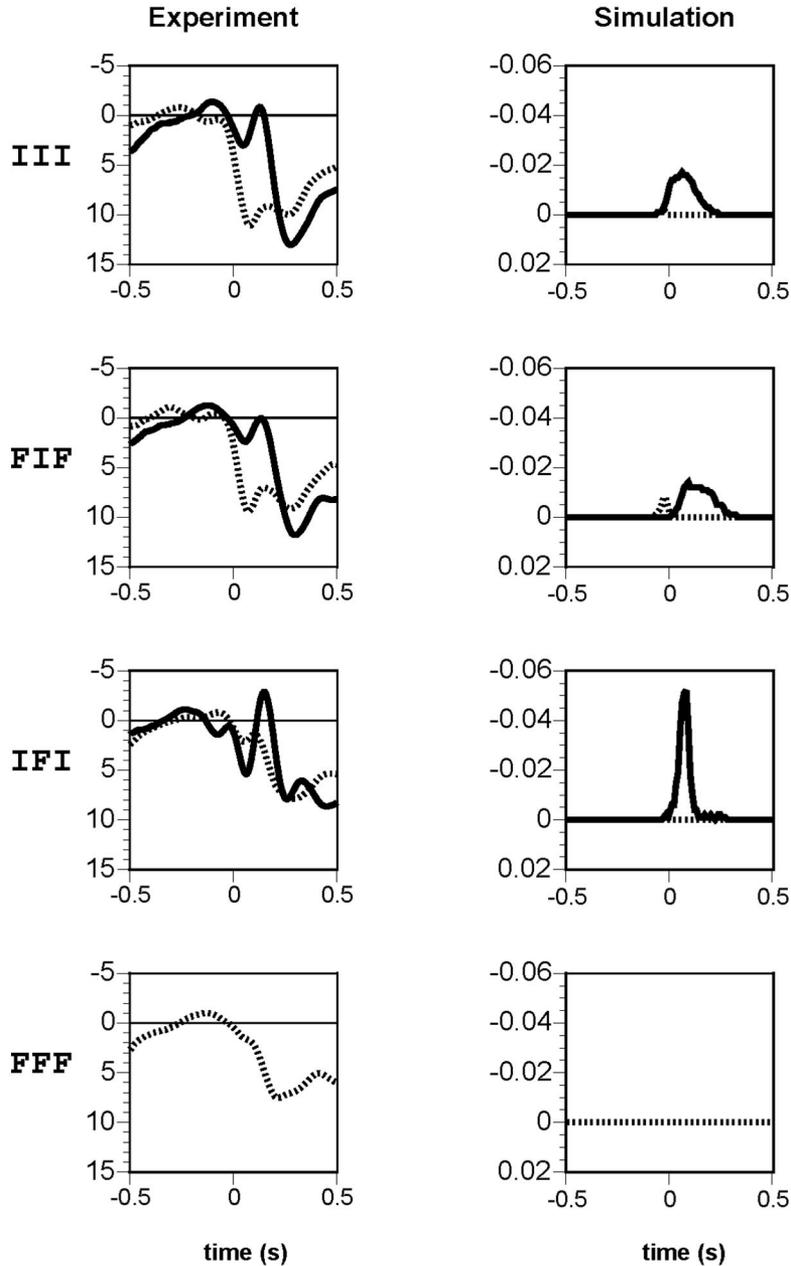


Figure 6. Event-related brain potential wave forms for correct (dotted lines) and incorrect (solid lines) trials, recorded at channel Cz. The left column shows experimental data ( $\mu V$ ), and the right column shows simulated data. Zero on abscissa indicates the time of response onset. The error-related negativity is the negativity following response onset on error trials. Note that, following convention, negative voltages are plotted upward. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition; FFF = frequent compatible condition.

*Task Component*

*Model*

In this section, we show that the task module captures the timing of the stimulus categorization and response generation

processes involved in the biased Eriksen flanker task and thus that the dynamics of the input to the monitor module are plausible. The construction of the task module was based on previous simulations of the Eriksen flanker task (Botvinick et al., 2001; Cohen et al., 1992; McClelland, 1992; Servan-

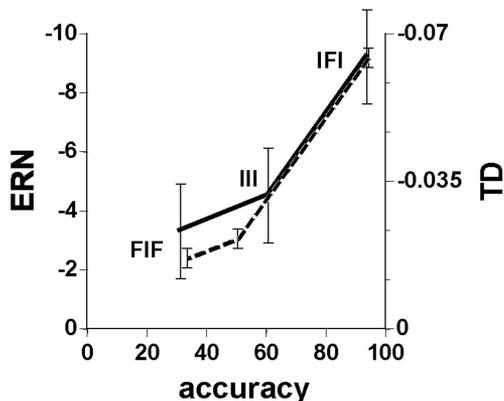


Figure 7. Average error-related negativity (ERN) amplitudes as a function of accuracy (percentage correct) for the empirical (ERN, in  $\mu\text{V}$ ) and simulated (temporal difference; TD) data. The solid line represents empirical data; the dashed line represents simulated data. IFI = frequent incompatible condition; III = infrequent compatible condition; FIF = infrequent incompatible condition.

Schreiber, 1990; Servan-Schreiber et al., 1998; Spencer & Coles, 1999; Yeung et al., 2004) and was guided by six central principles, described in the following sections. The first three principles reflect essential features of the model demanded by our test of the RL-ERN theory. The last three principles reflect important assumptions.

*Dynamic, continuous, and parallel processing.* As we have described, the primary objective of this study is to show that the monitor module can detect errors and produce ERNs on the basis of information provided to it by the task module, when the activity of the task module is free to wax and wane in a neurally and cognitively plausible fashion. Thus, the task module needed to reproduce the time when the target categorization and the response generation processes occurred on each trial. To implement this requirement, we adopted neural network principles that capture the continuous and dynamic nature of information processing in the brain (McClelland, 1992; Rumelhart & McClelland, 1986). Specifically, the task module was based on the interactive activation and competition model (McClelland & Rumelhart, 1988), in which processing units take on continuous activation values that evolve dynamically over time (see Appendix B).<sup>7</sup> These principles have already been used successfully to simulate performance on the Eriksen flanker task (Botvinick et al., 2001; Cohen et al., 1992; McClelland, 1992; Servan-Schreiber, 1990; Servan-Schreiber et al., 1998; Spencer & Coles, 1999; Yeung et al., 2004), so we expected that they could be extended in a straightforward manner to meet the demands of this study.

*Separate representations for category and response units.* The task module needed to have separate representations for the category and response units. This separation enables the monitor module to detect potentially inconsistent states (resulting from noise or prior expectations) between the category and response processes. Therefore, in contrast to the previous mod-

els of the Eriksen flanker task, our model contains an intermediate layer that intervenes between the perception and response layers (Figures 8 and 9; cf. Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999; Zhang et al., 1999). This categorization layer represents the system's evaluation of the identity of the center letter in the stimulus array, and the stimulus-response mappings are implemented by the flow of activity across the connections from the category layer to the response layer. Without such a layer, the monitor would not have been able to associate distinct values with a given stimulus, a given response, and their combination. Thus, each of the three input-output layers in the task module (perception, categorization, and response) can be considered a formally distinct level of cognitive processing. These levels are consistent with the functional organization of the brain, in which separate populations of neurons are associated with stimulus encoding, perceptual identification, and response production (Schall, 2002, 2003; Schall & Thompson, 1999).

*Modifiable weights.* The theory requires that the connections between the attention-response and response units be modifiable. In this way, the error signal can adjust the response bias applied to the response units. Although such connections were not used in this simulation, the theory also requires that the connections between the category layer and the response layer be modifiable, so that the task module can learn the appropriate stimulus-response mappings in tasks that demand such learning (Holroyd & Coles, 2002; Holroyd, Yeung, Coles, & Cohen, 2005). As we have described, this plasticity, which appears to be mediated by dopamine, seems to be characteristic of the neural systems involved in high-level motor control (for a review, see Schultz, 2002; see also Holroyd & Coles, 2002).

*Feed-forward inhibitory weights.* In contrast to previous models of the Eriksen flanker task, which depended on lateral inhibition between units, our model incorporates feed-forward inhibition. Recent work in our laboratory has demonstrated that neural networks simulate optimal performance of two-choice decision-making tasks when implemented with feed-forward inhibitory weights. These weights, which connect the decision and output layers of simple networks, must be equal in magnitude but opposite in sign to the feed-forward excitatory weights that connect those layers (Bogacz et al., 2005). Among the advantages of this mechanism is the fact that the inhibitory effects of one layer on another are immediate, whereas in the case of lateral inhibition they tend to be slow and less efficient. Furthermore, feed-forward connections are readily trained via reinforcement learning—one of the requirements of the task module—whereas a mechanism for training lateral connections with reinforcement learning has yet to be described.

*Attentional selection.* The attention layer is composed of three units: an attention-perception unit and two attention-response units (see Figure 9). In line with the model of attentional control proposed by Cohen, Dunbar, and McClelland

<sup>7</sup> The mathematics of these types of networks have been examined in detail by Grossberg (1978). In limiting cases, the networks are formally equivalent to accumulator models of decision making (e.g., Vickers, 1978), as described by Bogacz et al. (2005).

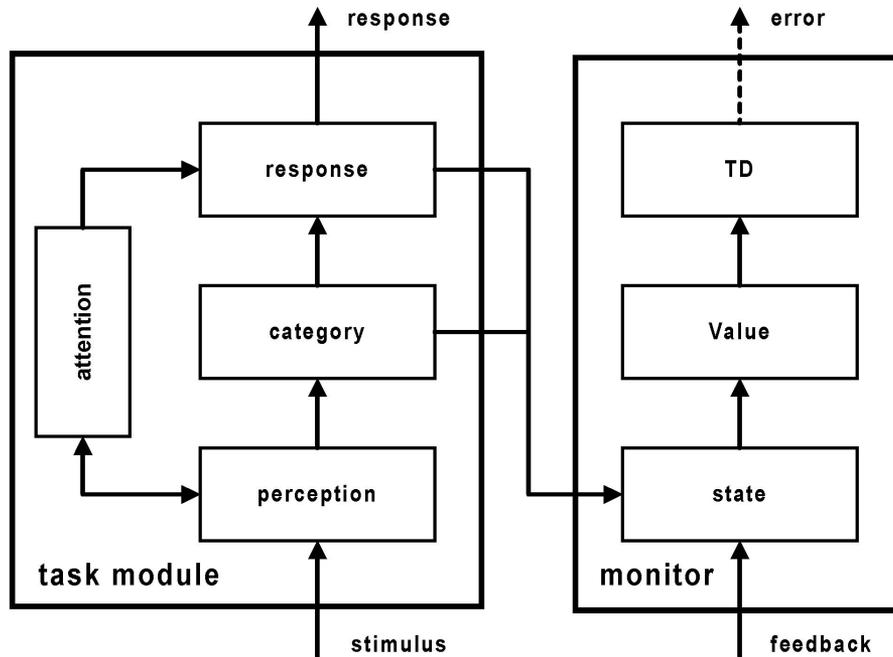


Figure 8. Model components. Boxes within the task and monitor modules correspond to layers of the network. Solid arrows indicate the direction of information transfer. The dashed line corresponds to the temporal difference (TD) signal produced by the monitor module.

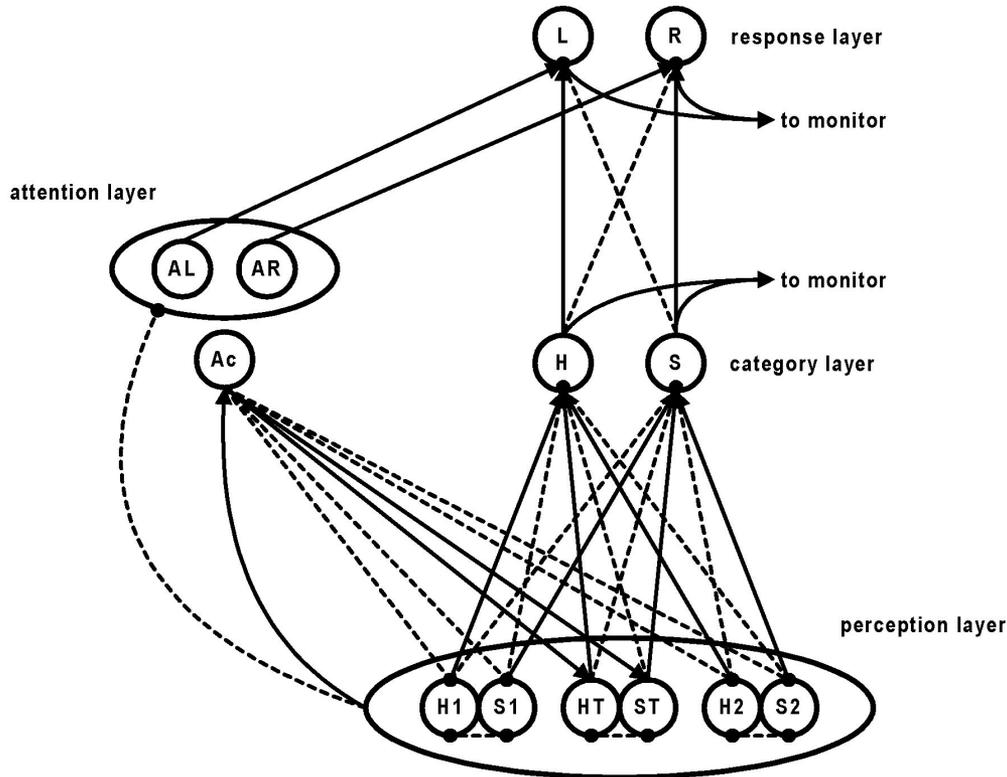
(1990), each of these units serves as “an additional source of input that provides contextual support for the processing of signals within a selected pathway” (p. 335). In our model, this principle is implemented as follows. First, the attention–perception unit excites the units in the perception layer that are associated with the center position of the array, and it inhibits units that correspond to the flanking distractor letters in the array. Conceptually, this attention unit implements a form of receptive field that amplifies the activity in its target region and suppresses activity outside that region. Second, the two attention–response units excite their corresponding response units, implementing the notion of a response set (as described in Simulation 6 of Cohen et al., 1990). An asymmetry in these weights favors one response type over the other, implementing the response bias corresponding to the frequencies of appearance of the target stimuli.

