Psychological Review
2007, Vol. 114, No. 2, 376–397

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.2.376

Inhibitory Control in Mind and Brain: An Interactive Race Model of Countermanding Saccades

Leanne Boucher, Thomas J. Palmeri, Gordon D. Logan, and Jeffrey D. Schall
Vanderbilt University

The stop-signal task has been used to study normal cognitive control and clinical dysfunction. Its utility is derived from a race model that accounts for performance and provides an estimate of the time it takes to stop a movement. This model posits a race between *go* and *stop* processes with stochastically independent finish times. However, neurophysiological studies demonstrate that the neural correlates of the go and stop processes produce movements through a network of interacting neurons. The juxtaposition of the computational model with the neural data exposes a paradox—how can a network of interacting units produce behavior that appears to be the outcome of an independent race? The authors report how a simple, competitive network can solve this paradox and provide an account of what is measured by stop-signal reaction time.

Keywords: stop-signal task, cognitive control, frontal eye field, cognitive modeling, stochastic decision models

The task of cognitive neuroscience is to bring behavioral and physiological data together to explain how mental computations are implemented in the brain. This task is difficult when behavioral and physiological data appear to contradict each other. In these situations, a new theory is required to resolve the contradiction. This article reports results from an endeavor to resolve a paradox in the behavioral and physiological analyses of the stop-signal task. For over 20 years, behavioral data have been modeled successfully in terms of a race between two independent processes that respond to the stop signal and the go signal (Logan & Cowan, 1984). However, the neural systems that control movements comprise layers of inhibitory interactions between neurons that implement movement inhibition and movement initiation (reviewed by Munoz & Schall, 2003). These two facts present a paradox: How can interacting neurons produce behavior that appears to be the outcome of independent processes? We present a new theory of performance in the stop-signal task—the *interactive race model*—which assumes that the *stop* and *go* processes are independent for most of their latent periods. After this latent period, a second stage occurs in which the stop process interacts strongly and briefly to interrupt the go process. The theory resolves the paradox and unifies behavioral and physiological perspectives on stop-signal task performance. More generally, our work illustrates a novel approach to bringing neurophysiological data to bear on quantitative computational model testing.

Leanne Boucher, Thomas J. Palmeri, Gordon D. Logan, and Jeffrey D. Schall, Department of Psychology, Vanderbilt University.

This work was supported by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience; National Science Foundation Grants BCS0218507 and BCS0446806; and National Institutes of Health Grants F32-EY016679, RO1-MH55806, RO1-EY13358, P30-EY08126, and P30-HD015052. We thank M. Paré for sharing data; J. Brown, C. Camalier, M. Leslie, R. Krauzlis, M. Paré, L. Pearson, E. Priddy, V. Stuphorn, and K. Thompson for comments; D. Shima for computer programming assistance; K. Reis for figures; and the Vanderbilt Advanced Center for Computing for Research and Education for access to the high-performance computing cluster (http://www.accre.vanderbilt.edu/research).

Correspondence concerning this article should be addressed to Leanne Boucher, Thomas J. Palmeri, Gordon D. Logan, or Jeffrey D. Schall, Department of Psychology, Vanderbilt University, Nashville, TN 37221. E-mail: leanne.boucher@vanderbilt.edu, thomas.j.palmeri@vanderbilt.edu, gordon.logan@vanderbilt.edu or jeffrey.d.schall@vanderbilt.edu

The Stop-Signal Task

The stop-signal task investigates the control of thought and action by probing subjects' ability to withhold a planned movement in response to an infrequent countermanding signal (see Figure 1a; e.g., Lappin & Eriksen, 1966; Logan, 1994; Logan & Cowan, 1984). Subjects are instructed to make a response as quickly as possible to a go signal (*no-stop-signal* trial). On a minority of trials, a stop signal is presented and subjects have to inhibit the previously planned response (*stop-signal trial*). Subjects' ability to inhibit the response is probabilistic due to variability in reaction times (RTs) of the stop and go processes and depends on the interval between the go-signal and stop-signal presentation, referred to as the *stop-signal delay* (SSD). A trial is labeled *signal inhibit* (or cancelled) if the subject inhibits the response that would have been produced otherwise. A trial is labeled as *signal respond* (or noncancelled) if the subject is unable to inhibit the response. Typically, as SSD increases, subjects' ability to inhibit the response decreases, so the probability of signal-respond trials increases. Plotting the probability of responding given a stop signal against SSD is described as the *inhibition function* and is illustrated in Figure 1. In addition to the inhibition function, other dependent measures include RTs on trials with no stop signal and RTs on trials in which a response was made despite the stop signal (i.e., the signal-respond trials).

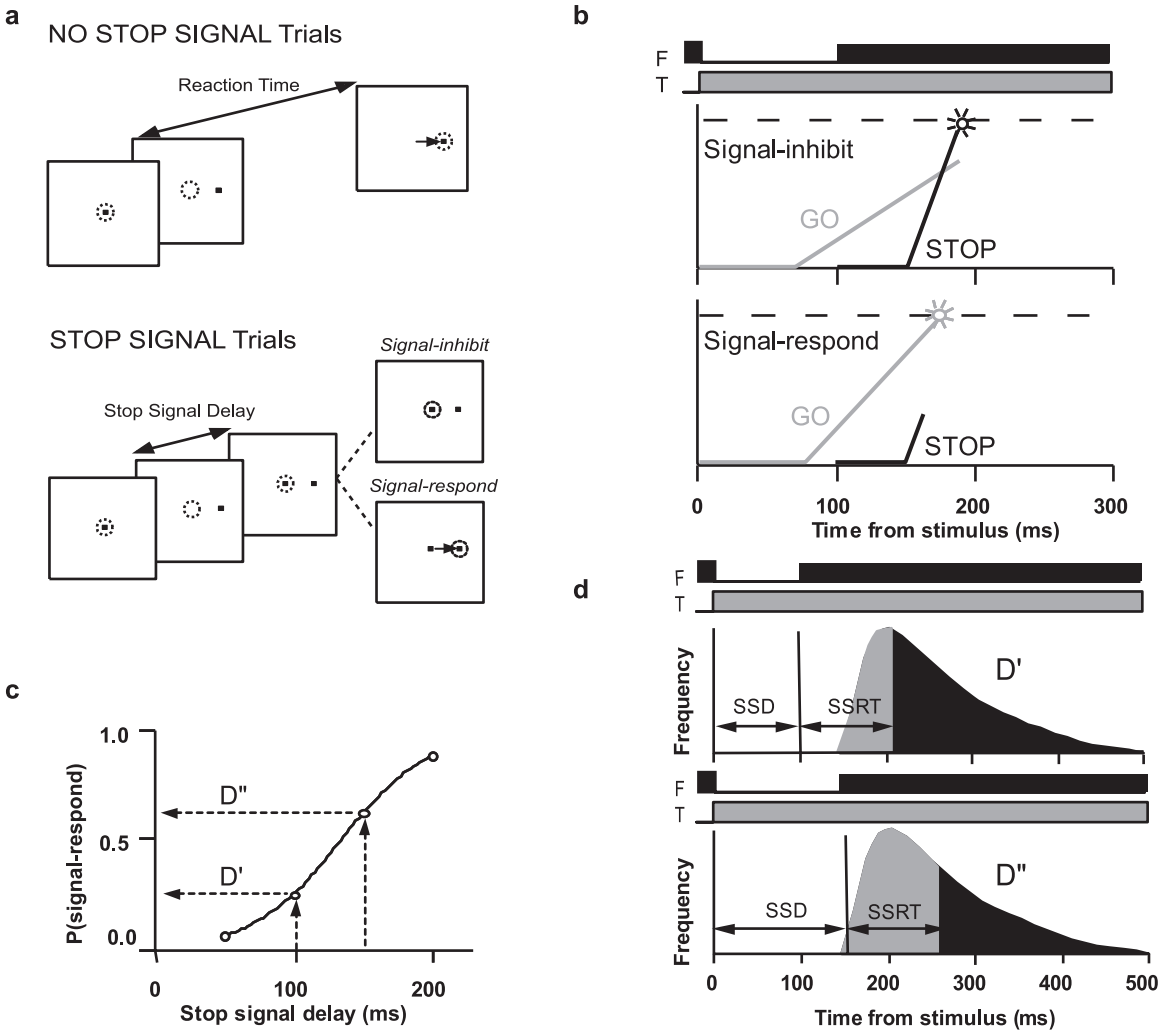

Figure 1. Saccade countermanding task. a: On no-stop-signal trials, monkeys were rewarded for shifting gaze to an eccentric visual target that appeared when the fixation point disappeared. On stop-signal trials, the fixation point reappeared after a variable delay following target presentation. Monkeys were positively reinforced if they inhibited the saccade when the stop signal appeared, but they produced signal-respond errors especially after longer stop-signal delays (SSDs). b: Race model outcome when $rt_{go} > rt_{stop}$ + SSD, resulting in a signal-inhibit trial (top panel), and when $rt_{go} < rt_{stop}$ + SSD, resulting in a signal-respond trial (bottom panel). Above each graph is a timeline marking the onset and offset of the fixation (F) and target (T). c: An idealized inhibition function plotting the proportion of signal-respond trials as a function of SSD. *D′* and *D″* indicate the proportion of signal-respond trials at SSDs of 100 ms and 150 ms, respectively. d: Schematic illustrating how stop-signal reaction time (SSRT) is calculated at two different SSDs by the integration method. See the text for details.

This task has been used extensively to study executive control and flexibility in behavior. Numerous experimental manipulations of the stop-signal task have yielded very similar results, demonstrating the generality of the task as an empirical model of self-control. For example, stop signals have ranged from visual (e.g., Lappin & Eriksen, 1966) to auditory (e.g., Logan, 1981; Logan, Cowan, & Davis, 1984; Osman, Kornblum, & Meyer, 1986, 1990) to tactile (e.g., Akerfelt, Colonius, & Diederich, 2005). Visual stop signals have been presented centrally or peripherally (e.g., Asrress & Carpenter, 2001; Cabel, Armstrong, Reingold, & Munoz, 2000). Responses have included key presses (e.g., Logan et al., 1984; Osman et al., 1986, 1990), typing responses (e.g., Logan, 1982), speech output (e.g., Ladefoged, Silverstein, & Papcun, 1973), arm movements (e.g., McGarry, Chua, & Franks, 2003; Slater-Hammel, 1960), hand squeezes (e.g., De Jong, Coles, & Logan, 1995; De Jong, Coles, Logan, & Gratton, 1990), hand movements (e.g., Boucher, Stuphorn, Logan, Schall, & Palmeri, in press), eye movements (e.g., Hanes & Carpenter, 1999; Logan & Irwin, 2000), and eye–head gaze shifts (e.g., Corneil & Elsley, 2005). Performance in this task has been investigated also in macaque monkeys (e.g., Hanes & Schall, 1995; Kornylo, Dill, Saenz, & Krauzlis, 2003) and rats (e.g., Eagle & Robbins, 2003; Feola, de Wit, & Richards, 2000).

Performance on the stop-signal task is qualitatively indistinguishable across all these stimulus modalities, effectors, and species. This generality implies that a unified account is possible. Such an account is provided by a race model, which assumes that performance is the outcome of a race between a go process responsible for initiating the movement and a stop process responsible for inhibiting the movement (Logan & Cowan, 1984; see also Becker & Jürgens, 1979; Lisberger, Fuchs, King, & Evinger, 1975; Logan, 1981; Olman, 1973). Since its mathematically explicit formulation over 20 years ago, this model has been tested extensively (e.g., Band, van der Molen, & Logan, 2003; Colonius, 1990; Colonius, Ozyurt, & Arndt, 2001; De Jong et al., 1990), and no major alternative theory has been proposed.