*Dynamics of attentional processing.* We assumed that the response bias can be triggered by coarse evaluation of the perceptual input, such that the response system can be activated by the mere presence of information in the input stream without identification of the specifics of that input. We implemented this assumption by providing external input to the attention–response units immediately on stimulus onset. In contrast, we assumed that the perceptual units become activated only after the external stimulus information has already undergone some low-level visual preprocessing. We assumed that this period corresponds to an early period of perceptual processing not accounted for by the model (e.g., Thorpe, Fize, & Marlot,

1996). Following the implementation of this idea in a previous simulation of the Eriksen flanker task (Spencer & Coles, 1999), we implemented this assumption by providing external input to the perception units after a variable period following stimulus onset. We also assumed that, once the stimulus information is available, that information should determine the outcome of further processing, and there should not be an influence of prior biases. This principle has been shown to be computationally optimal in speeded response tasks (Bogacz et al., 2005; see below for further discussion of this issue). We implemented this principle with inhibitory connections from all of the perception units to the attention–response units. This had the effect of shutting off the response bias as stimulus information passed through the input layer. Last, we assumed that perceptual input engages attention, which then induces the perceptual system to progressively focus on the target letter in the input stream. We implemented this idea with excitatory connections from all of the perception units to the attention–perception unit, which caused the unit to be activated by the perceptual input. Further specifics of the task module are presented in Figure 9 and outlined below.

*Input layer.* This layer is composed of five pairs of perception units, each pair corresponding to the *H* and *S* letters associated with the target and noise positions of the stimulus array. Each pair is subject to mutual inhibition, which ensures that either an *S* or an *H*, but not both, tends to be active for any position at any time.

*Category layer.* Each *H* and *S* unit in the input array excites the corresponding *H* or *S* unit in the category layer and inhibits the



*Figure 9.* The task module. Solid arrows indicate excitatory connections. Dashed lines indicate inhibitory connections. In the example shown, the *H* target is mapped to the left response, and the *S* target is mapped to the right response. Note that although the model simulates stimulus arrays with five letters (one target and four flankers), for clarity, only three letter positions are shown here. L = left response option; R = right response option; AL = left attention–response unit; AR = right attention–response unit; Ac = attention–perception unit; H = *H* category unit; S = *S* category unit; H1 and H2 = *H* noise perception units; HT = *H* target unit; ST = *S* target unit; S1 and S2 = *S* noise perception units.

opposing unit in that layer. This category layer represents the system’s evaluation of the identity of the target letter in the stimulus array.

*Response layer.* The category units excite and inhibit units in the response layer that corresponded to the left and right response options, thus implementing the stimulus–response mapping for the task. In the example shown in Figure 9, the *H* target was mapped to the left response option, and the *S* target was mapped to the right response option. A response is executed when the activation of a response unit crosses a specified threshold, followed by a variable interval. This interval accounts for a final period of the response generation processes that is associated with mechanical aspects of motor behavior.

We carried out the simulation by presenting to the network a succession of stimuli as they might occur in an actual experiment. A trial began when we supplied external input to the attention–response units by soft clamping their activity (i.e., by fixing the level of external input to the attention–response units while allowing their activity to fluctuate), which resulted in

boosted activation of the response unit that was mapped to the frequently occurring target. Then we soft clamped external input to the perception layer. Activation in this layer then propagated to the category layer. On compatible trials, this activity converged on the correct category unit in the category layer, but on incompatible trials, both the correct and incorrect category units received activation. Simultaneously, the perception units activated the attention–perception unit and inhibited the attention–response units. Activation of the attention–perception unit boosted the activity of the center perception units and suppressed the activity of the flanking perception units, ensuring that activity in the category layer converged on the correct category unit even on incompatible trials. This activity then passed to the response layer. Response execution occurred when the activity of one of the two response units crossed a specified threshold. Noise was added at each time step to the net input of each unit (see Appendix B), conferring some variability on the behavior of the model.

Errors occurred in two fashions. On some error trials, input from the attention–response units coupled with noise caused a

response unit to cross threshold before stimulus evaluation was complete, corresponding to the impulsive, fast-guess responses observed empirically in this task (Gratton et al., 1988). When the associated response was inconsistent with the stimulus on that trial, an error had occurred. On other error trials, activation from incompatible flankers in the perception layer activated the incorrect unit in the category layer before these channels were shut by the attention-perception unit. This activity then passed to the response layer, where it elicited an incorrect response. The simulation was run for 72,000 trials (4,800 trials for each of 15 simulated participants, corresponding to the empirical data), with the parameters given in Appendix B.

## Results

*Behavioral data.* As can be seen from Figures 2 and 3 (right column), the model accounts for the four salient features of the behavioral data we have described. First, nearly all the responses were correct on FFF trials. Second, fast responses on infrequent trials (III, FIF) tended to be incorrect. Third, a burst of errors occurred at intermediate RTs in the IFI condition. Fourth, relatively few responses were generated at intermediate RTs. These results are consistent with those of previous models of the Eriksen flanker task (Botvinick et al., 2001; Cohen et al., 1992; McClelland, 1992; Servan-Schreiber, 1990; Servan-Schreiber et al., 1998; Spencer & Coles, 1999; Yeung et al., 2004) but, in addition, capture the effects of the response bias associated with the unequal frequencies of appearance of the target stimuli in this version of the task.<sup>8</sup>

To compare the empirical and simulated RT distributions, we Vincentized the empirical and simulated distributions according to a procedure described by Ratcliff (1979). This procedure averages RTs at fixed quantiles across participants to produce a distribution representing that of the average participant. We then multiplied the resulting Vincentized distributions (decile probability levels) by the percentage of that trial type (correct, error) within the corresponding stimulus condition (III, FIF, IFI, FFF); this step ensured that a regression analysis would compare not only the shapes of the simulated and empirical RT distributions but also the proportion correct within each condition. The results of the regression analysis between the simulated and empirical Vincentized distributions confirmed that the two sets of distributions were highly similar,  $r(79) = .93$ ,  $p < .001$ , regression slope = .91, indicating that the task module successfully reproduced the behavioral data.

Thus, these results reflect the three periods of information processing seen in the empirical data (response bias, flanker processing, and target processing), which the model reproduced by simulating the combined influence of the response bias and the flanking letters on the response selection process. These processes are illustrated by the response unit activations in Figure 10. The effects of the response bias are evident in an early increase in activity of the incorrect response unit on infrequent (III, FIF) trials and in an early increase in the activity of the correct response unit on frequent (IFI, FFF) trials: This activity induced a high proportion of impulsive errors in the infrequent conditions and many fast correct responses in the frequent conditions. The influence of the flanker letters on the response units can be seen in the increase in

the incorrect unit activity and the dip in correct unit activity at intermediate RTs in the IFI condition: This momentary reversal in unit activations gave rise to the greater probability of error commission at intermediate RTs in this condition. Finally, the impact of information pertaining to the target letter can be seen in a second reversal of the response unit activity in the IFI condition and in the delay in the onset of the correct unit activity in the FIF condition relative to the III condition: This activity gave rise to the correct responses associated with slow RTs.

*P300 latency.* As we have noted, the latency of the P300 is thought to index the duration of the stimulus evaluation process. To confirm that the task module correctly reproduced the dynamics of the stimulus categorization and response generation processes and thus that the task module provided veridical input to the monitor module, we examined a measure of P300 latency defined by the time at which the associated category unit in the task module exceeded a threshold and the competing category unit in the task module was inactive. This corresponded to the onset time of binary task state units in the monitor module (which we describe below). Thus, we used the onset time of the target task state units in the monitor module as a measure of the duration of the stimulus evaluation process in the task module and compared this with the empirical data concerning the latency of the P300.

The model reproduces two commonly observed findings related to observations about P300 latencies that were also obtained in this experiment: Stimulus evaluation time was longer on incompatible correct trials than on compatible correct trials (95 ms  $\pm$  2 ms difference; Coles et al., 1985; Fournier et al., 1997; Scheffers & Coles, 2000), and in the infrequent conditions, stimulus evaluation time was longer on incorrect trials than on correct trials (22 ms  $\pm$  2 ms difference), as shown in Figure 5 (left panel). This result is consistent with the hypothesis that the longer the system takes to evaluate the stimulus, the more likely it is to make an error (Coles et al., 1985; Kutas et al., 1977; McCarthy, 1984; Scheffers & Coles, 2000; cf. Donchin et al., 1988). To explore this observation further, we defined early, intermediate, and late ranges of the RT distributions in the infrequent conditions (see above). We found that stimulus evaluation time was slower on error trials than on correct trials in the intermediate range of the RT distributions. As with the empirical data, furthermore, stimulus evaluation

<sup>8</sup> C. W. Eriksen and Schultz (1979) have also demonstrated that if the experimental stimuli consist of four possible target letters, two of which are mapped to the same response (e.g., *H* and *K*), then responses to incompatible stimulus arrays that are mapped to the same response (e.g., *HH-KHH*) are slower than responses to compatible stimulus arrays (e.g., *KKKKK*) by about 10 ms. In a separate simulation, we have adapted our model to this task and have reproduced this effect of stimulus incompatibility. The task module includes 20 perception units (1 for each of the four possible targets for each position in the stimulus array) and 4 category units. On stimulus incompatible trials, activation of the noise letters in the stimulus array inhibits activation of the category unit associated with the target, whereas on stimulus compatible trials, no such inhibition occurs. This differential inhibition delays the response generation process on stimulus incompatible trials relative to compatible trials. See Zhang et al. (1999) for a comparable finding.

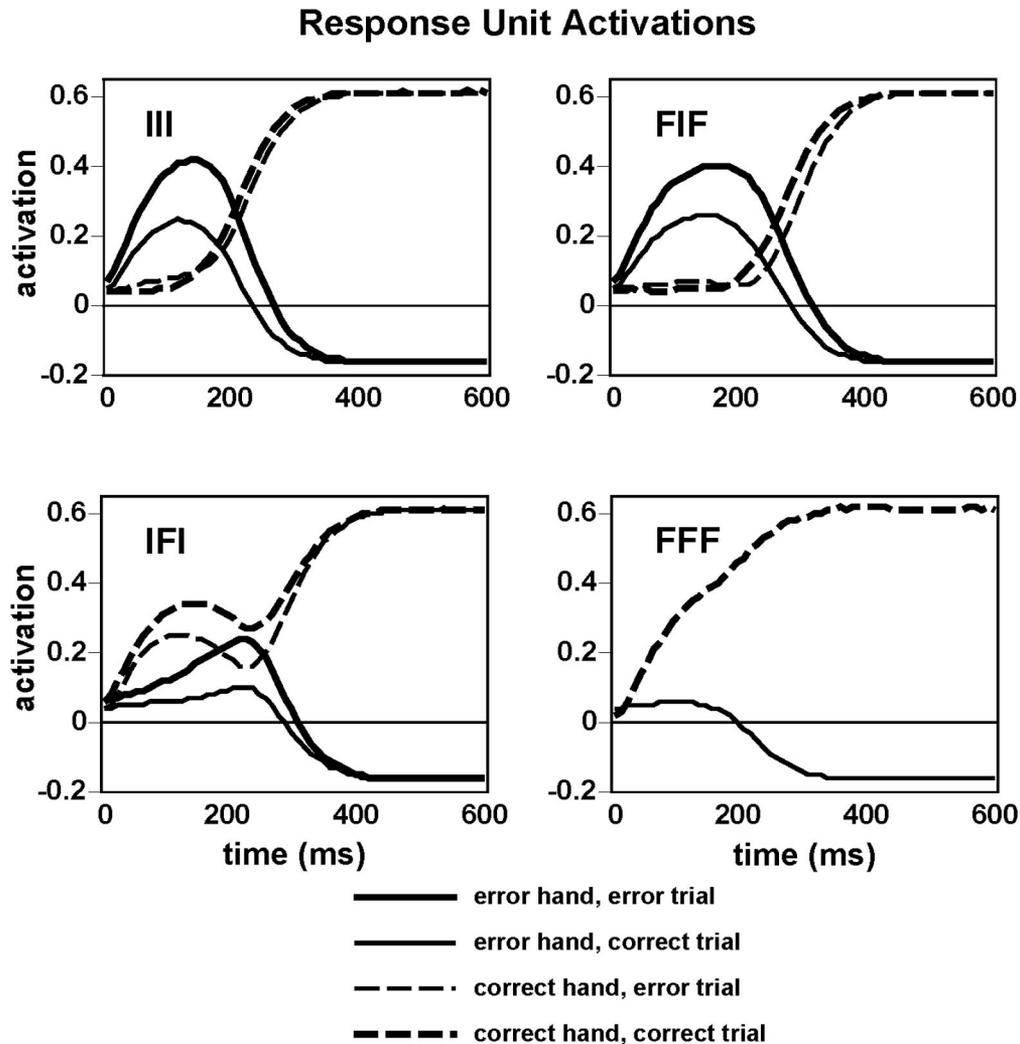


Figure 10. Simulated response unit activations, averaged across trials according to the correctness of each unit's associated mapping (error hand, correct hand) and to trial outcome (error trial, correct trial). Zero on abscissa indicates the time of stimulus onset. III = infrequent compatible condition; FIF = infrequent incompatible condition; IFI = frequent incompatible condition; FFF = frequent compatible condition.

time was slower on error trials than on correct trials for responses made during this intermediate period (Figure 5, right panel) compared with responses made on all trials (Figure 5, left panel). This result suggests that the impact of the speed of the stimulus evaluation process on response selection was greatest for trials during this period. Moreover, this period was characterized by response quiescence, when the flow of stimulus information to the motor system began to counteract the response bias. In the model, these simulated data occurred because the outcome of the response selection process during this period depended critically on the speed with which the motor system obtained stimulus-related information. The consistency between the empirical and simulated results suggests

that a similar mechanism may underlie human behavior in the task and provides further confirmation that the dynamics of the task module and the information it provides to the monitor module are plausible.