According to the race model, behavior is governed by the finish times of the go and stop processes on a given trial, indicated by rt_{go} and rt_{stop}, respectively. A response is initiated if $rt_{go} < rt_{stop} +$ SSD; a response is inhibited if $rt_{go} > rt_{stop} +$ SSD (see Figure 1b). Because the finish times for the go and stop processes are random variables (RT_{go} and RT_{stop}, respectively), inhibition is probabilistic, which gives rise to the inhibition function (see Figure 1c). (Note that we use the lowercase *rt* to refer to an RT on a particular trial and the uppercase *RT* to refer to the distribution of RTs.) This race model formulation can account for the distribution of signal-respond RTs by relating the proportion of signal-respond trials at each SSD (inhibition function) to the distribution of RT on trials with no stop signal in the following manner. When no stop signal is presented, the full RT_{go} distribution is produced. When a stop signal occurs, only a fraction of the RT_{go} distribution is produced because only the fastest RTs can escape inhibition. Said another way, only trials with a fast rt_{go} can finish before $rt_{stop} +$ SSD. As SSD increases, the proportion of longer latency trials that can escape inhibition increases because more time can elapse before the stop process finishes (see Figure 1d). The RTs of signal-respond trials share a common minimum, because the fastest samples from the RT_{go} distribution will always be produced regardless of SSD. However, the RTs of signal-respond trials will exhibit increasing maxima as SSD increases because the slowest samples from the RT_{go} distribution will be produced only at the longer SSDs. Figure 2b plots signal-respond RT distributions from a saccade stop-signal task as cumulative distribution functions. One can see how the signal-respond RT distributions exhibit a diverging pattern with the rightmost (i.e., slowest) distribution representing the no-stop-signal RTs and the leftmost (i.e., fastest) distribution representing the signal-respond RTs at the shortest SSD.

Beyond articulating an account of the functional mechanism responsible for performance, the race model provides a means for estimating the time needed to inhibit a response, referred to as the *stop-signal reaction time* (SSRT; Logan & Cowan, 1984). This is important because the response to the stop signal is not directly observable. If $rt_{stop} +$ SSD $< rt_{go}$, there is no overt response for which latency can be measured. If $rt_{stop} +$ SSD $> rt_{go}$, a response occurs, and we know that $rt_{stop} +$ SSD must have been longer than the latency of that response—however, we do not know how much longer. The race model provides several methods for estimating SSRT from the inhibition function and the no-stop-signal RT distribution (Colonius, 1990; De Jong et al, 1990; Logan & Cowan, 1984). The simplest method of estimating SSRT assumes that SSRT is a constant. The method is illustrated in Figure 1d. The finishing time of the stop process divides the no-stop-signal RT distribution into two parts, one in which rt_{go} is less than SSD + SSRT and one in which rt_{go} is greater than SSD + SSRT. The area under the no-stop-signal RT distribution corresponding to the first part equals the proportion of signal-respond trials at that SSD, and the area under the no-stop-signal RT distribution corresponding to the second part equals the proportion of signal-inhibit trials at that SSD. Thus, SSRT is estimated by finding the point that divides the no-stop-signal RT distribution into these two parts and subtracting SSD. Other methods of estimating SSRT do not assume that SSRT is constant (Colonius, 1990; De Jong et al., 1990; Logan & Cowan, 1984), but mathematical analysis (Logan & Cowan, 1984) and computer simulation (Band et al., 2003; De Jong et al. 1990) have shown that the assumption of constant SSRT does not bias the estimation substantially.

Measures of SSRT have been used extensively as indices of self-control across a variety of domains including human devel-

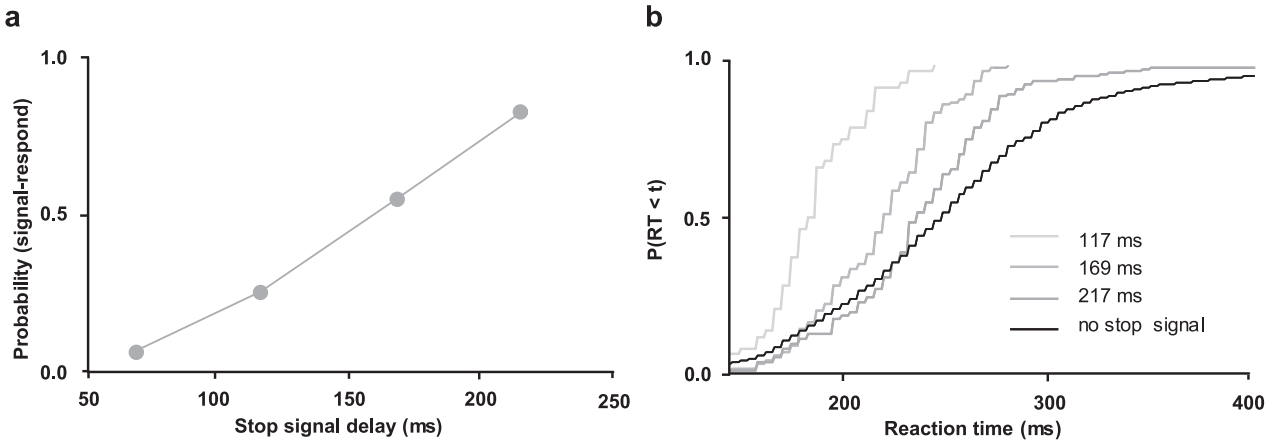

Figure 2. Observed behavioral data. a: Proportion of signal-respond trials as a function of stop-signal delay (SSD), that is, the inhibition function. b: Cumulative distribution of response times from no-stop-signal (black line) and signal-respond trials at different SSDs (lighter gray lines for progressively earlier SSDs).

opment (e.g., Bedard et al., 2002; Ridderinkhof, Band, & Logan, 1999; Williams, Ponesse, Schachar, Logan, & Tannock, 1999), aging (e.g., Kramer, Humphrey, Larish, Logan, & Strayer, 1994), and individual differences (e.g., Logan, Schachar, & Tannock, 1997; Miyake et al., 2000). Psychiatrists and clinical psychologists have used SSRT to assess deficits in cognitive control due to brain damage (Dimitrov et al., 2003), diseases of the brain such as Parkinson's disease (Gauggel, Rieger, & Feghoff, 2004), and psychopathologies such as schizophrenia (Badcock, Michie, Johnson, & Combrinck, 2002; Carter et al., 2003). SSRT measures have been particularly useful in understanding attention-deficit/hyperactivity disorder (ADHD; see Nigg, 2001, for review). Many studies have shown that SSRT is slower in children with ADHD than in psychiatric and nonpsychiatric control children (Armstrong & Munoz, 2003; Oosterlaan, Logan, & Sergeant, 1998; Schachar & Logan, 1990; Schachar, Mota, Logan, Tannock, & Klim, 2000; Schachar, Tannock, & Logan, 1993). Indeed, some researchers have suggested inhibitory control as the core deficit in ADHD; the other deficits being a consequence of a fundamental deficiency in inhibitory control (Barkley, 1997). SSRT is useful in clinical applications because it is a single datum that measures the inhibitory ability of a subject. SSRT can be correlated with measures of other abilities or used to distinguish the abilities of one group from those of another. All of this research rests on the validity of SSRT as a measure of inhibitory processing, which, in turn, relies on the validity of the race model. The race model's assumption that the stop and go processes are independent is critical. As documented below, certain forms of neural interaction may challenge the independence assumption in ways that could invalidate SSRT as a measure of inhibitory control. Thus, a major purpose of this article is to assess the validity of the race model and measures of SSRT that are derived from it.

Neural Correlates of Stopping

To make progress in linking cognitive processes with neurophysiology, researchers must identify the population of neurons that carry out a particular cognitive process. This mapping has been referred to as a *linking proposition*, and the population of neurons instantiating the linking proposition has been referred to as a *bridge locus* (Brindley, 1970; Schall, 2004; Teller, 1984; Teller & Pugh, 1983). Recent research using a saccade stop-signal task has identified a plausible bridge locus for the go and stop processes of the race model in particular neurons in the frontal eye fields (Hanes, Patterson, & Schall, 1998) and superior colliculus (Paré & Hanes, 2003), two structures responsible for saccadic preparation and generation. We identify the stop process with fixation cells and the go process with movement cells in frontal eye fields and superior colliculus.

Studies of the mechanisms controlling the initiation of visually guided saccades afford several advantages because so much is known about the neural processes producing saccadic eye movements (e.g., Carpenter, 1991; Girard & Berthoz, 2005; Munoz & Schall, 2003; Wurtz & Goldberg, 1989). Saccades are produced by a pulse of force that rapidly rotates the eyes followed by a step of force appropriate to resist the elastic forces of the orbit and maintain eccentric gaze. This pattern of force is exerted on the eyes by muscles innervated by neurons in the brainstem (Scudder, Kaneko, & Fuchs, 2002; Sparks, 2002; see Figure 3). Burst neurons innervate the extraocular motor neurons to provide the high-frequency burst of spikes necessary to produce saccadic eye movements. Different groups of burst neurons that discharge for saccades in different directions innervate different motor neurons

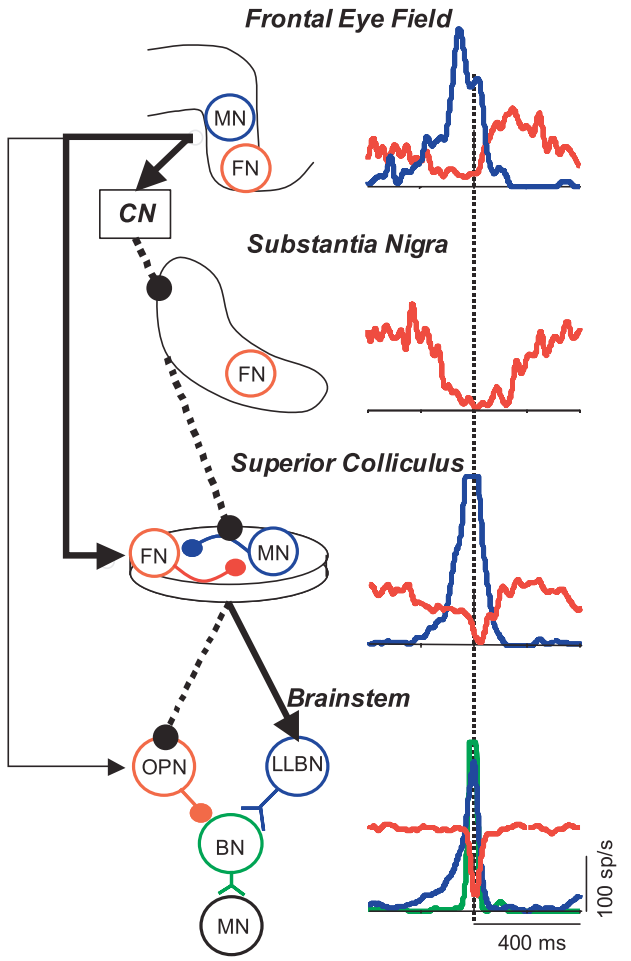

Figure 3. Saccade generator network. Saccades are controlled by a distributed network encompassing the frontal eye field, the basal ganglia including the substantia nigra, the superior colliculus, and circuits in the brainstem. Gaze-shifting and gaze-holding neurons are modulated concurrently throughout this network. Movement neurons (MN, blue) and fixation neurons (FN, red) are modulated reciprocally in the frontal eye field, substantia nigra, and superior colliculus. When the balance of activation of these neurons tips from gaze holding to gaze shifting, the omnipause neurons (OPN) in the brainstem are turned off, releasing inhibition on the burst neurons (BN) that activate the motor neurons (MN) that cause the extraocular muscles to contract rapidly to produce the saccade. The pattern of feedforward connectivity is illustrated on the left. The brainstem is innervated by the superior colliculus and less strongly by the frontal eye field to activate long-lead burst neurons (LLBN) and inhibit the omnipause neurons. The superior colliculus is inhibited by the substantia nigra. The inhibition from the substantia nigra is released through inhibition from the caudate nucleus (CN) that is activated by the frontal eye field. sp/s = spikes per second. Adapted from "Concurrent, Distributed Control of Saccade Initiation in the Frontal Eye Field and Superior Colliculus," by D. P. Munoz and J. D. Schall, 2003, In *The Superior Colliculus: New Approaches for Studying Sensorimotor Integration* (p. 177), New York: CRC Press. Copyright 2003 by CRC Press. Adapted with permission.