#### *Monitor Component*

The results we have presented indicate that the first objective of this study has been satisfied: We have developed a task module that simulates the dynamics of the stimulus categorization and response generation processes and that successfully captures important features of the empirical data. In the present section, we show that the monitor module is capable of detecting errors on the

basis of information provided to it by the task module. We further show that the task module is capable of using this error information to improve performance on the task at hand.

### Model

We considered the following four principles in the design of the monitor module.

*Consistency with original RL-ERN model.* The monitor module was derived from a series of recent adaptive critic models (i.e., models that involve monitors based on the method of temporal differences) of dopamine and the basal ganglia (e.g., Barto, 1995; Montague et al., 1996; Schultz et al., 1997). In particular, the model was derived from our previous RL-ERN simulation of the biased Eriksen flanker task, which, as we have noted, did not include a task module (Holroyd & Coles, 2002). As a guiding principle in the design of the monitor module, we tried to keep it as similar as possible to the design of the original model. That model contained units that (a) classified external states, (b) assigned good and bad values to those states, and (c) computed the TD signal (the change in the value of those states). These TD signals were used by the monitor to improve its value estimates. Among the state units were ones that separately represented the state of the category and response units in the task module and the state of external feedback as well as state conjunction units that were sensitive to combinations of the category and response units, such that values could be independently assigned to each possible conjunction. These features of the original model were all incorporated into the present model.

*No privileged access to information.* A guiding principle in the design of the model was that the monitor module should not have privileged access to information not available in explicit form to the task module. For this reason, the monitor described in this study contains a unit for each target (frequent and infrequent), but not (as in the original model) a unit for each full stimulus configuration (i.e., III, FIF, FFF, and IFI).

*Binary units and context sensitivity to task states.* The monitor units in this study act in an approximately binary fashion, ensuring rapid transitions in value between consecutive states and thus that the learning algorithm is well behaved.<sup>9</sup> The biological plausibility of such binary units is supported by the presence of cells in the basal ganglia that act in a binary fashion. As identified from in vitro studies in the rat, these spiny neurons have discrete on and off states that correspond to long periods of inactivity interrupted by short periods of activity lasting up to seconds (Wilson, 1995). These neurons appear to act as binary switches that maintain their state despite conflicting inputs (Gobbel, 1995), perhaps allowing for pattern classification of cortical activity (Houk, 1995). Furthermore, the behavior of these neurons reflects their dependency on task context, “encoding, rather separately, all individual task events occurring between the initial cues and the outcome of action” (Schultz, 1995, p. 37). This activity

is not sufficiently explained by the physical characteristics of the stimuli presented or the movements performed but depends on certain behavioral situations, certain conditions, or particular kinds of trials in a given task, thus showing relationships to the context in which the particular events occurred. (Schultz, Apicella, Romo, & Scarnati, 1995, p. 12)

For example, some cells in the basal ganglia fire only after an incorrect response (Kermadi & Joseph, 1995). All of these observations are consistent with the behavior of the monitor units in our model, which switch on and off in a binary fashion at key transition points between states throughout each trial. In particular, the combinatorial quality of the conjunctive units in the model (see below) is similar to the context dependence of the spiny neurons, which reflects sensitivity to higher order features of the task situation. Of course, this is not to suggest that other neural areas are not characterized by neurons with such properties nor that the basal ganglia are composed only of neurons with such properties; rather, it is to say that the basal ganglia contain a subset of neurons with properties that are sufficient to carry out the functions described here and implemented in the value units of our model.

*Flexible configuration.* We assumed that the monitor module units could be configured as needed, with a different set produced for each new task. Likewise, the initial weights of the value units were assumed to be programmable. These values were initialized at the start of the task, with some clamped such that they were kept constant throughout the task. The theory does not specify how this process might occur (however, we suspect the involvement of the prefrontal cortex acting on the basal ganglia via the loops that interconnect these brain areas; for a review, see Wise, Murray, & Gerfen, 1996). In general, a central and open question in executive control theory relates to how task instructions, as communicated by an experimenter, are encoded by the participant (E. K. Miller & Cohen, 2001; O’Reilly & Munakata, 2000). This problem is beyond the scope of our model.

The specifics of the three layers of the monitor module are as follows (see Figure 11).

*State layer.* This layer consists of three subsets of units: a set of *task state* units that detect states of different parts of the task module, a set of *conjunction state* units that form pairwise combinations of those task states, and a pair of *feedback state* units. The task state units monitor the state of the task module. This

<sup>9</sup> When the units have continuous activation functions, the learning algorithm is not well behaved. Imagine, for example, a value unit with a weight that fully predicts reward (i.e., equal to 1). Imagine further that the unit takes two cycles to reach maximum activation, outputting, for instance, 0.5 on the first cycle and 1.0 on the second cycle. This transition would produce a positive TD signal of 0.5 (value on the first cycle =  $0.5 \times 1.0$ ; value on the second cycle =  $1.0 \times 1.0$ ; TD signal = value on the second cycle – value on the first cycle). The positive TD signal would then increase the weight associated with the value of the unit—for instance, from 1.0 to 1.1. Such an increase in weight would occur every time that unit was activated, precipitating a weight explosion. This is to suggest not that the reinforcement learning algorithm could never be implemented in a continuous framework but rather that the mechanisms for doing so are only now being developed (e.g., Doya, 2000). Note also that the feed-forward inhibition from the state layer to the value layer provides a similar function. For example, without feed-forward inhibition, the values associated with the task state units and conjunction state units would sum together following both events, even if both of the activated value units individually predicted full reward. The transition between states would elicit a positive temporal difference error, which, in turn, would increase the values associated with the individual target and response states, giving rise to a weight explosion. Feed-forward inhibition from the state layer to the value layer ensures that this outcome does not occur.

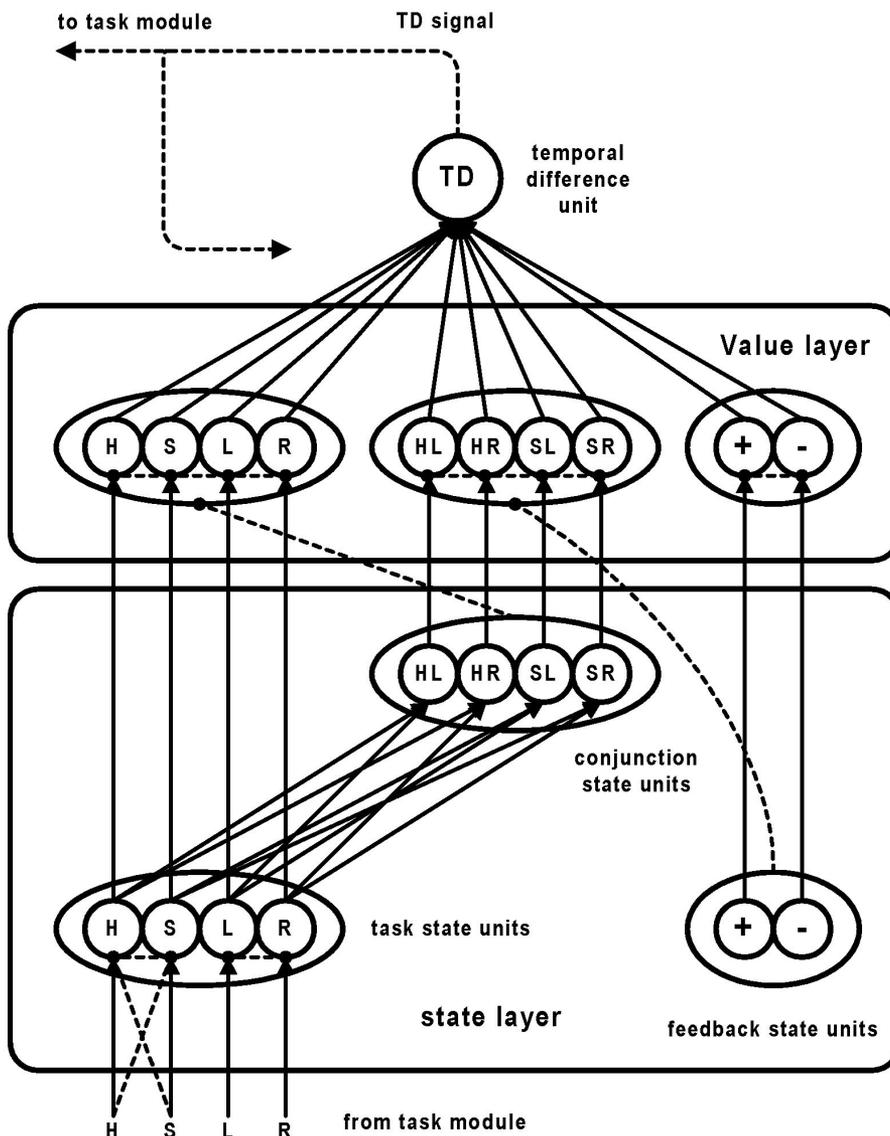


Figure 11. The monitor module. Plus and minus signs indicate units associated with positive and negative feedback, respectively. Solid arrows indicate excitatory connections; dashed lines indicate inhibitory connections. The dashed arrow indicates the TD signal produced by the monitor module. TD = temporal difference; H = H target unit; S = S target unit; L = left response unit; R = right response unit; HL = H target-left response unit; HR = H target-right response unit; SL = S target-left response unit; SR = S target-right response unit.

component consists of two pairs of units: one pair for target detection, and one pair for response detection. Thus, the task state component contains units for both target letters and for both response options. Each of these four units receives excitatory input from the corresponding unit in the task module. To ensure that the target task state units become activated only after the stimulus categorization process is complete, these units also receive inhibitory input from the competing units in the task module (see Figure 11). Thus, for example, the H categorization unit in the task module excites the H target task state and inhibits the S target task state units in the monitor module. Furthermore, because the cor-

rectness of the trial is defined by the identity of the first response, the monitor uses the first response detected to evaluate the correctness of the trial. For this reason, each unit within a pair strongly inhibits the other unit in that pair (see Figure 11). This strong inhibition, coupled with self-excitation (see Appendix C, Table C1), ensures that the first unit of each pair to be activated remains active until the end of the trial. Critically, if the task module generates a second response following an error (an error correction), the response detection unit activated by the initial response remains active, and the response detection unit associated with the second response remains inactive.

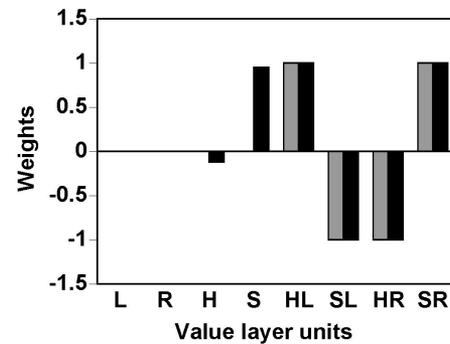
The conjunction units in this layer detected combinations of task module states, driven by input received from the task state units. Units existed for each combination of stimulus and response. Weights between the task state units and the conjunction state units were set such that each conjunction state unit was active only when it received input from both of its task state units. The remaining state units consist of a unit for positive feedback and a unit for negative feedback. Although they are not used in this study, these units serve to encode external feedback for tasks in which feedback is provided.

**Value layer.** This layer assigns a degree of value (goodness or badness) to each possible state of the system. The layer contains 10 units, each of which activate in response to the activation of a corresponding state unit and represent its value (as we describe in the *TD layer* section). It is important to note that a conjunction value unit is associated with each of the four conjunction state units, assigning a degree of value to each possible stimulus–response combination. These four value units compose the heart of the error-detection mechanism: Error detection corresponds to the activation of a conjunction unit assigned a negative value. No more than one value unit is active at any time, and the value of the feedback takes precedence over that of the conjunction units, which, in turn, take precedence over that of the individual state units (Holroyd & Coles, 2002). This relation is implemented with inhibitory projections from the conjunction state units to the value units corresponding to the task state units and from the feedback state units to the value units corresponding to the conjunction state units.

**TD layer.** Finally, the unit in the TD layer receives input from each of the value units. The strength of the weights between the value units and the TD unit represents the value of the given state. Positive weights indicate good states, and negative weights indicate bad states. As no more than one value unit is active at any time, the value of a given state of the network is determined solely by the weight associated with the currently active value unit. The TD unit computes the TD signal, which is the difference between the value of the state on the last cycle and the value of the state on the present cycle (Sutton, 1988; Appendix C).

In the simulation, the value weights were initialized with values corresponding to the definition of the task. Thus, the value weights corresponding to each of the conjunction units were initialized according to the stimulus–response instructions given to the participant (Figure 12, gray bars); weights associated with correct combinations of stimuli and responses were initialized to 1 (indicating good combinations), whereas weights associated with incorrect combinations of stimuli and responses were initialized to  $-1$  (indicating bad combinations). The remaining value weights were initialized to zero. The TD signal was carried from the TD unit to the value layer, where it updated the value weights such that the system's value estimates improved with exposure to the task (Sutton, 1988; Sutton & Barto, 1998; Appendix C).

The ERN was defined as the negative of the activity of the TD unit (Holroyd & Coles, 2002). We measured the amplitude of the simulated ERN from base to peak, as it would be measured in an ERP experiment. Specifically, we averaged the ERN data across trials by condition, which yielded three ERPs associated with error trials in the III, FIF, and IFI conditions. Then we determined the maximum (negative) value of the ERN within 150 ms (15 cycles) following the response. We also determined this value at the time



*Figure 12.* Value layer weights. Gray bars represent weights at the start of the simulation. Black bars represent weights at the end of the simulation. In this simulation, the *S* target occurred with high frequency, the *H* target occurred with low frequency, the *H* target was mapped to the left response, and the *S* target was mapped to the right response. L = left response unit; R = right response unit; H = *H* target unit; S = *S* target unit; HL = *H* target–left response unit; SL = *S* target–left response unit; HR = *H* target–right response unit; SR = *S* target–right response unit.

of the response, and we defined ERN amplitude as the difference between these two values. In practice, the value associated with the response was always close to zero, so ERN amplitude was just the maximum value of the negative TD signal (averaged across trials) within 150 ms following the response.

## Results

**ERN.** Trials were classified by the monitor as correct or incorrect according to the valence of the value unit associated with the conjunction state unit activated on that trial. Out of 72,000 simulated trials, the monitor misclassified only 1.1% as being either correct or incorrect, as a result of noise associated with the activity of the category layer units. Figure 6 (right column) presents response-locked averages of the simulated ERPs on correct and incorrect trials for each stimulus condition. On error trials, the ERN reached maximum amplitude within 100 ms after the response. This observation is important given that the empirical ERN reaches maximum amplitude within 100 ms following response onset (Falkenstein et al., 1990; Gehring et al., 1993). The result indicates that the monitor was able to detect errors even as they occurred.

Figure 7 presents the amplitude of simulated ERNs for the III, FIF, and IFI conditions. Consistent with the empirical data (Figures 6 and 7), the amplitude of the ERN was larger for the IFI condition than for the infrequent conditions. However, the model somewhat underestimated the amplitude of the ERN in the III and FIF conditions, a result also obtained in Yeung et al.'s (2004) simulation of these data. In fact, the error trials in the III and FIF conditions are relatively fast (about 200–250 ms), and it has been shown that the ERN on such trials overlaps in time with the N200 (Hajcak, Vidal, & Simons, 2004). Thus, the empirical measurements on these trials may overestimate the amplitude of the ERN because of contamination by the N200 (see also Coles et al., 2001). However, it is also possible that the amplitude of the actual ERN may not, in fact, be linearly related to the size of the simulated ERN in our model (i.e., to the TD error).

As in the original RL-ERN model, the amplitude of the simulated ERN depends on the change in value at the time of response generation. Because trials with frequent targets were normally correct, the imperative stimulus on frequent trials was associated with a large positive value ( $S$  in Figure 12). Thus, errors on those trials tended to result in a large change in value, from very good (following target categorization) to very bad (following error commission). In contrast, because trials with infrequent targets were less likely to be correct, the imperative stimulus on infrequent trials was associated with a small negative value ( $H$  in Figure 12). Therefore, errors on those trials tended to result in a smaller change in value, from bad (following target categorization) to worse (following error commission).

Note that the reduction in simulated ERN amplitude on infrequent trials relative to frequent trials (see Figure 7) is not an artifact of latency jitter across conditions (Coles et al., 2001). We matched error trials in the IFI condition with error trials in the FIF condition by RT (the III condition was not included because this condition did not contain enough slow error trials). As shown in Figure 13, the difference in ERN amplitude remained even when the IFI and FIF error trials were matched by RT.

**Posterror slowing.** Figure 4 presents simulated RTs for error trials and for correct trials immediately preceding and following the error trials. Correct trials were slower following errors than preceding errors. This result indicates that error commission tended to slow the response generation process on subsequent trials. In the model, this slowing was implemented by decreasing the strength of the connections between the attention–response and response units by an amount that was proportional to activity of the TD unit. Thus, the model uses principles of reinforcement learning to capture the common empirical finding of posterror slowing.

This aspect of the model yielded a testable prediction: The amount of slowing should be greater on trials with large ERNs than on trials with small ERNs. Specifically, the model predicts that the greatest slowing will occur on error trials with frequent targets, because error responses induce the largest ERNs on those trials (Figures 6 and 7). Figure 14 illustrates RTs on correct trials with frequent and infrequent targets, averaged according to whether they were preceded by error trials with frequent or infrequent targets. As can be seen from the figure, correct trials,

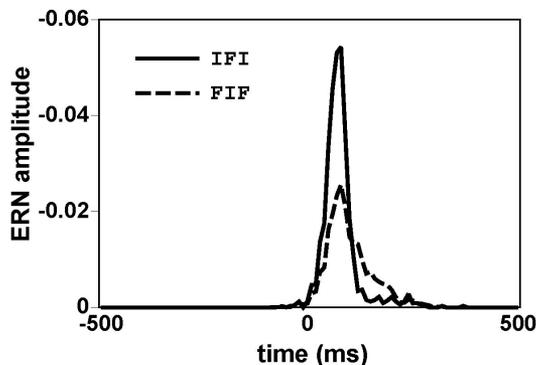


Figure 13. Simulated error-related negativities (ERNs) associated with the infrequent incompatible (FIF) and frequent incompatible (IFI) conditions matched by response time.

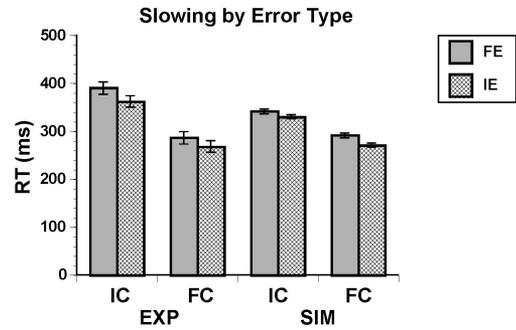


Figure 14. Average response times (RTs) on correct trials with frequent (FC) and infrequent (IC) targets, preceded by error trials with frequent (FE) and infrequent (IE) targets, for both the empirical (EXP) and the simulated (SIM) data.

irrespective of target frequency, were slower when preceded by error trials with frequent targets than by error trials with infrequent targets.

To investigate this possibility in the empirical data, we separated correct trials associated with frequent and infrequent targets according to whether these trials were preceded by error trials with frequent or infrequent targets (Figure 14). A two-way repeated measures ANOVA indicated that RT was faster on correct trials preceded by error trials with infrequent targets than by error trials with frequent targets,  $F(1, 14) = 35.3, p < .001$ . This result indicates that, indeed, errors on trials with frequent targets were associated with more posterror slowing than errors on trials with infrequent targets. The ANOVA also confirms that correct responses on frequent trials were faster than correct responses on infrequent trials,  $F(1, 14) = 129.0, p < .001$ , but there was no interaction between target frequency on the current (correct) trial and target frequency on the preceding (error) trial ( $p = .23$ ). These findings are consistent with a central claim of the RL-ERN theory: that the ERN reflects a control signal used by the motor system to adapt behavior and that the amplitude of the ERN is correlated with the size of the behavioral modification.

## General Discussion

In the present study, we have shown how an error detection mechanism, when implemented in a neurally and cognitively plausible connectionist architecture, can rapidly detect errors in a speeded RT task. As such, the model represents one of the first computationally and biologically plausible mechanisms for online error detection. The monitor detects errors by identifying states of the system associated with negative value—namely, by identifying particular combinations of stimuli and responses. A critical aspect of this mechanism is that it accounts for the latency of the ERN, which peaks within 100 ms following the error response, in addition to ERN amplitude. The mechanism does so on the basis of input from the task module that captures the timing of the stimulus categorization and response generation processes. We have also shown how the error signals produced by the error detection mechanism can be used to regulate performance on a trial-to-trial basis. The model accounts for the common finding that response generation slows following error commission (Laming, 1979; Rab-

bitt, 1966b), and it suggests that such slowing should be greater on trials following errors with large ERNs than on trials following errors with small ERNs.

This model was derived from previous work in which we showed that an error detection mechanism, based on these same principles of reinforcement learning, could use external feedback to learn internal representations of appropriate behavior, such that the system could learn to detect errors in the absence of feedback (Holroyd & Coles, 2002). That study illustrated how a task module could use the error signals produced by the error detection mechanism to improve performance on the task—for example, by learning stimulus–response mappings. The study also showed that the error detection mechanism accounted for the amplitudes of two ERP components associated with performance monitoring: the response ERN and the feedback ERN. In so doing, it integrated two apparently disparate phenomena—one waveform associated with internally generated responses, and another associated with externally provided feedback—into a common theoretical framework. For example, as the system internalized the appropriate stimulus–response mappings in a trial-and-error learning task, the amplitude of the feedback ERN decreased and the amplitude of the response ERN increased. This transfer occurred as the system came to rely less on external feedback and more on internal monitoring of behavior to detect errors. It is important to note that, in work described elsewhere, we have demonstrated that the current model replicates these findings of the original model (Holroyd, Yeung, et al., 2005). Taken together, this work provides a unified theoretical framework for understanding performance monitoring and reinforcement learning.

### *Evaluating the Theory*

One potential concern about our model is its complexity, relative to other, simpler models of similar phenomena (Figures 9 and 11). It has been shown that relatively simple three-layer feedforward neural networks with enough hidden units can approximate any well-behaved mathematical function (Hornik, 1991). Because of this property, it is sometimes argued that neural network models are unfalsifiable (e.g., Massaro, 1988). Therefore, it is a fair question to ask whether our model has elastically accommodated both the data and our preconceived ideas about how the system should behave. We believe that there are several reasons why this is not the case.

First, the model simulates a wide variety of empirical data, which include RT distributions for both correct trials and error trials, response slowing following errors, the amplitude and latency of the response ERN, and the relative latencies of the P300 on correct and error trials. In accounting for these observations, the model demonstrates that it can detect its own errors. More generally, the theory provides a unified account of the feedback ERN and the response ERN. These two ERP components might appear, at face value, to be very different empirical phenomena, so their unification in a single conceptual framework represents a potentially important theoretical advance. It should be noted that these are not trivial successes. First, the latency of the response ERN is determined by the relative timing of the stimulus categorization and response generation processes. On trials in which the stimulus categorization process is delayed with respect to the response (as

occurs on fast-guess error trials, e.g.), the latency of the response ERN could, in principle, be dissociated from the response. Second, the amplitude of the response ERN is determined both by the latency of the ERN and by the size of the prediction error: Small ERNs can be due either to latency jitter of larger ERNs or to small prediction errors. Third, the amplitudes of the response ERN and the feedback ERN are inversely related, so simulated changes to one component affect the other component. Together with the RT distributions, P300 latencies, and response slowing, these factors compose a set of multiple interdependent constraints that the model must satisfy simultaneously. That the model does so attests to the plausibility of the error detection mechanism.

Second, the architecture of the model is constrained by features of neurobiology that were not chosen specifically for the purpose of these simulations (for a discussion of this issue, see Seidenberg, 1993). For example, the model is based on principles of reinforcement learning, specifically the method of temporal differences, because the midbrain dopamine system is thought to implement principles of reinforcement learning and to carry TD signals. Furthermore, the model provides a unified theoretical account of the disparate phenomena that give rise to the data, from biology to behavior. In so doing, the model provides insight into those phenomena. For example, our model suggests that error detection is associated with the attribution of negative values to ongoing events, that this process occurs in the basal ganglia, and that the ERN reflects a learning process in the anterior cingulate cortex. These insights are something that a three-layer neural network with many hidden units could not provide, even if such a network could reproduce all of the empirical observations. In general, the complexity of the human brain suggests that overly simplistic theories of cognitive function are unlikely to be very informative (O'Reilly & Farah, 1999; O'Reilly & Munakata, 2000).

Third, the model and the theory from which it is derived make a wide variety of testable predictions. A number of predictions about the feedback ERN have already been confirmed in several experiments, and other empirical observations have been informed by the theory as well (for reviews, see Holroyd, Nieuwenhuis, Mars, & Coles, 2004; Nieuwenhuis, Holroyd, Mol, & Coles, 2004). For example, studies have confirmed the prediction that the amplitudes of the response ERN and feedback ERN are inversely related (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002). In the present study, we have made and confirmed an additional, new set of predictions (concerning the relation of the ERN to posterror slowing). Such predictions allow various parts of the theory to be falsified. For instance, observation of a normal ERN despite impairment of the midbrain dopamine system would provide strong evidence against the theory's central proposition: that ERN production depends on dopamine. Likewise, if the ERN were produced in a brain area other than the anterior cingulate cortex, this fact would falsify this aspect of the theory. In addition, if the ERN were shown not to behave as a temporal difference error, then the idea that the ERN indexes this kind of reinforcement learning signal would also be called into question.