that in turn innervate different muscles. The burst neurons are subject to potent monosynaptic inhibition from omnipause neurons. Omnipause neurons discharge tonically during fixation. Immediately prior to initiation of a saccade in any direction, omnipause neurons cease discharging, releasing the inhibition on the appropriate pools of burst neurons to produce the burst in the motor neurons necessary to shift gaze in the desired direction. Upon completion of the saccade, omnipause neurons reactivate to reinstate inhibition on the burst neurons. Omnipause neurons are not modulated during the ~100-ms period before saccade initiation when preparatory processes that affect RTs occur (Everling, Paré, Dorris, & Munoz, 1998).

The neural events preceding activation of the brainstem saccade generator occur in a circuit distributed through particular areas of the frontal lobe (e.g., Bruce & Goldberg 1985; reviewed by Schall, 1997), the basal ganglia (e.g., Hikosaka & Wurtz, 1983a, 1983b, 1983c; reviewed by Hikosaka, Takikawa, & Kawagoe, 2000), cerebellum (e.g., Kase, Miller, & Noda, 1980; reviewed by Thier, Dicke, Haas, Thielert, & Catz, 2002), and superior colliculus (Schiller & Körner, 1971; Sparks, 1975; Wurtz & Goldberg, 1972; reviewed by Munoz, Dorris, Paré, & Everling, 2000; Munoz & Schall, 2003; see Figure 3). This circuit conveys to the brainstem saccade generator signals controlling where and when to shift gaze. The superior colliculus is organized in a topographic map of saccade direction and amplitude. The frontal eye field has a rougher map of saccade amplitude, and the frontal eye field and superior colliculus are connected topographically. Thus, the direction and amplitude of the saccade produced is dictated by the location in the map of the active population of neurons. However, neurons in frontal eye field and superior colliculus have broad movement fields, so many neurons contribute to the generation of any saccade by pooling activity through vector averaging.

The stop-signal task is ideal for investigating the neural control of movement initiation because it specifies the criteria a neuron must meet to be identified as contributing to controlling saccade initiation. First, the activity in trials when a saccade is made (no-stop-signal or signal-respond trials) must be different from that in trials when no saccade is made (signal-inhibit trials). Second, in stop-signal trials, the activity should begin along the trajectory that would lead to saccade initiation, but on presentation of the stop signal, the activity must be modulated away from that trajectory, and this modulation must occur within the SSRT. Neurons with movement-related and fixation-related activity in frontal eye field and superior colliculus satisfy both of these requirements (Hanes et al., 1998; Paré & Hanes, 2003). On signal-respond trials, the neural activity of movement and fixation neurons is indistinguishable from that observed in no-stop-signal trials (see Figure 4a). On signal-inhibit trials, the activity of movement neurons initially increases, but following presentation of the stop signal, the activity rapidly decreases within SSRT (see Figure 4b). Because the movement neuron does not reach threshold, no saccade is made. Fixation neuron activity exhibits the complementary pattern of modulation (see Figures 4c and 4d). Note that the time of modulation of movement- and fixation-related activity on signal-inhibit trials occurs within the SSRT, confirming the primary role of movement and fixation neurons in controlling the initiation of saccades.

The identification of the go process with the activity of long-lead presaccadic movement-related neurons in the frontal eye field, caudate nucleus, superior colliculus, and brainstem seems evident and justified. The go process cannot be identified with the activity of motor neurons or the burst neurons in the brainstem because the modulation of these neurons occurs during the ballistic period of movement production, during which the saccade cannot be inhibited. Numerous studies have demonstrated that a progressive commitment to saccade initiation occurs with growing activation of the long-lead movement neurons in frontal eye field (Hanes et al., 1998; Hanes & Schall, 1996), caudate nucleus (Lauwereyns, Watanabe, Coe, & Hikosaka, 2002), superior colliculus (Dorris & Munoz, 1998; Paré & Hanes, 2003), and brainstem (Kaneko, 2006). In fact, saccades are initiated when the activity of movement neurons reaches a fixed threshold, so the speed of growth of long-lead presaccadic movement neuron activity is a strong predictor of saccade initiation time (Dorris & Munoz, 1998; Hanes & Schall, 1996). Finally, when saccades are withheld during this saccade countermanding task, movement neurons with activity that had been growing toward the threshold exhibit a reduction of discharge rate within SSRT (Hanes et al. 1998; Paré & Hanes, 2003; see Figure 4b).

For saccade initiation to be cancelled, the go process must be interrupted by the stop process. Therefore, the stop process for saccade countermanding with a foveal stop signal must be identified with the activity of neurons that can inhibit the activity of long-lead presaccadic movement neurons. As illustrated in Figure 3, many studies have provided evidence for neurons with activity reciprocal to movement neurons that can inhibit movement neurons; these neurons are referred to collectively as *fixation neurons*. The stop process cannot be identified with the activity of omnipause neurons in the brainstem because these neurons are modulated only during the ballistic period of movement production, during which the saccade cannot be inhibited; in fact, according to all current models of saccade generation, the ballistic phase of saccade production starts only when omnipause neurons turn off. Fixation neurons and omnipause neurons are functionally distinct. In particular, fixation neurons exhibit modulation of activity as much as 100 ms before saccade initiation that may or may not result in saccade initiation; omnipause neurons modulate only ~10 ms before saccade initiation, with a saccade initiated each time the omnipause neurons turn off (e.g., Everling et al., 1998). Many studies have demonstrated, before saccade initiation, a gradual reduction of activation of fixation neurons in frontal eye field (Hanes et al., 1998; Segraves & Goldberg, 1987), substantia nigra (e.g., Hikosaka & Wurtz, 1983c), and superior colliculus (Dorris, Paré, & Munoz, 1997; Everling et al., 1998; Munoz & Wurtz, 1993a; Paré & Hanes, 2003). Finally, when saccades are withheld in this saccade countermanding task, recordings of single-unit activity demonstrate that fixation neurons with activity that had been reduced during preparation for saccade initiation exhibit an increased discharge rate within SSRT (see Figure 4d; Hanes et al., 1998; Paré & Hanes, 2003).

The reciprocity of the activation of movement and fixation neurons naturally suggests the hypothesis that they are in a mutually inhibitory relationship, and evidence consistent with this has been obtained. It must be emphasized that effectively all of the experiments performed to date have addressed the question of how saccades or eye–head gaze shifts are terminated, which is not necessarily the same as how saccades or gaze shifts are prevented from being initiated (reviewed by Guitton, Bergeron, & Choi, 2004; Sparks, 2002). Electrical stimulation of the rostral superior

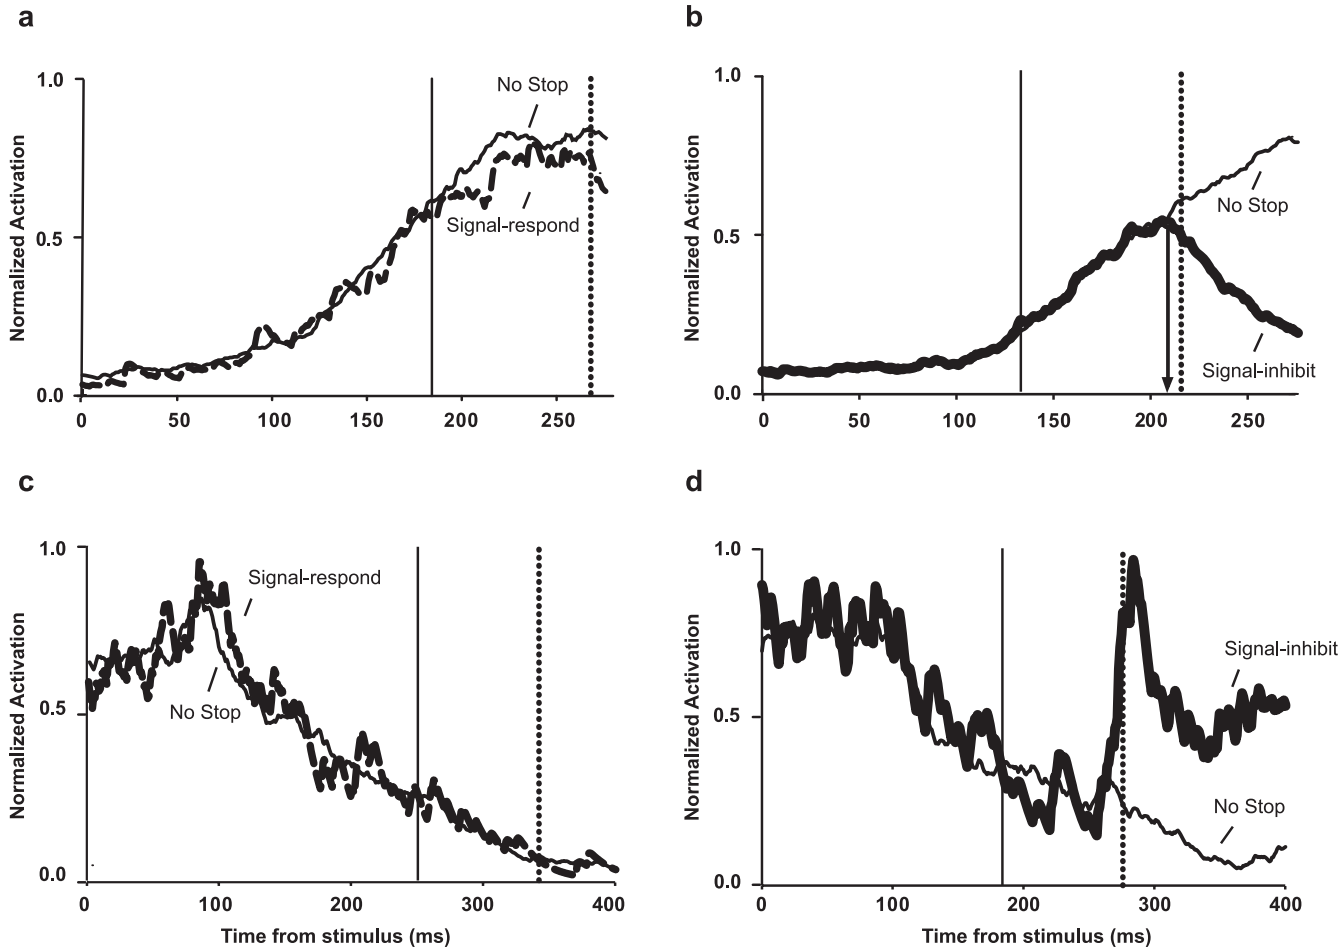

Figure 4. Movement and fixation neuron activity during saccade countermanding. One stop-signal delay (SSD) is displayed, although the neural profiles at all SSDs are similar. a: Average normalized activity from 12 frontal eye field movement neurons from Monkey A on signal-respond trials (thick dashed line, $N = 43$) and latency-matched no-stop-signal trials (thin line, $N = 399$). The activity was not significantly different. The solid vertical line marks stop-signal presentation time. The dotted vertical line indicates stop-signal reaction time determined from the performance while the neural activity was monitored. b: Average normalized activity from 12 frontal eye field movement neurons on signal-inhibit trials (thick line, $N = 121$) and latency-matched no-stop-signal trials (thin line, $N = 477$). The activity was significantly different; the growth of movement-related activity was interrupted immediately before the stop-signal reaction time, marked by the downward arrow. This time is referred to as the cancel time. c: Average normalized activity from 6 frontal eye field fixation neurons from Monkeys A and C on signal-respond trials (thick dashed line, $N = 40$) and latency-matched no-stop-signal trials (thin line, $N = 6$). The activity was not significantly different. d: Average normalized activity from 11 frontal eye field fixation neurons from Monkeys A and C on signal-inhibit trials (thick line, $N = 18$) and latency-matched no-stop-signal trials (thin line, $N = 111$). The activity was significantly different; the growth of fixation-related activity increased immediately before the stop-signal reaction time.