Fourth, the RL-ERN model and the response conflict model (Yeung et al., 2004) are the only existing models of high-level error detection. Together with the work of Spencer and Coles

(1999), which simulated motor potentials in the Eriksen flanker task, these are, to our knowledge, the only computational studies that have simulated ERP indices of cognitive phenomena in addition to behavioral measures. This fact offers the promise that these models can be used to make contrasting, quantitative predictions about behavioral and electrophysiological findings. Furthermore, the models provide a benchmark by which future theories of error detection can be evaluated.

### *Comparison With Other Theories of the ERN*

The RL-ERN theory shares much in common with both the conflict monitoring theory and another theory of error detection, the mismatch theory. Here we compare the RL-ERN theory with each of these theories.

#### *Mismatch Theory*

As articulated by Coles et al. (2001), the mismatch theory holds that the monitor is a comparator that compares the actual response executed with the desired response. Comparators are common in the neurobiology of motor control literature (for reviews, see Desmurget & Grafton, 2000; Kawato, 1999; Wolpert & Ghahramani, 2000). The concept describes a mechanism that determines the discrepancy between two values, typically by subtracting one from the other; the computation assumes that the two values are of like kind, such that the results of the comparison are meaningful. Although the mismatch theory of the ERN has yet to be formalized in a computational model, these considerations suggest a comparison between the activation states associated with the motor system.

Because the monitor module in the RL-ERN theory can detect errors, it might be construed as a comparator. However, in contrast to the mechanism we have described, the monitor in this theory detects errors by identifying states of the system and assigning values to those states. The states can include features of both the internal and the external environments and combinations of those features. The monitor can detect response errors by recognizing that certain stimulus and response combinations are bad. This distinction between the mismatch theory and the RL-ERN theory is not trivial, because the theories make different predictions. For example, the monitor module in the RL-ERN theory can detect error feedback by recognizing these stimuli as being bad. In contrast, the monitor module in the mismatch theory cannot detect error feedback, because error detection in the theory depends entirely on comparisons of the internal states of the task module.

The mismatch theory has been criticized for an apparent inability to account for the timing of the ERN (Yeung et al., 2004). Specifically, any system that can identify errors during the process of error commission, as is the case with the mechanism that produces the ERN, must do so very rapidly, within a few moments of the onset of the error. This issue has not been addressed by the mismatch theory, nor was it addressed in the original RL-ERN simulations (Holroyd & Coles, 2002), and it raises the following question: Given that the monitor has detected the error and has therefore identified the target by the time the error is committed, why does the task module not use that information to avoid making the error in the first place? It seems that the monitor module—

whether it is the comparator in the mismatch theory or the value detector in the RL-ERN theory—enjoys privileged access to important task-related information.

However, the logic motivating that question is not correct. In fact, the participants are required to respond so fast that a certain percentage of impulsive errors is inevitable (15% in the case of this study), regardless of the ability of the monitor module to detect those errors. In speeded RT tasks, these errors occur because the response is produced before the response system has fully processed information related to the external stimulus. Thus, a more appropriate question might be the following: Given that the task module produces impulsive errors, how is the monitor module capable of detecting those errors so quickly?

The RL-ERN model answers this question in the following way. Optimally, once stimulus information becomes available to the response system, this information should override the influence of prior biases (Bogacz et al., 2005). It is intuitive that the response should be produced around the time of stimulus categorization: Responses that occur too soon or too late following stimulus categorization tend to be, respectively, too inaccurate or too slow. As we have described, this principle was implemented with inhibition of the attention–response units by the perceptual units, which suppressed the tendency of the system to produce impulsive responses after the stimulus had been evaluated. Thus, the simulated participants responded around the time of stimulus categorization, just quickly enough such that impulsive responses could occur even as the stimulus categorization process was nearing completion. As a result, the error detection system tended to receive stimulus- and response-related information at about the same time, during the response generation process or shortly thereafter. The ERN was hence produced at the time of error commission but too late to do anything about it.

#### *Conflict Monitoring Theory*

Another account of performance monitoring holds that the anterior cingulate cortex is sensitive to response conflict (Botvinick et al., 2001). The conflict monitoring theory is often described in opposition to the mismatch theory of the ERN, but in fact the two theories appear to share much in common at an algorithmic level (Yeung et al., 2004). According to the theory, the anterior cingulate cortex detects the simultaneous activation of incompatible response options, a position supported by an accumulating body of functional neuroimaging data (for a review, see Botvinick et al., 2001). The theory also proposes that the ERN reflects response conflict on error trials in the period immediately following the incorrect response and thus can serve to detect errors. This latter concept has recently been implemented in a computational model of the Eriksen flanker task by Yeung et al. (2004). In that study, the response conflict mechanism proposed by Botvinick et al. (2001) was embedded into the original connectionist model of the Eriksen flanker task (Servan-Schreiber, 1990; Servan-Schreiber et al., 1998). Among other things, Yeung et al. (2004) demonstrated that postresponse conflict on error trials accounted for the ERN amplitudes observed in the biased Eriksen flanker task (see Figure 7). They also showed that preresponse conflict on cor-

rect trials accounted for the amplitude of another component of the ERP, the N200 (see also Nieuwenhuis, Yeung, van den Wildenberg, & Ridderinkhof, 2003).

The conflict model of the ERN enjoys several advantages over the RL-ERN model, not least of which is its computational elegance. Whereas the RL-ERN model depends on a relatively complicated apparatus for predicting ERN amplitude (see Figure 11), the conflict model predicts ERN amplitude by simply multiplying together the activity of the two response units. Another strength of the conflict model is its generality, because it accounts for both N200 data and functional neuroimaging data regarding the anterior cingulate cortex, which are associated with conflict on correct trials, as well as the ERN data. Conversely, the RL-ERN model exhibits some strengths of its own. In particular, the RL-ERN model unifies in one theoretical framework the response and feedback ERNs (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002; see also Holroyd, Coles, & Nieuwenhuis, 2002), whereas the conflict monitoring hypothesis (like the mismatch hypothesis) has not yet accounted for the feedback ERN. Second, the RL-ERN model links that phenomenon to an extremely rich literature on learning theory (Kaelbling et al., 1996; Sutton & Barto, 1998) and to how those principles are implemented at the neural level. This connection holds out the promise that the ERN can be used to investigate the neural mechanisms underlying reinforcement learning in humans.

From this summary, it is apparent that the RL-ERN and conflict monitoring theories account for partially overlapping but also partially complementary sets of data. Both provide an account of error detection and its relation to the response ERN. The RL-ERN theory can also explain the feedback ERN but not conflict-related phenomena on correct trials (i.e., the N2 and the anterior cingulate cortex hemodynamic activity), whereas the reverse is true for the conflict monitoring theory. In view of

this, it is worth considering how the specific mechanisms proposed by each theory relate to one another. In the conflict theory, monitoring involves the detection of simultaneous activation of incompatible representations (technically, the computational energy of the representations involved in task processing; Botvinick et al., 2001). Most applications of this theory to date, in particular those addressing the ERN, have focused on response conflict. That is, they have focused on conflict between response representations, computed as the product of the activity of competing (i.e., mutually inhibitory) response units in the model. In the current model, however, there is no lateral inhibition between response units; thus, conflict cannot be measured in this layer alone (the computation of energy has no meaning for units that are not connected to one another). However, we can measure conflict by computing the energy of representations in the category and response layers together, as indicated in Figure 15. In this case, conflict occurs when there is coactivation of incompatible category and response units. Notice that when representations in the category and response layers are incompatible, the sign of the energy is positive, which can be used as a signal that an error has occurred, much as the conjunction units in the RL-ERN monitor's value layer indicate an error. Figure 15 shows the homology between the measurement of conflict and the value assigned to corresponding conjunctions of category and response layer states in the RL-ERN monitor.

From this perspective, it appears that, at least at an abstract level, the monitors in the conflict and RL-ERN theories share some important similarities. Both are responsive to the compatibility of representations activated in the category and response layers of the task module. However, the theories are not identical. In particular, there are potentially important differences in the dynamics of how the stimulus-response compatibility is

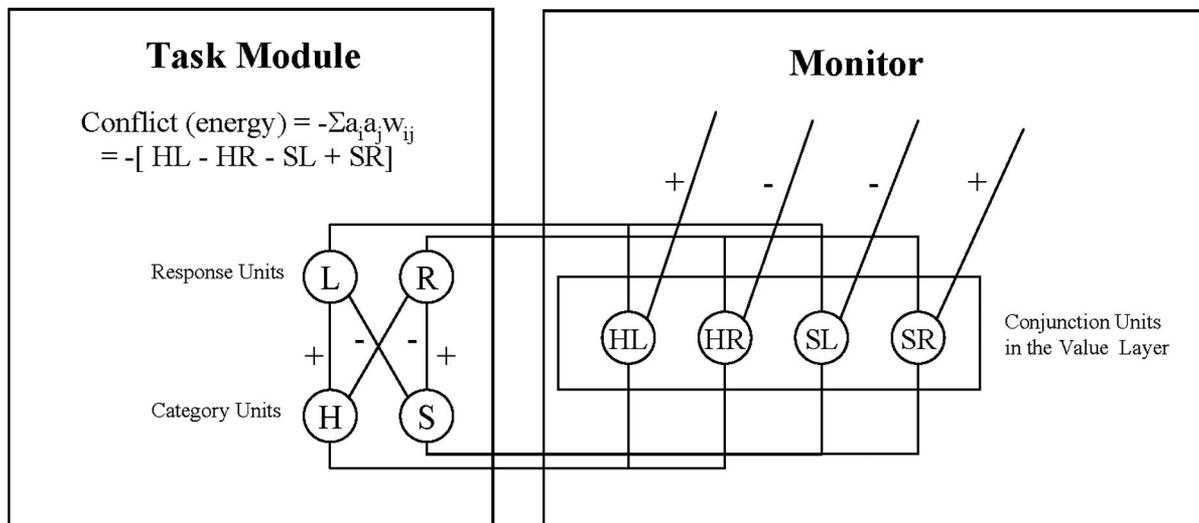


Figure 15. A comparison of the response conflict mechanism and the reinforcement learning of the error-related negativity mechanism for error detection. HL = H target-left response unit; HR = H target-right response unit; SL = S target-left response unit; SR = S target-right response unit; L = left response unit; R = right response unit; H = H target unit; S = S target unit.

computed. As we have noted, the conflict theory assumes that energy is computed instantaneously—that is, as the product of simultaneous activity among processing units. Conflict is high when two incompatible units are simultaneously active (e.g., *S* and *L*) but low if only one is active at a time (e.g., the activity of one precedes the other). In contrast, in the RL-ERN model, incompatibilities can be detected asynchronously. This capacity results from the fact that as soon as a category unit in the task module exceeds a critical threshold (e.g., *S*), the corresponding task state unit is immediately activated and remains so until the conclusion of the trial. This activity, coupled with later activation of the task state unit for an incompatible response (e.g., *L*), activates a conjunction unit associated with negative value (e.g., the *SL* unit), signaling an error.

Strictly construed, a conflict detection mechanism should not be sensitive to such asynchronous activation of incompatible representations. There is no direct evidence regarding whether the activation of stimulus category and response representations occurs asynchronously in the human brain. However, the present model shows that such an assumption, together with the mechanisms implemented in the RL-ERN model, is sufficient to account for both the behavioral and the electrophysiological data we have addressed. It remains a challenge for the conflict theory to demonstrate that, under alternative assumptions (e.g., greater temporal overlap of category and response representations), it can provide an equally satisfactory account of these data. In any event, it is encouraging that the formal specification of these theories now allows them to be compared in a precise and direct way, which promises to provide guidance on the design of empirical studies that may adjudicate between them.

The conflict and RL-ERN theories also differ in three other important respects. First, like the mismatch comparator theory, the conflict model cannot detect incorrect responses on the basis of other sources of error information, such as external feedback, without modification to the theory. Second, the mechanisms for ERN production are different in the two theories. In the conflict theory, the ERN is associated with the simultaneous activation of incompatible response channels. In the RL-ERN theory, the ERN is associated with a prediction error in reward. It is not yet clear how to reconcile these two positions, if possible. Third, the conflict theory proposes that conflict, and thereby errors, are detected by the anterior cingulate cortex, but the RL-ERN theory proposes that errors are detected by the basal ganglia. To the extent that the error detection process can be localized to either of these areas (e.g., with lesion studies), future research may be able to decide between these assertions of the two theories.

Thus, a subject for future research is the direct comparison of the two models. By instantiating the RL-ERN model in an architecture with continuous activation functions, these simulations indicate how response conflict (as determined by the simultaneous activation of the two response channels) and the TD signal (as determined by the change in value) could be studied in a single system. We suspect that the two theories may yield similar predictions in a wide variety of tasks. In addition, we expect that the two mechanisms are dissociable, providing complementary functions. For example, the reinforcement learning component of a combined conflict–TD model could apply feedback information for behavioral adaptation, whereas

the conflict component could direct response selection in situations with multiple, equipotent response options. In this event, it is an empirical matter to determine how and when the two components contribute to the ERN (see also Holroyd, 2004; Holroyd & Yeung, 2003).

### Conclusion

In this study we have presented an error detection mechanism that can (a) rapidly detect errors on the basis of information associated with both internal response processes and external stimuli, (b) produce error signals that can be used for behavioral control in speeded RT tasks, and (c) reproduce the amplitude and latency of an electrophysiological phenomena associated with performance monitoring—the response ERN. This error detection mechanism is implemented in a biologically plausible model that captures the dynamics of the stimulus categorization and response generation processes and that is consistent with the functional neuroanatomy of the associated neural systems (especially the anterior cingulate cortex, the basal ganglia, and the midbrain dopamine system). Furthermore, in other work, we have shown that this system can (a) detect errors on the basis of external performance feedback, (b) use the external feedback to learn to detect errors in the absence of feedback, (c) provide error information to a behavioral system that can use it to learn appropriate stimulus–response mappings and/or to regulate the response production process, and (d) reproduce the amplitude and latency of a second electrophysiological phenomenon associated with performance monitoring—the feedback ERN (Holroyd & Coles, 2002; see also Holroyd et al., 2005). When taken together, this work presents an initial step toward the development of a theoretical foundation for a cognitive neuroscience of error processing. An important aspect of future work is to compare and contrast this theoretical framework with comparable theories of error processing and the anterior cingulate cortex (e.g., Botvinick et al., 2001; Yeung et al., 2004).

### References

- Angel, R. W. (1976). Efference copy in the control of movement. *Neurology*, *26*, 1164–1168.
- Bao, S., Chan, V. T., & Merzenich, M. M. (2001, July 5). Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature*, *412*, 79–83.
- Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 215–232). Cambridge, MA: MIT Press.
- Bates, A. T., Kiehl, K. A., Laurens, K. R., & Liddle, P. F. (2002). Error-related negativity and correct response negativity in schizophrenia. *Clinical Neurophysiology*, *113*, 1454–1463.
- Bogacz, R., Brown, E. T., Moehlis, J. M., Hu, P., Holmes, P., & Cohen, J. D. (2005). *The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks*. Manuscript submitted for publication.
- Botvinick, M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.
- Braver, T. S., Barch, D. M., Keys, B. A., Carter, C. S., Cohen, J. D., Kaye, J. A., et al. (2001). Context processing in older adults: Evidence for a

- theory relating to cognitive control in healthy aging. *Journal of Experimental Psychology: General*, *130*, 746–763.
- Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience*, *19*, 10502–10511.
- Brown, L. L., Schneider, J. S., & Lidsky, T. I. (1997). Sensory and cognitive functions of the basal ganglia. *Current Opinion in Neurobiology*, *7*, 157–163.
- Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., & Rosen, B. R. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proceedings of the National Academy of Sciences, USA*, *99*, 523–528.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processing: A parallel distributed account of the Stroop effect. *Psychological Review*, *97*, 332–361.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, *99*, 45–77.
- Cohen, J. D., Servan-Schreiber, D., & McClelland, J. L. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology*, *105*, 239–269.
- Coles, M. G. H., De Jong, R., Gehring, W. J., & Gratton, G. (1991). Continuous versus discrete information processing: Evidence from movement-related potentials. In C. H. M. Brunia, G. Mulder, & M. N. Verbaten (Eds.), *Event-related brain research* (EEG Suppl. 42, pp. 260–269). Amsterdam: Elsevier.
- Coles, M. G. H., Gratton, G., Bashore, T. R., Eriksen, C. W., & Donchin, E. (1985). A psychophysical investigation of the continuous flow model of human information processing. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 529–553.
- Coles, M. G. H., & Rugg, M. D. (1995). Event-related brain potentials: An introduction. In M. D. Rugg & M. G. H. Coles (Eds.), *Electrophysiology of mind* (pp. 1–26). Oxford, England: Oxford University Press.
- Coles, M. G. H., Scheffers, M. K., & Fournier, L. (1995). Where did you go wrong? Errors, partial errors, and the nature of human information processing. *Acta Psychologica*, *90*, 129–144.
- Coles, M. G. H., Scheffers, M. K., & Holroyd, C. B. (2001). Why is there an ERN/Ne on correct trials? Response representations, stimulus-related components, and the theory of error-processing. *Biological Psychology*, *56*, 191–206.
- Contreras-Vidal, J. L., & Schultz, W. (1999). A predictive reinforcement model of dopamine neurons for learning approach behavior. *Journal of Computational Neuroscience*, *6*, 191–214.
- Crino, P. B., Morrison, J. H., & Hof, P. R. (1993). Monoaminergic innervation of cingulate cortex. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 285–310). Boston: Birkhauser.
- Cummings, J. L. (1993). Frontal-subcortical circuits and human behavior. *Archives of Neurology*, *50*, 873–880.
- Davis, K. L., Kahn, R. S., Ko, G., & Davidson, M. (1991). Dopamine in schizophrenia: A review and reconceptualization. *American Journal of Psychiatry*, *148*, 1474–1486.
- Daw, N. D., Kakade, S., & Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Networks*, *15*, 603–616.
- Daw, N. D., & Touretzky, D. S. (2002). Long-term reward prediction in TD models of the dopamine system. *Neural Computation*, *14*, 2567–2583.
- De Bruijn, E. R. A., Hulstijn, W., Verkes, R. J., Ruijt, G. S. F., & Sabbe, B. G. C. (2004). Drug-induced stimulation and suppression of action monitoring. *Psychopharmacology*, *177*, 151–160.
- Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, *5*, 303–305.
- Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, *4*, 423–431.
- Devinsky, O. (1983). Neuroanatomy of Gilles de la Tourette's syndrome. *Annals of Neurology*, *40*, 508–514.
- Devinsky, O., Morrell, M. J., & Vogt, B. A. (1995). Contributions of anterior cingulate cortex to behaviour. *Brain*, *118*, 279–306.
- Diggles, V. A. (1987). Rapid error corrections: Behavioral evidence for central monitoring of efference. *Journal of Human Movement Studies*, *13*, 433–460.
- Donchin, E. (1984). Dissociation between electrophysiology and behavior—A disaster or a challenge? In E. Donchin (Ed.), *Cognitive psychophysiology* (pp. 107–118). Hillsdale, NJ: Erlbaum.
- Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, *11*, 355–372.
- Donchin, E., Gratton, G., Dupree, D., & Coles, M. (1988). After a rash action: Latency and amplitude of the P300 following fast guesses. In G. C. Galbraith, M. L. Kietzman, & E. Donchin (Eds.), *Neurophysiology and psychophysiology: Experimental and clinical applications* (pp. 173–188). Hillsdale, NJ: Erlbaum.
- Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? *Neural Networks*, *12*, 961–974.
- Doya, K. (2000). Reinforcement learning in continuous time and space. *Neural Computation*, *12*, 219–245.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, *15*, 495–506.
- Dreher, J.-C., & Burnod, Y. (2002). An integrative theory of the phasic and tonic modes of dopamine modulation in the prefrontal cortex. *Neural Networks*, *15*, 583–602.
- Dreher, J.-C., Guigon, E., & Burnod, Y. (2002). A model of prefrontal cortex dopaminergic modulation during the delayed alternation task. *Journal of Cognitive Neuroscience*, *14*, 853–865.
- Dum, R. P., & Strick, P. L. (1993). Cingulate motor areas. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 415–441). Boston: Birkhauser.
- Duncan-Johnson, C. C., & Donchin, E. (1982). The P300 component of the event-related brain potential as an index of information processing. *Biological Psychology*, *24*, 1–52.
- Egelman, D. M., Person, C., & Montague, P. R. (1998). A computational role for dopamine delivery in human decision-making. *Journal of Cognitive Neuroscience*, *10*, 623–630.
- Elliott, R., & Dolan, R. J. (1998). Activation of different anterior cingulate foci in association with hypothesis testing and response selection. *Neuroimage*, *8*, 17–29.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*, 143–149.
- Eriksen, C. W., Coles, M. G. H., Morris, L. R., & O'Hara, W. P. (1985). An electromyographic examination of response competition. *Bulletin of the Psychonomic Society*, *23*, 165–168.
- Eriksen, C. W., & Schultz, D. W. (1979). Information processing in visual search: A continuous flow conception and experimental results. *Perception & Psychophysics*, *25*, 249–263.
- Falkenstein, M., Hielscher, H., Dziobek, I., Schwarzenau, P., Hoormann, J., Sundermann, B., & Hohnsbein, J. (2001). Action monitoring, error detection, and the basal ganglia: An ERP study. *NeuroReport*, *12*, 157–161.
- Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1990). Effects of errors in choice reaction tasks on the ERP under focused and divided attention. In C. Brunia, A. Gaillard, & A. Kok (Eds.), *Psychophysiological brain research* (pp. 192–195). Tilburg, the Netherlands: Tilburg University Press.
- Falkenstein, M., Hoormann, J., Christ, S., & Hohnsbein, J. (2000). ERP

- components on reaction errors and their functional significance: A tutorial. *Biological Psychology*, 51, 87–107.
- Farwell, L. A., Martinerie, J. M., Bashore, T. R., Rapp, P. E., & Goddard, P. H. (1993). Optimal digital filters for long-latency components of the event-related brain potential. *Psychophysiology*, 30, 306–315.
- Ferron, A., Thierry, A. M., Le Douarin, C., & Glowinski, J. (1984). Inhibitory influence of the mesocortical dopaminergic system on spontaneous activity or excitatory response induced from the thalamic mediodorsal nucleus in the rat medial prefrontal cortex. *Brain Research*, 302, 257–265.
- Fournier, L. R., Scheffers, M. K., Coles, M. G. H., & Adamson, A. (1997). The dimensionality of the flanker compatibility effect: A psychophysiological analysis. *Psychological Research*, 60, 144–155.
- Gabriel, M. (1993). Discriminative avoidance learning: A model system. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 478–523). Boston: Birkhauser.
- Gao, W.-J., Krimer, L. S., & Goldman-Rakic, P. S. (2001). Presynaptic regulation of recurrent excitation by D1 receptors in prefrontal circuits. *Proceedings of the National Academy of Sciences, USA*, 98, 295–300.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science*, 4, 385–390.
- Gehring, W. J., & Willoughby, A. R. (2002, March 22). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science*, 295, 2279–2282.
- Gemba, H., Sasaki, K., & Brooks, V. B. (1986). Error potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neuroscience Letters*, 70, 223–227.
- Gobbel, J. R. (1995). A biophysically-based model of the neostriatum as a dynamically reconfigurable network. In M. Boden & L.-E. Niklasson (Eds.), *Current trends in connectionism: Proceedings of the second Swedish conference on connectionism, Skovde, Sweden*. Hillsdale, NJ: Erlbaum.
- Gratton, G., Coles, M. G. H., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, 55, 468–484.
- Gratton, G., Coles, M. G. H., Sirevaag, E. J., Eriksen, C. W., & Donchin, E. (1988). Pre- and poststimulus activation of response channels: A psychophysiological analysis. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 331–344.
- Grossberg, S. (1978). A theory of visual coding, memory, and development. In E. L. J. Leeuwenberg & H. F. J. M. Buffart (Eds.), *Formal theories of visual perception* (pp. 7–26). New York: Wiley.
- Gurden, H., Takita, M., & Jay, T. M. (2000, November 15). Essential role of D1 but not D2 receptors in the NMDA receptor-dependent long-term potentiation at hippocampal-prefrontal cortex synapses *in vivo*. *Journal of Neurosciences*, 20, RC106.
- Hajcak, G., Vidal, F., & Simons, R. F. (2004). Difficulties with easy tasks: ERN/Ne and stimulus overlap. In N. M. Ullsperger & M. Falkenstein (Eds.), *Errors, conflicts, and the brain: Current opinions on performance monitoring* (pp. 204–211). Leipzig, Germany: MPI of Cognitive Neuroscience.
- Harrison, P. J. (2000, September 16). Dopamine and schizophrenia—proof at last? *Lancet*, 356, 958–959.
- Higgins, J. R., & Angel, R. W. (1970). Correction of tracking errors without sensory feedback. *Journal of Experimental Psychology*, 84, 412–416.
- Holroyd, C. B. (2004). A note on the N200 and the feedback ERN. In M. Ullsperger & M. Falkenstein (Eds.), *Errors, conflicts, and the brain: Current opinions on performance monitoring* (pp. 211–218). Leipzig, Germany: MPI of Cognitive Neuroscience.
- 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.
- Holroyd, C. B., Coles, M. G. H., & Nieuwenhuis, S. (2002, May 31). Medial prefrontal cortex and error potentials. *Science*, 296, 1610–1611.
- Holroyd, C. B., Dien, J., & Coles, M. G. H. (1998). Error-related scalp potentials elicited by hand and foot movements: Evidence for an output-independent error-processing system in humans. *Neuroscience Letters*, 242, 65–68.
- Holroyd, C. B., Larsen, J. T., & Cohen, J. D. (2004). Context dependence of the event-related brain potential associated with reward and punishment. *Psychophysiology*, 41, 245–253.
- Holroyd, C. B., Nieuwenhuis, S., Mars, R., & Coles, M. G. H. (2004). Anterior cingulate cortex, selection for action, and error processing. In M. Posner (Ed.), *Cognitive neuroscience of attention* (pp. 219–231). New York: Guilford Press.
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (2003). Reward prediction errors are reflected in the event-related brain potential. *NeuroReport*, 18, 2481–2484.
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R., Coles, M. G. H., & Cohen, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience*, 7, 497–498.
- Holroyd, C. B., Praamstra, P., Plat, E., & Coles, M. G. H. (2002). Spared error-related potentials in mild to moderate Parkinson's disease. *Neuropsychologia*, 40, 2116–2124.
- Holroyd, C. B., & Yeung, N. (2003). Alcohol and error processing. *Trends in Neurosciences*, 26, 402–404.
- Holroyd, C. B., Yeung, N., Coles, M. G. H., & Cohen, J. D. (2005). *Internalizing representations of correct behavior from feedback*. Manuscript in preparation.
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4, 251–257.
- Houk, J. C. (1995). Information processing in modular circuits linking basal ganglia and cerebral cortex. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 3–9). Cambridge, MA: MIT Press.
- Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 249–270). Cambridge, MA: MIT Press.
- Ito, S., Stuphorn, V., Brown, J. W., & Schall, J. D. (2003, October 3). Performance monitoring by anterior cingulate cortex during saccade countermanding. *Science*, 302, 120–122.
- Jasper, H. H. (1958). The ten twenty electrode system of the international federation. *Electroencephalography and Clinical Neurophysiology*, 10, 371–375.
- Johannes, S., Wieringa, B. M., Nager, W., Müller-Vahl, K. R., Dengler, R., & Münte, T. F. (2002). Excessive action monitoring in Tourette syndrome. *Journal of Neurology*, 249, 961–966.
- Johnson, R., Jr. (1988). The amplitude of the P300 component of the event-related potential: Review and synthesis. *Advances in Psychology*, 3, 69–137.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718–727.
- Kermadi, I., & Joseph, J. P. (1995). Activity in the caudate nucleus of monkey during spatial sequencing. *Journal of Neurophysiology*, 74, 911–933.
- Keselman, H. J., & Rogan, J. C. (1980). Repeated measures *F* tests and psychophysiological research: Controlling the number of false positives. *Psychophysiology*, 17, 499–503.
- Kimura, M., & Matsumoto, N. (1997). Nigrostriatal dopamine system may