colliculus, where fixation cells are concentrated, stops eye–head gaze shifts in cats (Paré & Guitton, 1998) and saccades in monkeys (Gandhi & Keller, 1999). This influence could be exerted through more than one circuit. First, fixation neurons could prevent gaze saccades through projections to omnipause neurons that inhibit the saccade generator (e.g., Buttner-Ennever, Horn, Henn, & Cohen, 1999; Takahashi, Sugiuchi, Izawa, & Shinoda, 2005; Yoshida, Iwamoto, Chimoto, & Shimazu, 2001). Second, fixation neurons could inhibit movement neurons elsewhere in the superior colliculus (Munoz & Istvan, 1998; Takahashi et al., 2005). However, it must be acknowledged that the evidence is not decisive in support of the hypothesis that the movement neurons in the superior colliculus are gated by fixation neurons. For example, the movement neuron bursts tend to end just before or at the end of gaze shifts (Freedman & Sparks, 1997; Waitzman, Ma, Optican, & Wurtz, 1991), but fixation neurons exhibit peak activity ~50 ms after gaze shifts (Choi & Guitton, 2006; Everling et al., 1998; Munoz & Wurtz, 1993a). Furthermore, experimental deactivation of the rostral superior colliculus results in excessive saccade initiation, but the saccades do stop (Munoz & Wurtz, 1993b). The

body of evidence suggests that fixation neurons may not contribute to controlling saccade execution, but they are necessary for controlling saccade initiation, which is, by definition, just what the stop process is supposed to do. The evidentiary and conceptual gaps that prevent a more definite conclusion that the stop process corresponds to the activity of gaze-holding fixation neurons highlight the need for and guide the focus of further neuroanatomical and neurophysiological research to map these circuits.

Models of Saccade Production and RT

Sophisticated models of saccade generation have been formulated (for a review, see Girard & Berthoz, 2005). Models of the brainstem circuits that generate saccades include a trigger signal that releases the inhibition of omnipause neurons on burst neurons to initiate saccades, but these models do not specify any characteristics of the trigger signal. The input to the brainstem arrives from the superior colliculus and frontal eye field. Models of superior colliculus contribution to saccade preparation and generation have also been formulated that address several issues beyond the scope of this article (reviewed by Girard & Berthoz, 2005). For our purposes, we note that lateral inhibition is a central characteristic of many models of saccade production (Arai, Keller, & Edelman, 1994; Bozis & Moschovakis, 1998; Das, Gandhi, & Keller, 1995; van Opstal & van Gisbergen, 1989). The model of Trappenberg, Dorris, Klein, and Munoz (2001) implemented short-range excitation and long-range inhibition across the map of saccades in the superior colliculus to account for the systematic variation of saccade initiation times in different conditions. Saccade initiation time is influenced in this model by the balance of inhibition between fixation and movement neurons. For example, the model produces the reduction of saccade latency consequent to removal of a fixation point (gap effect) by reduced activation of fixation neurons leading to reduced inhibition of movement neurons, allowing them to trigger the saccade sooner. However, a recent comprehensive review of saccade models concluded that the dynamics of these lateral interactions required more study to resolve the diverse and in some cases contradictory results (Girard & Berthoz, 2005).

Models that include cortical and basal ganglia circuits implement gating of saccade initiation by selective disinhibition of the superior colliculus. This disinhibition has been conceived of mainly in terms of target selection or working memory (e.g., J. W. Brown, Bullock, & Grossberg, 2004; Dominey & Arbib, 1992; Dominey, Arbib, & Joseph, 1995). In the model of J. W. Brown et al. (2004), the indirect pathway through the basal ganglia implements a trainable stopping process used when a delay intervenes before a saccade. This stopping process operates by increasing the activity in the subthalamic nucleus that produces additional activation of the substantia nigra pars reticulata, thus increasing the inhibition on the saccade production circuit. Only a few models of saccade production have incorporated elements of frontal eye field functional organization (J. W. Brown et al., 2004; Mitchell & Zipser, 2003). These models use a fixation signal to specify the timing of saccade initiation.

A significant shortcoming of all of these models is the lack of stochastic elements capable of accounting for the range of saccade latencies. Without stochastic elements, none of the aforementioned models can account for performance in the stop-signal task. This is the central concern of our work, which seeks to build a bridge between neural network models of saccade generation formulated at a circuit level and cognitive models of decision making formulated at a computational level (e.g., S. Brown & Heathcote, 2005; Carpenter & Williams, 1995; Link, 1975; Nosofsky & Palmeri, 1997; Ratcliff, 1978; Usher & McClelland, 2001). Properly formulated, such a model can help translate more concretely between neural circuit level models of saccade production and higher level models of eye movement production during tasks such as reading (e.g., Engbert, Nuthmann, Richter, & Kliegl, 2005; McDonald, Carpenter, & Shillcock, 2005; Reichle, Rayner, & Pollatsek, 2003).

A Paradox Motivating the Interactive Race Model

The identification of the go and stop processes with movement and fixation neurons exposes a paradox. Given the overwhelming evidence indicating that saccades are produced by a network of mutually inhibitory gaze-shifting and gaze-holding neurons, how can such interacting neural units produce behavior that appears to be the outcome of a race between processes with independent finishing times? This paradox must be resolved because the validity of SSRT as a measure of cognitive control derives entirely from the validity of the independent race model.

The independent race model is almost universally accepted in the literature on stop-signal studies. If the basic assumption of independence is false, then this entire body of work may require reinterpretation. In particular, the validity of estimates of SSRT relies on the assumption that the stop process has a specific duration that can be measured through application of the race model to behavior. The duration of the stop process is the SSRT. If the activity of movement-related and fixation-related neurons instantiates the go and stop processes and these neurons interact over an extended period of time, then the stop process would not have a definite finishing time and so could not be described validly in terms of an RT measure like SSRT. If SSRT is not a valid measure of inhibitory control, then the utility of the stop-signal paradigm in studying clinical disorders and development, as well as normal movement preparation and control, would be called into question. Thus, an important outcome of this work is to establish a solid foundation for the measurement and interpretation of SSRT by obtaining a more secure understanding of how movement preparation is interrupted.

It is critical to note that the independent race model is concerned exclusively with the finish times of the go and stop processes. The race model makes no assumptions about the processes by which the RTs are generated beyond assuming they are stochastically independent. This has been regarded as a virtue because it makes the race model applicable to any distribution of stop and go times, provided they are stochastically independent (Logan & Cowan, 1984). This study is concerned with only how the finish times of the go and stop processes can be generated in the neural circuit producing saccadic eye movements.

We created a simple network consisting of two noisy accumulators identified as a go unit and a stop unit. We contrasted the performance of an independent race model in which the go and stop units raced to a fixed threshold (see Figure 5a) with the performance of an interactive race model in which the go and stop units interacted through inhibitory connections (see Figure 6a). We