- contribute to behavioral learning through providing reinforcement signals to the striatum. *European Neurology*, 38(Suppl. 1), 11–17.
- Kopp, B., & Rist, F. (1999). An event-related brain potential substrate of disturbed response monitoring in paranoid schizophrenic patients. *Journal of Abnormal Psychology*, 108, 337–346.
- Kutas, M., McCarthy, G., & Donchin, E. (1977, August 19). Augmenting mental chronometry: The P300 as a measure of stimulus evaluation time. *Science*, 197, 792–795.
- Laming, D. R. J. (1979). Choice reaction performance following an error. *Acta Psychologica*, 43, 199–224.
- Leuthold, H., & Sommer, W. (1998). Postperceptual effects and P300 latency. *Psychophysiology*, 35, 34–46.
- Levine, T. R., & Hullett, C. R. (2002). Eta squared, partial eta squared, and misreporting of effect size in communication research. *Human Communication Research*, 28, 612–625.
- Lewis, B. L., & O'Donnell, P. (2000). Ventral tegmental area afferents to the prefrontal cortex maintain membrane potential up states in pyramidal neurons via D1 dopamine receptors. *Cerebral Cortex*, 10, 1168–1175.
- Lippold, O. C. J. (1967). Electromyography. In P. H. Venables & I. Martin (Eds.), *A manual of psychophysiological methods* (pp. 245–297). Amsterdam: North-Holland.
- Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, 108, 393–434.
- Luria, A. R. (1973). *The working brain*. New York: Basic Books.
- Mars, R. B., de Bruijn, E. R. A., Hulstijn, W., Miltner, W. H. R., & Coles, M. G. H. (2004). What if I told you: “You were wrong”? Brain potentials and behavioral adjustments elicited by feedback in a time-estimation task. In M. Ullsperger & M. Falkenstein (Eds.), *Errors, conflicts, and the brain: Current opinions on performance monitoring* (pp. 129–134). Leipzig, Germany: MPI of Cognitive Neuroscience.
- Masaki, H., Takasawa, N., & Yamazaki, K. (2000). An electrophysiological study of the locus of the interference effect in a stimulus-response compatibility paradigm. *Psychophysiology*, 37, 464–472.
- Massaro, D. W. (1988). Some criticisms of connectionist models of human performance. *Journal of Learning and Memory*, 27, 213–234.
- Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology*, 57, 203–220.
- Mathalon, D. H., Fedor, M., Faustman, W. O., Gray, M., Askari, N., & Ford, J. M. (2002). Response-monitoring dysfunction in schizophrenia: An event-related brain potential study. *Journal of Abnormal Psychology*, 111, 22–41.
- McCarthy, G. (1984). Stimulus evaluation time and P300 latency. In E. Donchin (Ed.), *Cognitive psychophysiology: Event-related potentials and the study of cognition* (Vol. 1, pp. 254–287). Hillsdale, NJ: Erlbaum.
- McCarthy, G., & Donchin, E. (1981, January 2). A metric for thought: A comparison of P300 latency and reaction time. *Science*, 211, 77–80.
- McClelland, J. L. (1992). Toward a theory of information processing in graded, random, and interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience—A Silver Jubilee volume* (pp. 655–688). Cambridge, MA: MIT Press.
- McClelland, J. L., & Rumelhart, D. E. (1988). Interactive activation and competition. In J. L. McClelland & D. E. Rumelhart (Eds.), *Explorations in parallel distributed processing: A handbook of models, programs, and exercises* (pp. 11–47). Cambridge, MA: MIT Press.
- McClure, S. M., Daw, N. D., & Montague, P. R. (2003). A computational substrate for incentive salience. *Trends in Neurosciences*, 26, 423–428.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.
- Miller, J. (1988). Discrete and continuous models of human information processing: Theoretical distinctions and empirical results. *Acta Psychologica*, 67, 191–257.
- Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a “generic” neural system for error detection. *Journal of Cognitive Neuroscience*, 9, 788–798.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, 16, 1936–1947.
- Nakahara, H., Doya, K., & Hikosaka, O. (2001). Parallel cortical-basal ganglia mechanisms for acquisition and execution of visuomotor sequences—A computational approach. *Journal of Cognitive Neuroscience*, 13, 626–647.
- Nieuwenhuis, S., Holroyd, C. B., Mol, N., & Coles, M. G. H. (2004). Reinforcement-related brain potentials from medial frontal cortex: Origins and functional significance. *Neuroscience & Biobehavioral Reviews*, 28, 441–448.
- Nieuwenhuis, S., Ridderinkhof, K. R., Talsma, D., Coles, M. G. H., Holroyd, C. B., Kok, A., & Van der Molen, M. W. (2002). A computational account of altered error processing in older age: Dopamine and the error-related negativity. *Cognitive, Affective, and Behavioral Neuroscience*, 2, 19–36.
- Nieuwenhuis, S., Yeung, N., Holroyd, C. B., Schurger, A., & Cohen, J. D. (2004). Sensitivity of electrophysiological activity from medial frontal cortex to utilitarian and performance feedback. *Cerebral Cortex*, 14, 741–747.
- Nieuwenhuis, S., Yeung, N., van den Wildenberg, W., & Ridderinkhof, K. R. (2003). Electrophysiological correlates of anterior cingulate function in a go/nogo task: Effects of response conflict and trial-type frequency. *Cognitive, Affective, and Behavioral Neuroscience*, 3, 17–26.
- Niki, H., & Watanabe, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Research*, 171, 213–224.
- Ohlsson, S. (1996). Learning from performance errors. *Psychological Review*, 103, 241–262.
- O'Reilly, R. C., & Farah, M. J. (1999). Simulation and explanation in neuropsychology and beyond. *Cognitive Neuropsychology*, 16, 49–72.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience*. Cambridge, MA: MIT Press.
- Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive, and cognition interface. *Nature Reviews Neuroscience*, 2, 417–424.
- Picard, N., & Strick, P. L. (1996). Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex*, 6, 342–353.
- Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson, R., Jr., et al. (2000). Guidelines for using human event-related potentials to study cognition: Recording standards and publication criteria. *Psychophysiology*, 37, 127–152.
- Porrino, L. J. (1993). Cortical mechanisms of reinforcement. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 445–460). Boston: Birkhauser.
- Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–423). Cambridge, MA: MIT Press.
- Procyk, E., Tanaka, Y. L., & Joseph, J. P. (2000). Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience*, 3, 502–508.
- Rabbitt, P. M. A. (1966a, October 22). Error correction time without external error signals. *Nature*, 212, 438.
- Rabbitt, P. M. A. (1966b). Errors and error corrections in choice-response tasks. *Journal of Experimental Psychology*, 71, 264–272.
- Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, 86, 446–461.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical*

- conditioning II: Current research and theory (pp. 64–99). New York: Appleton-Century-Crofts.
- Reynolds, J. N. J., Hyland, B. I., & Wickens, J. R. (2001, September 6). A cellular mechanism of reward-related learning. *Nature*, *413*, 67–70.
- Reynolds, J. N. J., & Wickens, J. R. (2002). Dopamine-dependent plasticity of corticostriatal synapses. *Neural Networks*, *15*, 507–521.
- Richardson, N. R., & Gratton, A. (1998). Changes in medial prefrontal cortical dopamine levels associated with response-contingent food reward: An electrochemical study in rat. *Journal of Neuroscience*, *18*, 9130–9138.
- Ridderinkhof, K. R., de Vlugt, Y., Bramlage, A., Spaan, M., Elton, M., Snel, J., & Band, G. P. H. (2002, December 13). Alcohol consumption impairs detection of performance errors in mediofrontal cortex. *Science*, *298*, 2209–2211.
- Rougier, N. P., & O'Reilly, R. C. (2002). Learning representations in a gated prefrontal cortex model of dynamic task switching. *Cognitive Science*, *26*, 503–520.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1). Cambridge, MA: MIT Press.
- Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society of London B, Biological Sciences*, *357*, 1073–1082.
- Schall, J. D. (2003). Neural correlates of decision processes: Neural and mental chronometry. *Current Opinion in Neurobiology*, *13*, 182–186.
- Schall, J. D., Stuphorn, V., & Brown, J. (2002). Monitoring and control of action by the frontal lobes. *Neuron*, *36*, 309–322.
- Schall, J. D., & Thompson, K. G. (1999). Neural selection and control of visually guided eye movements. *Annual Review of Neuroscience*, *22*, 241–259.
- Scheffers, M. K., & Coles, M. G. H. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 141–151.
- Schultz, W. (1995). The primate basal ganglia between the intention and outcome of action. In M. Kimura & A. M. Graybiel (Eds.), *Functions of the cortical-basal ganglia loop* (pp. 31–48). Tokyo: Springer-Verlag.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*, 241–263.
- Schultz, W., Apicella, P., Romo, R., & Scarnati, E. (1995). Context-dependent activity in primate striatum reflecting past and future behavioral events. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 11–27). Cambridge, MA: MIT Press.
- Schultz, W., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Seidenberg, M. S. (1993). Connectionist models and cognitive theory. *Psychological Science*, *4*, 228–235.
- Servan-Schreiber, D. (1990). *From physiology to behavior: Computational models of catecholamine modulation of information processing*. Unpublished doctoral dissertation, Carnegie Mellon University.
- Servan-Schreiber, D., Bruno, R. M., Carter, C. S., & Cohen, J. D. (1998). Dopamine and the mechanisms of cognition: Part I. A neural network model predicting dopamine effects on selective attention. *Biological Psychiatry*, *43*, 713–722.
- Shidara, M., & Richmond, B. J. (2002, May 31). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, *296*, 1709–1711.
- Shima, K., & Tanji, J. (1998, November 13). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, *282*, 1335–1338.
- Singer, H. S., Butler, I. J., Tune, L. E., Seifert, W. E., & Coyle, J. T. (1982). Dopaminergic dysfunction in Tourette syndrome. *Annals of Neurology*, *12*, 361–366.
- Spencer, K. M., & Coles, M. G. H. (1999). The lateralized readiness potential: Relationships between human data and response activation in a connectionist model. *Psychophysiology*, *36*, 364–370.
- Spencer, K. M., Vila Abad, E., & Donchin, E. (2000). On the search for the neurophysiological manifestation of recollective experience. *Psychophysiology*, *37*, 494–506.
- Sporns, O., & Alexander, W. H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks*, *15*, 761–774.
- Stuphorn, V., Taylor, T. L., & Schall, J. D. (2000, December 14). Performance monitoring by the supplementary eye field. *Nature*, *408*, 857–860.
- Stuss, D. T., & Knight, R. T. (2002). *Principles of frontal lobe function*. Oxford, England: Oxford University Press.
- Suri, R. E. (2002). TD models of reward predictive responses in dopamine neurons. *Neural Networks*, *15*, 523–533.
- Suri, R. E., & Schultz, W. (2001). Temporal difference model reproduces anticipatory neural activity. *Neural Computation*, *13*, 841–862.
- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9–44.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: MIT Press.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thorpe, S., Fize, D., & Marlot, C. (1996, June 6). Speed of processing in the human visual system. *Nature*, *381*, 520–522.
- Tzschentke, T. M. (2000). The medial prefrontal cortex as part of the brain reward system. *Amino Acids*, *19*, 211–219.
- Ullsperger, M., & Von Cramon, D. Y. (2003). Error processing using external feedback: Specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by fMRI. *Journal of Neuroscience*, *23*, 4308–4314.
- Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999, January 22). The roles of locus coeruleus in the regulation of cognitive performance. *Science*, *283*, 549–554.
- Verleger, R. (1997). On the utility of P3 latency as an index of mental chronometry. *Psychophysiology*, *34*, 131–156.
- Vickers, D. (1978). An adaptive module for simple judgment. In J. Requin (Ed.), *Attention and performance* (Vol. 7, pp. 599–618). Hillsdale, NJ: Erlbaum.
- Vogt, B. A., Vogt, L. J., Nimchinsky, E. A., & Hof, P. R. (1997). Primate cingulate cortex chemoarchitecture and its disruption in Alzheimer's disease. In F. E. Bloom, A. Bjorklund, & T. Hokfelt (Eds.), *Handbook of chemical neuroanatomy: Vol. 13. The primate nervous system, Part I* (pp. 455–528). Amsterdam: Elsevier.
- Waelti, P., Dickinson, A., & Schultz, W. (2001, July 5). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*, 43–48.
- Wilson, C. J. (1995). The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 29–50). Cambridge, MA: MIT Press.
- Winer, B. J., Brown, D. R., & Michels, K. M. (1991). *Statistical principles in experimental design*. New York: McGraw-Hill.
- Wise, S. P., Murray, E. A., & Gerfen, C. R. (1996). The frontal cortex-basal ganglia system in primates. *Critical Reviews in Neurobiology*, *10*, 317–356.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, *3*, 1212–1217.
- Yang, C. R., & Seamans, J. K. (1996). Dopamine D1 receptor actions in layers

V-V1 rat prefrontal cortex neurons in vitro: Modulation of dendritic-somatic signal integration. *Journal of Neuroscience*, 16, 1922–1935.