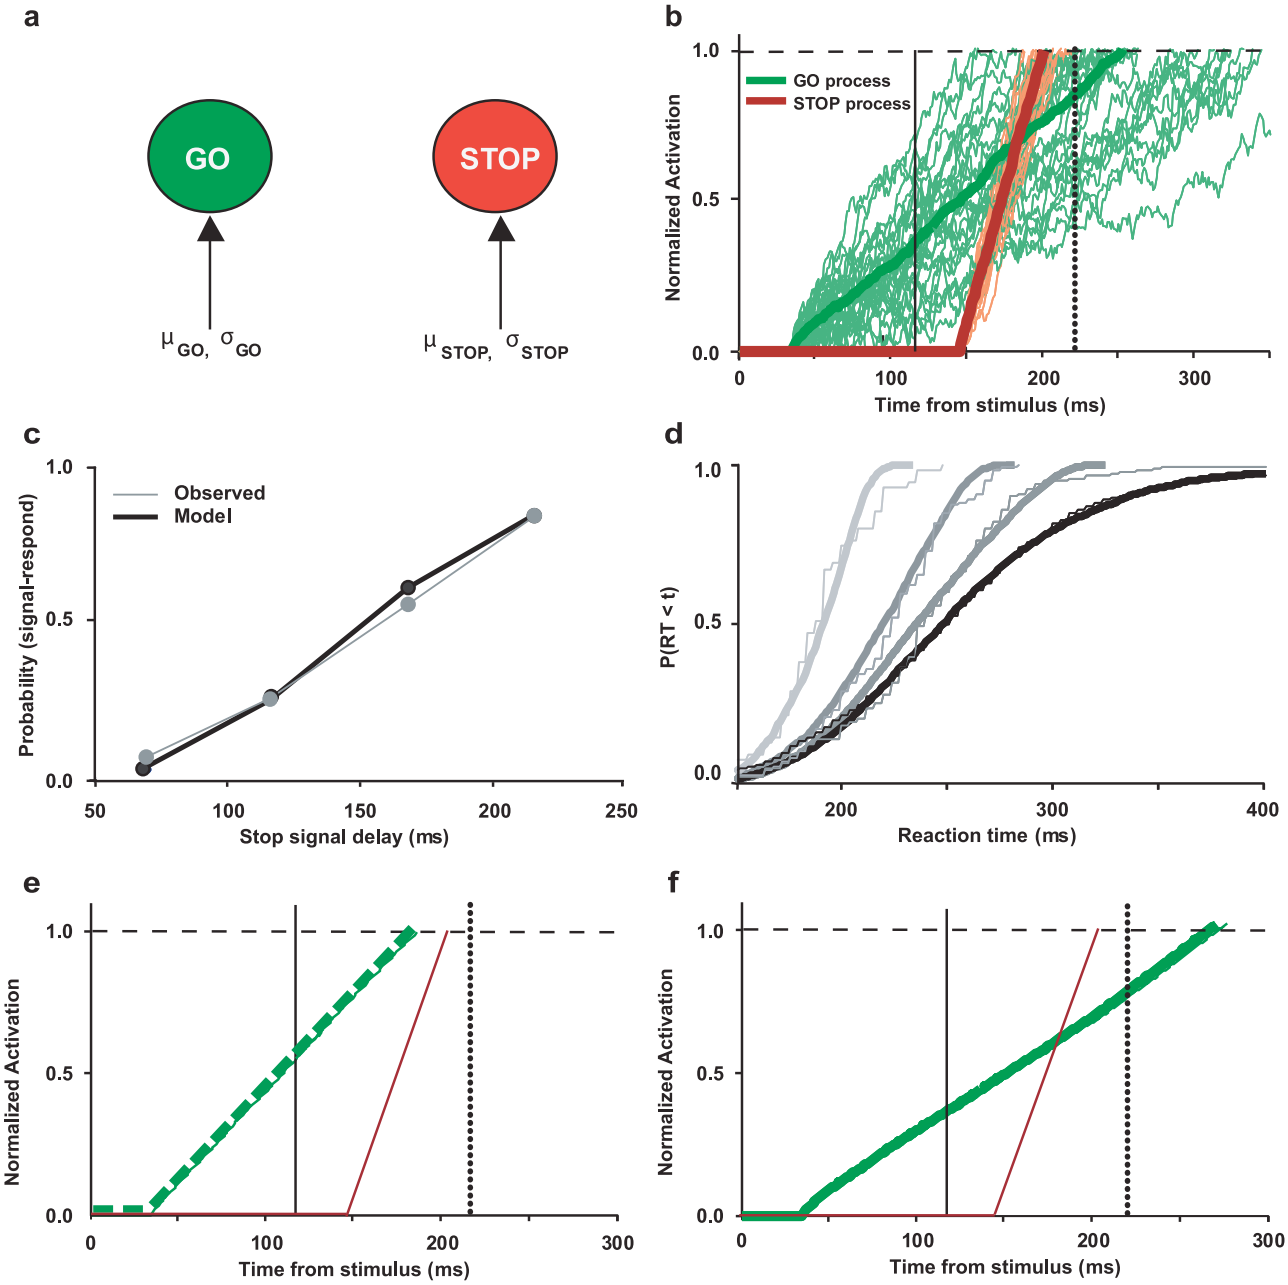

Figure 5. Independent race model. Simulated stop-signal trials were classified as signal respond if the go unit reached threshold before the stop unit and as signal inhibit if the stop unit reached threshold before the go unit. One stop-signal delay (SSD) is displayed, although the activation functions at all SSDs are similar. a: Independent race model architecture. b: Thirty trials of simulated go (green) and stop (red) unit activation with an SSD of 117 ms (solid vertical line); thick lines are mean activation functions, and the threshold is the dashed horizontal line. The vertical dotted line marks stop-signal reaction time (SSRT). c: Observed (gray line) and simulated (black line) inhibition functions. d: Observed (thin lines) and simulated (thick lines) reaction time distributions from no-stop-signal (black line) and signal-respond trials with progressively longer SSDs (progressively darker gray lines). As expected, the independent race model simulated performance very well. e: Average normalized go unit (green) and stop unit (red) activation functions on signal-respond (thick dashed line) and latency-matched no-stop-signal trials (thin solid line). Stop-signal presentation is indicated by the solid vertical line; SSRT is indicated by the dotted vertical line. f: Average normalized go unit (green) and stop unit (red) activation functions on signal-inhibit (thick solid line) and latency-matched no-stop-signal trials (thin solid line) with stop-signal presentation and SSRT indicated. On simulated signal-inhibit trials, the go unit reaches threshold because nothing interrupts it.

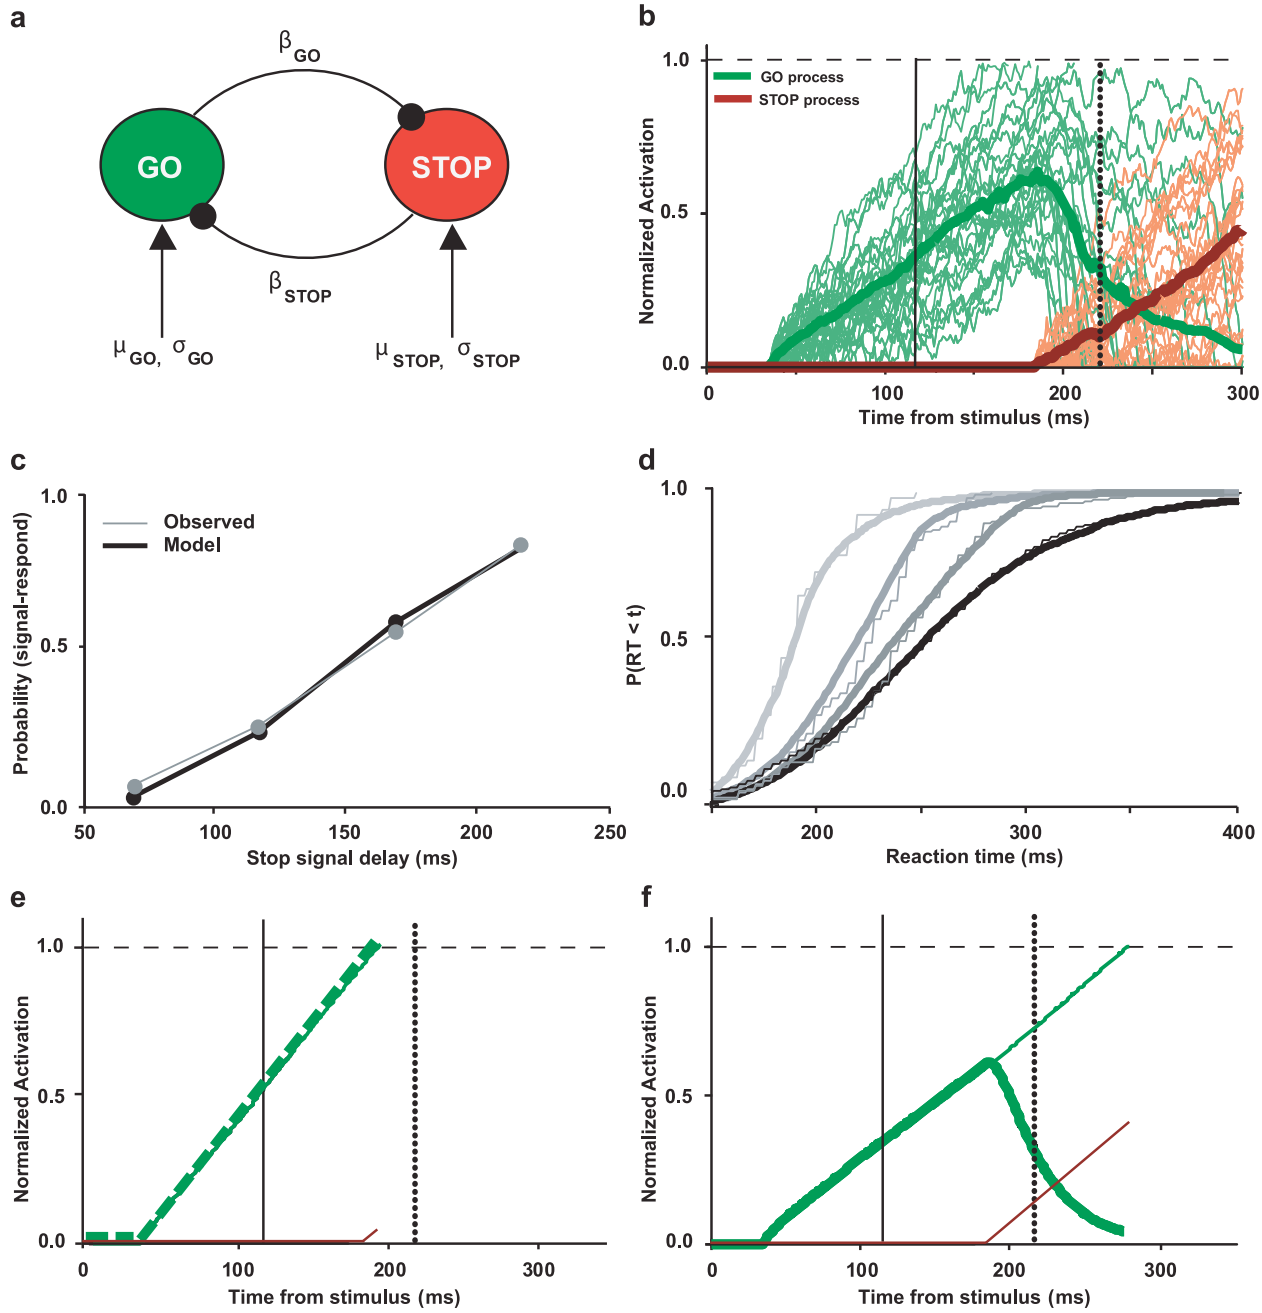

Figure 6. Interactive race model. Simulated stop-signal trials were classified as signal respond if the go unit reached threshold and as signal inhibit if the stop unit prevented the go unit from reaching threshold. a: Interactive race model architecture. b: Thirty trials of simulated go (green) and stop (red) unit activation with a stop-signal delay (SSD) of 117 ms (solid vertical line); thick lines are mean activation functions, and the threshold is the dashed horizontal line. The vertical dotted line marks stop-signal reaction time (SSRT). c: Observed (gray line) and simulated (black line) inhibition functions. d: Observed (thin lines) and simulated (thick lines) reaction time distributions from no-stop-signal (black line) and signal-respond trials with progressively longer SSDs (progressively darker gray lines). e: Average normalized go unit (green) and stop unit (red) activation functions on signal-respond (thick dashed line) and latency-matched no-stop-signal trials (thin solid line). Stop-signal presentation is indicated by the solid vertical line; SSRT is indicated by the dotted vertical line. f: Average normalized go unit (green) and stop unit (red) activation functions on signal-inhibit (thick solid line) and latency-matched no-stop-signal trials (thin solid lines) with stop-signal presentation and SSRT indicated. Note the absence of go unit modulation on signal-respond trials and the pronounced modulation of the go unit before SSRT on signal-inhibit trials.

first found the best quantitative fit of the alternative model architectures to behavioral data; these data were collected from macaque monkeys performing a saccade stop-signal task while neurophysiological data were collected from the frontal eye field (Hanes et al., 1998). We then assessed the correspondence of the go unit activation in the models with the pattern of movement-related neural activation; these neurophysiological recordings were made at the same time as the behavioral data were collected on which the models were optimized. We also assessed the correspondence of the stop unit activation in the models with the pattern of fixation-related activation. However, for this comparison we included fixation-related neurons recorded in the superior colliculus during a saccade stop-signal task (Paré & Hanes, 2003) because, being rare, the number of fixation-related neurons in the frontal eye fields was too small for meaningful analysis. In other words, the models were fit to the behavioral data, and then the activation of the go and stop model units was compared with the patterns of activity of movement and fixation neurons. The activity of purely visually responsive neurons was not incorporated into this model analysis because those neurons in frontal eye field and superior colliculus do not influence saccade initiation (Hanes et al., 1998; Paré & Hanes, 2003).

This approach was taken for two reasons. First, prior research has shown that the race model provides excellent quantitative accounts of behavioral data. Any new model motivated by neurophysiology must fit the behavioral data as well as the race model irrespective of its account of neural data. Second, fitting the models to only the behavioral data allowed us to assess how well they predict important aspects of the underlying neurophysiology rather than simply fitting the neurophysiology with additional free parameters. Thus, the activity of movement and fixation neurons served as logically converging evidence for selecting among competing models.

To satisfy the behavioral constraints, a viable network architecture must account for the RT distributions in no-stop-signal and signal-respond trials and for the proportion of signal-respond trials at each SSD (see Figures 2a and 2b). Consequently, a successful architecture must produce simulated data yielding an SSRT that is indistinguishable from that derived from the observed data. To satisfy the neural constraint, the average activation of the units in the network must correspond qualitatively and quantitatively to the pattern of activity of neurons that were recorded as the behavioral data were being collected (see Figure 4). Qualitatively, a successful model must produce go unit activation accumulating to reach threshold on trials in which a saccade is made in no-stop-signal or signal-respond trials and a decrease in activation after an initial accumulation when saccades are inhibited on signal-inhibit trials. Conversely, stop unit activation must remain off when saccades are made in no-stop-signal trials or signal-respond trials and become activated when saccades are withheld on signal-inhibit trials. Quantitatively, the time of this modulation relative to SSRT, referred to as *cancel time*, must fall within the range of such times obtained from neurons in the frontal eye field and superior colliculus. For movement-related neurons, the average cancel time measured was 8.5 ms before SSRT in the frontal eye field (Hanes et al., 1998) and 10 ms before SSRT in the superior colliculus (Paré & Hanes, 2003). For fixation-related neurons, the average cancel time was 0.2 ms after SSRT (not significantly different from 0) in the frontal eye field (Hanes et al., 1998) and 13 ms before SSRT in the superior colliculus (Paré & Hanes, 2003).

Behavior, Neurophysiology, and Computational Modeling

Behavioral Measures

The data fit by the model were collected in the first neurophysiological study using the stop-signal task (Hanes et al., 1998). These data were collected from the frontal eye field, but qualitatively and in important respects quantitatively identical data have been obtained in a subsequent study of the superior colliculus (Paré & Hanes, 2003). Two monkeys were trained in a saccade stop-signal task for the frontal eye field study. Each trial began with the presentation of a central fixation spot. After a variable delay, the fixation point was extinguished and a target appeared either within a neuron's receptive field or opposite that location. Monkeys were rewarded if they shifted gaze to the target location. On a fraction of trials (average of 29% for Monkey C and 25% for Monkey A), the fixation point reilluminated, and monkeys were rewarded for maintaining fixation. The interval between fixation point disappearance and reappearance was the SSD.

The behavioral data fit by the model constitute a select subset of all the data reported in Hanes et al. (1998). The following inclusion criteria were applied: (a) The behavioral data were collected while a clear movement-related neuron was recorded, (b) the inhibition function included at least three SSDs and spanned from less than 0.2 at the shortest SSD to more than 0.7 at the longest SSD, and (c) more than 100 trials were collected at each SSD. Adopting these strict criteria ensured the consistency but reduced the quantity of data, so we combined neural and behavioral data across sessions and neurons to obtain a sufficiently large data set to evaluate competing models. The activity of fixation cells recorded in the frontal eye field and superior colliculus during other sessions with comparable performance was also compared with the model unit activation.

On the basis of these criteria, we culled 1,976 trials from 5 movement neurons for Monkey C and 2,075 trials from 12 movement neurons for Monkey A. Although the number of neurons may seem small, the number of trials is more important because model parameters are optimized on the basis of the behavioral data and not the neural data. SSDs were 69, 117, 169, and 217 ms for Monkey C and 84, 101, 134, 184, 201, and 234 ms for Monkey A. The data fit by the models were the RT distributions for no-stop-signal trials, the RT distributions for signal-respond trials at each SSD, and the inhibition function (i.e., the proportion of signal-respond trials at each SSD; see Figure 2).

We calculated the SSRT for each monkey separately using the integration method (Logan, 1994; Logan & Cowan, 1984). For each SSD, SSRT was defined as the RT at which the integral of the no-stop-signal RT distribution equaled the proportion of signal-respond trials at that SSD minus the SSD. SSRT was then averaged over SSD for each monkey. In the data set fit by the model, SSRTs were 71 ms and 94 ms for Monkeys A and C, respectively.

Neurophysiological Measures

Neural activation functions were derived from spike trains that were converted to spike density functions as described previously

(Hanes et al., 1998). It is important to note that we are not fitting the fine-scale fluctuations of neural activity; rather, we are assessing qualitatively the pattern of modulation and quantitatively the time of the major modulations in the different types of trials. This approach is based on the premise that a population of neurons with a specific function respond in generally the same way, but each neuron in that population has idiosyncrasies probably derived from incidental variation in local circuit connectivity (e.g., Marder & Goaillard, 2006). For example, all movement-related neurons increase their activity before a movement, but one neuron may have a stronger burst while another may grow at a steady rate. To average across neurons, we first normalized the spike density function of each neuron by dividing its activity by the peak firing rate in the interval from 20 ms before to 50 ms after saccade initiation on no-stop-signal trials. Go unit activation was compared with movement neuron activity. Stop unit activation was compared with fixation neuron activity.

For both neurons and model units, activation on signal-inhibit or signal-respond trials was compared with the activity of a subset of *latency-matched* no-stop-signal trials. No-stop-signal trials with $rt > \text{SSD} + \text{SSRT}$ were compared with signal-inhibit trials, because according to the race model, the saccade would have been inhibited had the stop signal been presented. No-stop-signal trials with $rt < \text{SSD} + \text{SSRT}$ were compared with signal-respond trials, because according to the race model, the saccade would have been initiated even if the stop signal had occurred.

For the go unit, as for movement neurons, *cancel time* was defined as the time at which activation on signal-inhibit trials significantly diverged from the activation on no-stop-signal trials relative to SSRT. This required measuring the difference between the time-varying activation in these two trial types. In the original neurophysiological study, a particular measurement procedure and criterion were used that afforded the necessary balance between sensitivity and robustness (Hanes et al., 1998; see also Paré & Hanes, 2003). To provide comparable measurements for this study, we used the same procedure and criterion. Specifically, cancel time was defined as the instant the difference in the activation on signal-inhibit and latency-matched no-stop-signal trials became significant, that is, exceeded the largest difference expected by chance. The baseline distribution of differences expected by chance was measured in the period before any modulation was possible, that is, from 600 ms before target onset to SSD for the neural data and the time between the start of the go unit activation and the start of the stop unit activation for the model. The criterion for a significant departure from random variation was when the difference first surpassed 2 *SD* of the baseline difference provided the difference surpassed 6 *SD* within the next 50 ms.

Analyses of neural data show an effect of noise on estimates of cancel time: the noisier the activation functions, the later the time of the difference measured by the statistical criteria (Hanes et al., 1998). The model unit activation functions have considerably less variability than neural spike density functions because of the tremendously large number of trials that are simulated to generate stable model predictions. Consequently, we subsampled the activation of the go and stop units in a manner comparable to the sampling of neural activity in the physiology experiments via the following procedure. We simulated the model with 20–50 trials at each SSD to mimic the number of trials typically obtained in the physiology experiments. We then calculated the average activation functions of the simulated trials and determined the time of modulation at each SSD for the go unit (see Figure 7a). Cancel time was the difference between the time of modulation of the go unit and SSRT, with negative and positive differences measuring times before or after SSRT, respectively. We repeated this procedure 500 times and calculated the average cancel time at each SSD as the measure of the go unit cancel time; this was then statistically comparable to the physiological data.

Because activation in the stop unit remained at zero until $\text{SSD} + D_{\text{stop}}$ (the stop unit delay), the criteria used for determining the cancel time of the go unit were not appropriate because there was no noise in the activation.[1] Therefore, we calculated the time at which the t value of the stop unit activation became significantly greater than zero. To do this, we successively entered the stop unit activation into a t test with $df = N - 1$, where N was equal to the number of units entered into the t test. Cancel time was determined by the time at which it became significant at the $p < .05$ level, provided the difference remained significant for at least 50 ms.

Model Specifics

The independent race model is concerned with the finish times of the go and stop processes and does not specify the way in which the finish times are achieved. Here, we extend the race model by formally describing the processes by which the finish times arise. Our approach was to analyze the performance of a family of simple networks consisting of a go unit and a stop unit that are each noisy accumulators. These accumulators race toward a common threshold. Depending on the network architecture, rules about which unit reached its threshold first determine whether a stop-signal trial is signal inhibit or signal respond. Our main comparison was between models with independent versus interactive network architectures. In the independent model, the activation of one unit has no effect on the activation of the second unit. This model architecture is one particular instantiation of the independent race model of Logan and Cowan (1984; see also Hanes & Carpenter, 1999). We chose this instantiation because it allows a direct comparison with the interactive race model in which the go and stop units are mutually inhibitory.

[1] Gaze-holding neurons are characterized by higher discharge rates during periods of visual fixation, hence the name fixation neurons. This attribute was not included in our modeling efforts, which focused on understanding how a stop process can interrupt a go process. Every model has boundaries, and we are confident that these efforts could be extended to account for the activity of fixation neurons during fixation given prior work demonstrating conditions for stability of mutual inhibitory networks (e.g., Machens, Romo, & Brody, 2005; Wong & Wang, 2006). In order to directly compare the interactive race model with the independent model, we attempted to keep everything about those two models as equivalent as possible. Indeed, this is a recommended practice in computational model comparison. If the competing models differ along a large number of dimensions, it is difficult to determine which aspects of competing models are responsible for the success or failure of a model. By keeping the competitors as comparable as possible, we are able to explain why the interactive race model provides a better account of both neurophysiology and behavior than the independent race model. A more thoroughly fleshed out neural instantiation of the interactive race model would need to include a number of details that we did not include in our modeling efforts.