Yeung, N., Botvinick, 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., & Sanfey, A. (2004). Independent coding of reward magnitude and valence in the human brain. *Journal of Neuroscience*, 24, 6258–6264.

Zhang, H., Zhang, J., & Kornblum, S. (1999). A parallel distributed processing model of stimulus-stimulus and stimulus-response compatibility. *Cognitive Psychology*, 38, 386–432.

## Appendix A

### Recording and Data Analysis Procedures

#### Electrophysiological Recording

The electroencephalogram (EEG) was recorded with disposable Ag/AgCl electrodes. Scalp electrodes were placed according to the 10–20 system (Jasper, 1958) in locations Fz, Cz, Pz, and Oz. Additionally, electrodes were placed at C3' and C4', which are defined as situated 4 cm to either side of the midline at Cz. An electrode was placed on the right mastoid, and a ground electrode was placed on the forehead. These electrodes were referenced to the left mastoid, and their impedances were less than 5 k $\Omega$ . The vertical and horizontal electrooculograms (EOGs) were collected with bipolar-referenced electrodes placed above and below the right eye and on the outer canthi of the left and right eye, respectively; EOG impedances were less than 10 k $\Omega$ . Both EEG and EOG signals were amplified (Model 7P122, Grass Instruments, Quincy, MA) and filtered through a passband of 0.02–35.00 Hz (3 dB octave roll off). The electromyogram (EMG) was recorded with electrodes placed on the flexors of both arms (Lippold, 1967); impedances were under 30 k $\Omega$ . EMG signals were rectified (preamplifiers, Model 7P3B, Grass Instruments; 0.5 amplitude low-frequency cut-off at 0.3 Hz) and integrated (full-wave rectification and a time constant of 0.05 s). The EEG, EMG, and EOG signals were sampled at 200.0 Hz. Experimental control software developed at the Illinois laboratory was used for data acquisition and stimulus display.

#### Data Analysis

Response time (RT) and accuracy were determined offline from EMG onset with a computer algorithm described elsewhere (Holroyd, Praamstra, et al., 2002). Response accuracies and average RTs for correct and incorrect responses were computed for each participant and each stimulus condition. In addition, RT histograms (10-ms bins) were constructed for correct and incorrect responses, separately for each of the four stimulus conditions.

We rereferenced the EEG to linked mastoid electrodes offline by subtracting, for each sample on each trial, half of the activity recorded at the right mastoid from the activity recorded at each scalp electrode. Ocular artifact was removed from the EEG with the eye movement correction procedure described in Gratton, Coles, and Donchin (1983). For each electrode, we baseline corrected the EEG by subtracting from each data sample the average activity of that electrode during a 200-ms baseline period preceding stimulus onset. The single trial data were filtered with a 1–10 Hz (error-related negativity; ERN) and 0–6 Hz (P300; cf. Farwell, Martinerie, Bashore, Rapp, & Goddard, 1993; Fournier et al., 1997; Scheffers & Coles, 2000) passband via the Interactive Data Language digital filter algorithm.

We derived stimulus-locked and response-locked event-related brain potential (ERP) averages for each participant, channel, and condition by averaging the data across trials according to stimulus onset and response onset, respectively. For display, we averaged the ERPs across participants. The latency of the P300 was determined from the single trial ERPs recorded at Pz (Kutas et al., 1977; Picton et al., 2000; Spencer, Vila Abad, & Donchin, 2000); for each trial, the algorithm identified the latency of the most positive value between 280 ms and 1,300 ms following stimulus onset. ERN amplitude was determined from the response-locked ERP averages for each participant with an algorithm that identified the peak negativity recorded at channel Cz between 0 and 200 ms following the response. This algorithm also identified the onset of the component and computed its base-to-peak magnitude.

All statistical comparisons involved paired *t* tests or analysis of variance designs with repeated measures; effect sizes were evaluated with  $\eta_p^2$  (Levine & Hullett, 2002). The Greenhouse–Geisser correction for repeated measures (Keselman & Rogan, 1980) and the arc-sine transformation (Winer, Brown, & Michels, 1991) were applied where appropriate. Error bars in graphical representations represent within-subject confidence intervals (Masson & Loftus, 2003).

## Appendix B

### Task Module

Each trial of the simulation consisted of 150 cycles, each cycle corresponding to 10 ms. On each cycle, the output  $a_j$  of each unit  $j$  was thresholded such that activations less than zero were set equal to zero. The net input to each unit  $i$  was then given by

$$\text{net}_i = \sum_j w_{ij} a_j + 0.28 \text{ extinput}_i + \text{noise}, \quad (\text{B1})$$

where  $w_{ij}$  was the strength of the weight of the connection from unit  $j$  to unit  $i$  (see Table B1); noise was a random value, sampled on

each cycle for each unit, from a Gaussian distribution with a mean of 0.00 and a standard deviation of 0.05; and  $\text{extinput}_i$  was the external input to unit  $i$ . The external input to the attention–response units was 1.0 for cycles greater than 20 and was 0.0 otherwise. The external input to all of the perception units was  $-0.2$  for all cycles (tonic inhibition) plus, for the perception units activated on that trial, 0.8 beginning at cycle  $\sigma$ , where  $\sigma$  was a random value from a Gaussian distribution with a mean of 37.00 and a standard deviation of 5.00. The external input to all of the other units was zero for all cycles.

(Appendix continues)

Table B1

Task Module Weight Matrix

Weight	HF <sub>1</sub>	HF <sub>2</sub>	HT	HF <sub>3</sub>	HF <sub>4</sub>	SF <sub>1</sub>	SF <sub>2</sub>	ST	SF <sub>3</sub>	SF <sub>4</sub>	AC	AL	AR	HC	SC	LC	RC
1. HF <sub>1</sub>						$\gamma$					$9\gamma$						
2. HF <sub>2</sub>							$\gamma$				$9\gamma$						
3. HT								$\gamma$			$2\gamma$						
4. HF <sub>3</sub>									$\gamma$		$9\gamma$						
5. HF <sub>4</sub>										$\gamma$	$9\gamma$						
6. SF <sub>1</sub>	$\gamma$										$9\gamma$						
7. SF <sub>2</sub>		$\gamma$									$9\gamma$						
8. ST			$\gamma$								$2\gamma$						
9. SF <sub>3</sub>				$\gamma$							$9\gamma$						
10. SF <sub>4</sub>					$\gamma$						$9\gamma$						
11. AC	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$	$2.5\alpha$							
12. AL	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$							
13. AR	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$	$7\gamma$							
14. HC	$4\alpha$	$4\alpha$	$4\alpha$	$4\alpha$	$4\alpha$	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$							
15. SC	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$	$0.9\gamma$	$4\alpha$	$4\alpha$	$4\alpha$	$4\alpha$	$4\alpha$							
16. LC																	
17. RC														$1.4\alpha$	$\beta$	$-\beta$	$\beta$

*Note.* Cells indicate weights from units  $j$  (columns) to units  $i$  (rows). Shown is the case in which the  $H$  target is mapped to the left response, the  $S$  target is mapped to the right response, and the  $S$  target is frequent. Weights:  $\alpha = 0.08$ ;  $\gamma = -0.12$ ;  $\beta = 3.45\alpha$ ; empty cells = 0. HF<sub>1</sub>, HF<sub>2</sub>, HF<sub>3</sub>, and HF<sub>4</sub> are  $H$  noise perception units; SF<sub>1</sub>, SF<sub>2</sub>, SF<sub>3</sub>, and SF<sub>4</sub> are  $S$  noise perception units; HT =  $H$  target perception unit; ST =  $S$  target perception unit; AC = attention-perception unit; AL = left attention-response unit; AR = right attention-response unit; HC =  $H$  category unit; SC =  $S$  category unit; LC = left response unit; RC = right response unit.

The change in output of each unit was governed by the following equations (McClelland & Rumelhart, 1988):

$$\begin{aligned} &\text{If } (\text{net}_i > 0), \\ &\Delta a_i = (1.0 - a_i)\text{net}_i - 0.1(a_i + 0.1). \\ &\text{Otherwise,} \\ &\Delta a_i = (a_i + 0.2)\text{net}_i - 0.1(a_i + 0.1). \end{aligned} \quad (\text{B2})$$

The output of each unit was then updated as

$$a_i = a_i + \Delta a_i. \quad (\text{B3})$$

Outputs greater than 1.0 were set equal to 1.0, and outputs less than  $-0.2$  were set equal to  $-0.2$ .

The responding hand was determined by the first cycle (RT') at which the activity of either response unit crossed a threshold (0.42) between

$25 - \theta$  cycles and the end of the trial. The parameter  $\theta$  accounted for the period of activity associated with the mechanical aspects of the response generation process and was determined at random at the start of each trial from the equation

$$\theta = 5 + 2 * \text{exponential}, \quad (\text{B4})$$

where exponential was a sample from an exponential distribution with a mean of 1. The RT of the responding hand was then given by

$$\text{RT} = \text{RT}' + \theta. \quad (\text{B5})$$

Weights  $w_{ra}$  between each attention-response unit  $a$  and its corresponding response unit  $r$  were updated on each time step according to

$$w_{ra} = w_{ra} + \chi \alpha \delta, \quad (\text{B6})$$

where  $\delta$  was the temporal difference error at that time step (see Appendix C),  $\chi = 0.196$  if  $\delta \geq 0$  and  $\chi = 0.98$  if  $\delta < 0$ , and  $\alpha = .08$  (Table B1).

## Appendix C

### Monitor Module

The net input to each unit  $i$  from each unit  $j$  was

$$\text{net}_i = \sum_j w_{ij} a_j, \quad (\text{C1})$$

where  $a_j$  was the output of unit  $j$  and  $w_{ij}$  was the strength of the weight between units  $i$  and  $j$  (Table C1). The new output  $a_i$  of each unit  $i$  was then computed as

$$a_i = \frac{1}{(1 + e^{-g(\text{net}_i + b)})}, \quad (\text{C2})$$

Table C1

Critic Module Weight Matrix

Weight	H	S	L	R	Hs	Ss	Ls	Rs	HL	HR	SL	SR	F+	F-	V1	V2	V3	V4	V5	V6	V7	V8	
Hs	1	-1			1	-1																	
Ss	-1	1			-1	1																	
Ls			1				1	-1															
Rs				1			-1	1															
HL					1		1		3														
HR					1			1		3													
SL						1	1				3												
SR						1		1				3											
V1					3				-3	-3	-3	-3				-3							
V2						3			-3	-3	-3	-3			-3								
V3									3				-3	-3				-3	-3	-3			
V4										3			-3	-3			-3	-3	-3				
V5											3		-3	-3			-3	-3	-3				
V6												3	-3	-3			-3	-3	-3				
V7													3										-3
V8														3									
TD																	1	-1	-1	1			

Note. Cells indicate weights from units  $j$  (columns) to units  $i$  (rows). Shown is the case in which the  $H$  target is mapped to the left response and the  $S$  target is mapped to the right response. H =  $H$  categorization unit; S =  $S$  categorization unit; L = left response unit; R = right response unit; Hs =  $H$  target state unit; Ss =  $S$  target state unit; Ls = left response state unit; Rs = right response state unit; HL =  $H$  target-left response unit; HR =  $H$  target-right response unit; SL =  $S$  target-left response unit; SR =  $S$  target-right response unit; F+ = positive feedback unit; F- = negative feedback unit; V1 = value unit for  $H$  target; V2 = value unit for  $S$  target; V3 = value unit for  $H$  target-left response; V4 = value unit for  $H$  target-right response; V5 = value unit for  $S$  target-left response; V6 = value unit for  $S$  target-right response; V7 = value unit for positive feedback; V8 = value unit for negative feedback; TD = temporal difference unit.

where  $g$  was the gain (1 for the TD unit; 100,000 for all of the other units) and  $b_i$  was the bias of unit  $i$  (Table C2).

The TD signal  $\delta$  was determined by the change in activation of the TD unit over two successive cycles:

$$\delta^t = a_{TD}^t - a_{TD}^{t-1}, \tag{C3}$$

where  $t$  was the current cycle and TD indicates the TD unit (Sutton, 1988; Sutton & Barto, 1998).

The change in weights  $\Delta V_j$  between each value unit  $j$  and the TD unit was determined by

$$\Delta V_j = 0.8a_j^{t-2}\delta^t, \tag{C4}$$

and the weights were updated according to

$$V_j = V_j + \Delta V_j. \tag{C5}$$

Table C2  
Critic Unit Biases

Unit	No. units	Bias
Target state detection units	2	-0.70
Response state detection units	2	-0.42
Conjunction units	4	-1.95
Feedback units	2	-0.95
Value units	10	-0.95
TD unit	1	0.00

Note. TD = temporal difference.

Received November 28, 2002  
Revision received January 4, 2005  
Accepted January 14, 2005 ■