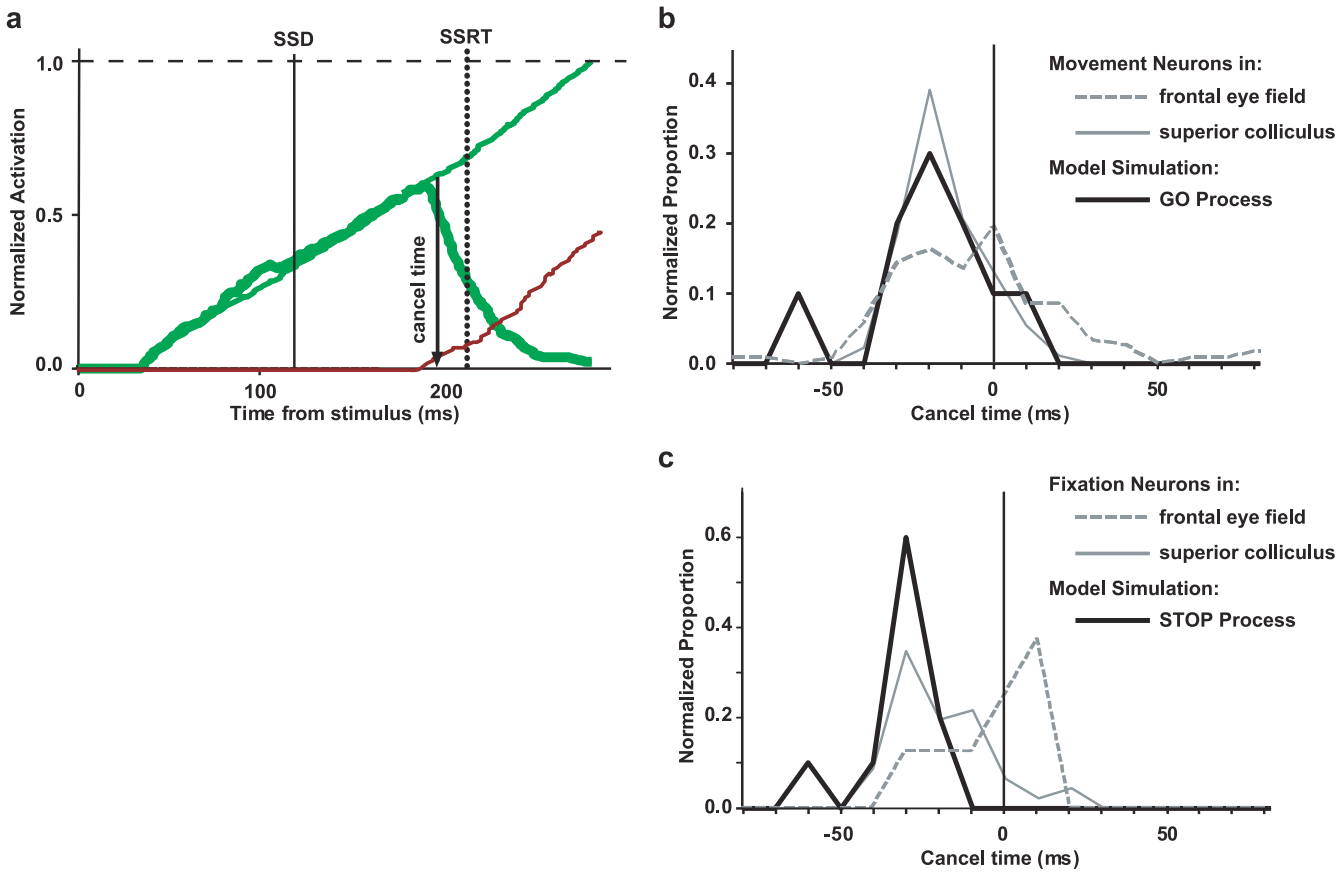

Figure 7. Cancel time. a: One run of the simulation with fewer trials to mimic the number of trials recorded for a typical neuron. Average normalized go unit (green) and stop unit (red) activation on signal-inhibit (thick solid line) and latency-matched no-stop-signal trials (thin solid lines) with stop-signal delay (SSD) and stop-signal reaction time (SSRT) indicated. Cancel time is indicated by the downward arrow. b: Probability distribution of cancel times of the go unit in the optimized interactive race model ($N = 10$) compared with the probability distributions of cancel times measured for movement-related neurons in frontal eye field ($N = 119$) and superior colliculus ($N = 92$). c: Probability distribution of cancel times of the stop unit in the optimized interactive race model ($N = 10$) compared with the probability distributions of cancel times measured for movement-related neurons in frontal eye field ($N = 8$) and superior colliculus ($N = 46$).

Formally, the go and stop units accumulate activation according to the following stochastic differential equations (Usher & McClelland, 2001):

$$da_{go}(t) = \frac{dt}{\tau}[\mu_{go} - k \cdot a_{go}(t) - \beta_{stop} \cdot a_{stop}(t)] + \sqrt{\frac{dt}{\tau}}\xi_{go}; \quad (1)$$

$$da_{stop}(t) = \frac{dt}{\tau}[\mu_{stop} - k \cdot a_{stop}(t) - \beta_{go} \cdot a_{go}(t)] + \sqrt{\frac{dt}{\tau}}\xi_{stop}. \quad (2)$$

These equations specify the change in unit activation (da_{go} and da_{stop}) within a time step dt (note that dt/τ was set equal to 1). The mean growth rates of the go and stop units are given by the μ_{go} and μ_{stop} parameters, respectively. ξ is a Gaussian noise term with a mean of zero and a variance of σ^2_{go} or σ^2_{stop}. The leakage parameter, k, prevents the activation from increasing without bound. Our initial investigation determined that the leakage term could be set to zero. However, we discuss a version of the model with leakage but with no inhibition below.

The key element of the interactive race model was the amount of inhibition between the go and stop units specified by the β parameters. β_{go} specifies the inhibition of the go unit on the stop unit, and β_{stop} specifies the inhibition of the stop unit on the go unit. The amounts of inhibition depend on the instantaneous activation levels (a_{go} or a_{stop}), causing a unit with a low activation to have a small inhibitory effect on the other unit. Together, these parameters specify the increment of activation on each time step toward a response threshold. The response threshold was fixed at 1,000 for both units, but other values produce the same results because the other parameters scale accordingly. The activations were rectified to be greater than zero as specified in the model of Usher and McClelland (2001).

Other parameters in the models capture the times for other stages of processing. Stimulus encoding that occurs before go unit and stop unit activation was instantiated as a constant delay between stimulus presentation and activation in the go unit and between stop-signal presentation and activation in the stop unit.

Thus, each model unit activation is set to zero until stimulus encoding is complete. The stop unit delay (D_{stop}) is a free parameter of the model. We chose to constrain the go unit delay (D_{go}) separately for each monkey on the basis of the average time at which the sample of movement neurons began to increase discharge rate in response to target presentation. We wanted to compare go unit activity with movement neuron activity to evaluate the neural predictions of the models, so it was necessary for the activation to start at times that were constrained by the neural data. The average was derived from the activity of 10 trials, sorted by RT to control for RT-dependent differences in activation growth, aligned on target presentation (Hanes & Schall, 1996). For this study, the beginning of activation was the latest time meeting the following criteria: (a) Spike density values increased significantly according to a Spearman correlation over an interval ranging from −20 ms to 20 ms, (b) the spike density at that time was less than the midpoint between the baseline and the threshold measured in the interval 200 ms before target onset, and (c) the correlation remained nonsignificant for 200 ms before the time. In this way, D_{go} values were set to 80 ms and 35 ms for Monkeys A and C, respectively. Finally, we assumed a ballistic interval of 10 ms preceding the eye movement after the go unit accumulated activation to its threshold ($go_{ballistic}$). This estimate corresponds to the time prior to a saccade during which the omnipause neurons cease firing and thus allows the eyes to move (e.g., Everling et al., 1998; Scudder et al., 2002). The time between the onset of activation of the go unit and that activation hitting threshold is defined as go_{RT}. Consequently, the simulated go unit finish time for each trial (rt_{go}) is equal to $D_{go} + go_{RT} + go_{ballistic}$.

We formulated independent and interactive models within the above general architecture. By constraining particular parameters, we defined special cases that could be compared quantitatively with a more general model using hierarchical model testing techniques. In particular, the presence or absence of inhibition between the go and stop units created models with two distinct architectures. In the *independent race* architecture (see Figure 5a), inhibition between the go and stop units was absent ($\beta_{go} = \beta_{stop} = 0$), so the outcome of a trial was dictated by a simple race of the go and stop units to their respective thresholds. If the go unit reached threshold first so that $rt_{go} < rt_{stop} + SSD$, then a signal-respond trial was produced; if the stop unit reached threshold first so that $rt_{stop} + SSD < rt_{go}$, then a signal-inhibit trial occurred.

In the *interactive race* architecture (see Figure 6a), the stop and go units were mutually inhibitory, so a signal-respond trial occurred if the go unit reached threshold, and a signal-inhibit trial occurred if the stop unit prevented the go unit from reaching threshold. Models with alternative architectures were assessed on their ability to satisfy the converging constraints of accounting for the essential characteristics of both behavioral and neural data obtained from monkeys performing the saccade stop-signal task (Hanes et al., 1998; Paré & Hanes, 2003).[2]

Models were fit quantitatively to the behavioral data with a search routine that found parameters that optimized the fit between model predictions and observed data. Following recommendations by Ratcliff and Tuerlinckx (2002), we optimized the fit of each model to the behavioral data by minimizing a Pearson chi-square statistic,

$$\chi^2 = \sum_i \sum_j \frac{(o_{ij} - p_{ij})^2}{p_{ij}}. \quad (3)$$

The first summation over i indexes the various SSD conditions and the no-stop-signal condition. Within each condition, a particular observation or prediction can fall into one of several bins. The o_{ij} and p_{ij} tally the number of observations and predictions that fall in each bin, respectively, where each bin is indexed by j. On no-stop-signal trials, an observation could fall into one of five RT bins (quintiles) in the RT distribution; a model prediction could also fall into one of those same bins. On stop-signal trials, for each SSD, an observation could fall either into one of five RT bins on signal-respond trials or into an inhibition bin on signal-inhibit trials; similarly, a model prediction could also fall into one of those same bins. If there were fewer than 40 RTs observed in a distribution at a particular SSD, then the entire distribution was considered a single bin, which effectively counted only the number of signal-respond versus signal-inhibit trials. We chose this particular fit statistic because it allowed us to fit quantitatively and simultaneously both the RT distributions on no-stop-signal and signal-respond trials and the proportion of signal-inhibit trials that make up the inhibition function.

The use of a chi-square statistic allowed us to evaluate statistically the fits of special cases of models with particular parameters constrained using the logic of hierarchical model testing. Testing special cases allowed us to evaluate rigorously what aspects of a model are necessary to account for the observations and how different parameter settings trade off against one another. For example, we could set a parameter to zero or set a parameter equal to another parameter. In both cases, we have constrained one parameter of the more general model. If $\chi^2_{general}$ is the fit of the general model with N parameters and $\chi^2_{special}$ is the fit of a special case with $N - M$ parameters (where M is the number of constrained parameters), then the special case is judged to fit significantly worse than the general model if the difference in chi-squares ($\chi^2_{general} - \chi^2_{special}$) is greater than the critical chi-square with M degrees of freedom at a .05 alpha level. Statistical model comparison techniques such as this are now regularly used to evaluate cognitive models (Ratcliff & Tuerlinckx, 2002).

Model parameters were optimized by minimizing chi-square using the subplex algorithm (Rowan, 1990), a more efficient variant of the Nelder–Mead simplex method (Nelder & Mead, 1965) that is well suited for optimizing stochastic models (Bogacz & Cohen, 2002). With stochastic simulation models, hill-climbing algorithms such as subplex can often settle in a local minimum rather than the global minimum. This was mitigated by starting the hill-climbing algorithm at a minimum of 40 different starting points for each model fit conducted.[3]

[2] Working simulations of the independent race model and the interactive race model can be found at http://www.psy.vanderbilt.edu/faculty/palmeri/psyrev07_model/model.m

[3] Given the time-intensive nature of these parameter optimizations, the simulations were run on a near-supercomputer computing cluster supported by the Vanderbilt Advanced Center for Computing for Research and Education.

Modeling Results

We first evaluate the ability of the independent race architecture to account for the observed data. It is necessary to establish the baseline performance of this architecture to relate our findings to the large literature applying the abstract independent race model to behavior obtained in the stop-signal task (Band et al., 2003; Logan, 1994). Furthermore, viable alternative model architectures must account for the behavioral data at least as well as the independent race model because this race model has been tested extensively and has demonstrated consistent success in describing countermanding behavior for over 20 years.

The independent race architecture predicted very well the inhibition function and the RT distributions on no-stop-signal and signal-respond trials (see Figures 5c and 5d). The SSRT, best-fitting parameters, and chi-square goodness of fit for the independent architecture are given in Table 1. In particular, note how each predicted RT distribution shares a common minimum and how the maximum RT increases with increasing SSD, as was observed in the behavioral data. This fanning pattern of RT distributions is a characteristic feature of RT data in the stop-signal task (e.g., Osman et al., 1986). In addition, the independent race model was able to account for the SSRT. Although this model accounted for the behavioral data quite well, it failed altogether in accounting for the neural data. This architecture cannot account for the modulation of movement-related neurons because nothing interrupts the go unit on stop-signal trials in which the stop unit reaches threshold first, that is, signal-inhibit trials (see Figures 5b and 5f). In signal-inhibit trials, the go unit reaches the threshold, albeit after the stop process. Mechanistically, if the go unit reaches its threshold, a saccade must be initiated, but neither human nor monkey subjects make such erroneous saccades on signal-inhibit trials. Of course, in its original formulation, the race model stated that if the stop process finished first, the trial was classified as signal inhibit because the go process was not permitted to finish, but this formalism did not specify a mechanism. The goal of our modeling effort was to determine a mechanism that could shut off the go process so that it does not reach threshold on signal-inhibit trials but preserve the essential characteristics of stop-signal task performance.

One way to turn off the go process is to make the stop unit inhibit the go unit. Thus, the interactive race architecture is defined by mutual inhibition between the go and stop units that allows the

Table 1
Best-Fitting Model Parameters for Monkeys A and C

Parameter	Independent	Interactive	$D_{stop} = 0$	$\mu_{stop} = \mu_{go}$	$\beta_{stop} = \beta_{go}$	$\beta_{stop} = \beta_{go}$, $\mu_{stop} = \mu_{go}$
			Monkey A			
μ_{go}	5.09	5.08	5.18	5.08	5.14	2.26
σ_{go}	26.38	26.24	26.42	26.24	26.27	31.82
μ_{stop}	50.24	5.07	25.96	**5.08**	33.68	**2.26**
σ_{stop}	40.17	26.34	21.30	**26.24**	40.47	**31.82**
β_{go}	**0.000**	0.005	0.000	0.005	0.024	0.009
β_{stop}	**0.000**	0.111	0.003	0.113	**0.024**	**0.009**
D_{stop}	51	51	**0**	51	51	51
χ^2	128.80	120.94	150.61	120.99	125.60	2684.90
Avg go cancel time (ms)	—	−20	−28	−20	−18	−111
Avg stop cancel time (ms)	—	−31	−81	−31	−33	−189
$Stop_{interrupt}$ (ms)	—	23	63	23	27	112
$SSRT_{model}$ (ms)	80	82	76	82	81	234
			Monkey C			
μ_{go}	4.64	4.63	4.59	4.63	4.60	1.16
σ_{go}	20.26	20.43	21.11	20.42	20.55	48.55
μ_{stop}	17.67	4.62	10.14	**4.63**	29.73	**1.16**
σ_{stop}	15.58	20.41	14.95	**20.42**	23.11	**48.55**
β_{go}	**0.000**	0.010	0.013	0.010	0.023	12.586
β_{stop}	**0.000**	0.434	0.029	0.435	**0.023**	**12.586**
D_{stop}	29	67	**0**	67	62	31
χ^2	57.24	50.64	139.97	50.65	53.56	1690.70
Avg go cancel time (ms)	—	−10	−44	−10	−2	−15
Avg stop cancel time (ms)	—	−19	−89	−19	−32	−162
$Stop_{interrupt}$ (ms)	—	21	49	21	38	169
$SSRT_{model}$ (ms)	97	94	91	93	95	219

Note. Chi-square values measure the goodness of fit between the simulated and observed reaction time (RT) distributions on no-stop-signal trials, RT distributions on signal-respond trials at each stop-signal delay (SSD), and proportion of signal-inhibit trials (one-inhibition function) at each SSD. The go and stop cancel times are the times relative to stop-signal RT (SSRT) at which the go or stop activation functions on signal-inhibit and latency-matched no-stop-signal trials become significantly different. $Stop_{interrupt}$ is the time interval between the start of the stop process (SSD + D_{stop}) and the time it takes to cancel the movement, the cancel time. Cancel time and $Stop_{interrupt}$ were averaged across SSDs. Note that there are no cancel times or $Stop_{interrupt}$ values for the independent model because nothing interrupts the go unit in this model to produce these values. $SSRT_{model}$ is calculated by the integration method. Values in boldface type indicate constrained parameters (see the text for more details).

inhibition from the stop unit to prevent the go unit from reaching threshold and triggering a movement on signal-inhibit trials. We found that with suitable parameters, the interactive race architecture could account for the behavioral data just as well as the independent race architecture (see Figure 6 and Table 1). As with the independent architecture, the inhibition function, the fanning of the RT distributions on no-stop-signal and signal-respond trials, and the SSRTs are all well predicted. In contrast to the independent race architecture, this interactive race architecture can account for the pattern and timing of neural modulation. Specifically, go unit activation reaches threshold on no-stop-signal trials and signal-respond trials, but it is reduced substantially on signal-inhibit trials before the behavioral measure of SSRT (see Figures 6e and 6f). Thus, the activation of the go unit corresponds precisely to the pattern of activation of movement neurons in frontal eye field (Hanes et al., 1998) and superior colliculus (Paré & Hanes, 2003) during this task (see Figure 4).

The reduction of activation in the go unit resulted from inhibition from the stop unit on the go unit (β_{stop}). On these signal-inhibit trials, the stop unit became active just before SSRT had elapsed. The reduction of activation of the go unit did not happen on signal-respond trials because the stop unit failed to become active before the SSRT. This modulation of the stop unit before SSRT on signal-inhibit trials, but not signal-respond trials, corresponds to the modulation of fixation neurons in frontal eye field (Hanes et al., 1998) and superior colliculus (Paré & Hanes, 2003; see Figure 4). Thus, the interactive race architecture parameterized to fit the performance accomplishes this with go and stop unit activation having the presence and timing of modulation corresponding to what has been observed in movement neurons and fixation neurons.

This qualitative conclusion can be evaluated quantitatively by examining the cancel time of the go and stop units—the time at which modulation occurred on signal-inhibit trials. Analyses of neural data show an effect of noise on estimates of cancel time: the noisier the activation functions, the later the time of the difference measured by the statistical criteria (Hanes et al., 1998). The model unit activation functions have considerably less variability than neural spike density functions owing to the tremendously large number of trials that are simulated in order to generate stable model predictions. This fact has implications for the calculation of cancel time, so we first needed to subsample the activation of the go and stop units in a manner comparable to the original sampling of neural activity in the physiology experiments. We did so using the following procedure: We simulated the model with a range of 20–50 trials at each SSD to mimic the number of trials typically obtained in the single-cell physiology experiments that provided our data. We then calculated the average activation functions of the simulated trials and determined the time of modulation at each SSD for the go and stop units (see Figure 7a). We calculated cancel time as the difference between the time of modulation in the go and stop units and SSRT; negative and positive cancel times are instances in which the modulation occurred before or after SSRT, respectively. We repeated this procedure 500 times and took the average cancel time at each SSD to obtain an estimate of the go unit and stop unit cancel time that was statistically comparable to the physiological data. We found that the distribution of go unit and stop unit cancel times derived from the interactive race architecture overlaps with the distribution of cancel times measured from neurons in frontal eye field (Hanes et al., 1998) and superior colliculus (Paré & Hanes, 2003; see Figures 7b and 7c). Thus, quantitatively, the interactive race model yields a predicted distribution of cancel times that are very similar to the observed distribution of neural cancel times. The data from fixation neurons in the superior colliculus are more reliable than those obtained from frontal eye field.

Model Exploration

The interactive race architecture prevails over the independent race architecture in that it can account for both the behavioral and the neural data. We wanted to test the boundary conditions of the interactive model to understand better which aspects of the model are necessary to account for the behavior and neurophysiology. Hierarchical model testing showed that two features of the interactive architecture were necessary to account for both behavioral and neural data. First, the interaction between the stop and go units must occur well after presentation of the stop signal ($D_{stop} \gg 0$). Second, the inhibition of the stop process on the go process must be very potent. These requirements of the interactive race architecture are considered in turn.

First, to investigate the necessity of a delayed interaction, we set $D_{stop} = 0$, activating the stop unit immediately with stop-signal presentation, and allowed all other parameters to vary freely. With $D_{stop} = 0$, the fit to the behavioral data was significantly worse than that of the general interactive model for both monkeys ($ps < .001$; see Table 1). This occurred because, without a delay, the stop unit interacts with the go unit for a prolonged period of time, producing longer than observed simulated RTs on signal-respond trials and a premature cancel time. Thus, in order for the interactive race model to account for stop-signal performance, activation of the stop unit must be delayed for a substantial amount of time after the stop signal occurs (i.e., $D_{stop} \gg 0$).

The prolonged RTs are characteristic of violations of the independence assumption of the race model (Band et al., 2003). Note, however, that these violations appear uniformly across all SSDs, whereas violations of the race model from observable data occur at the shortest SSDs (Boucher et al., in press; Hanes & Carpenter, 1999; Özyurt, Colonius, & Arndt, 2003). Only a small number of trials contribute to the mean RT at these short SSDs, and these outlier signal-respond RTs are often much longer than an average saccadic RT. We believe these responses occur on rare trials when subjects successfully inhibit the movement and then after a typical RT produce the movement anyway, so they are outside the race model framework.

Second, to investigate the necessity of potent inhibition of the go unit by the stop unit, we constrained either the growth rate such that $\mu_{stop} = \mu_{go}$ and $\sigma_{stop} = \sigma_{go}$ or the inhibition such that $\beta_{stop} = \beta_{go}$. We found a clear trade-off between μ_{stop}, μ_{go} and β_{stop}, β_{go} (see Table 1). If $\mu_{stop} = \mu_{go}$, then the optimized β values were always $\beta_{stop} \gg \beta_{go}$. If $\beta_{stop} = \beta_{go,}$ then the optimized μ values were always $\mu_{stop} \gg \mu_{go}$. These results occurred because the potency of the stop unit is a function of both the magnitude of the inhibition of the stop unit on the go unit (β_{stop}) and the rate of growth of the stop unit (μ_{stop}, σ_{stop}). To explore this trade-off, we equated simultaneously the growth rate ($\mu_{stop} = \mu_{go}$ and $\sigma_{stop} = \sigma_{go}$) and the inhibition ($\beta_{stop} = \beta_{go}$). This architecture fit the observed data significantly worse than the general interactive

model ($ps < .001$). It failed because the trade-off between μ_{stop} and β_{stop} could not occur. For example, if μ and β were too high, the go unit reached threshold too early before the stop unit, resulting in premature RTs, and the inhibition of the go unit on the stop unit was too potent for the stop unit to accumulate sufficiently, resulting in a lack of sensitivity to SSD. Conversely, if μ and β were too low, the go unit reached threshold too late to account for the RT data, and the stop unit was too weak to interrupt the go unit sufficiently often enough. Thus, for the interactive race model to account for stop-signal performance, the stop unit must inhibit the go unit much more than vice versa.

General Discussion

This theoretical investigation furnishes a deeper understanding of SSRT that justifies its use in clinical and developmental studies, reveals a previously unappreciated characteristic of the neural circuitry responsible for gaze control, and makes new predictions about the stages of processing making up SSRT, which are amenable to experimental manipulation. More generally, our work exemplifies an approach to mapping cognitive processes onto neural process through quantitative testing of stochastic models of behavior constrained by the characteristics of the neural mechanisms instantiating the cognitive process. This approach can readily be used in other domains to assess competing models. We consider each of these points in turn.

What Does SSRT Measure?

Using hierarchical model testing, we determined that the interactive race model alone could satisfy both behavioral and physiological constraints if and only if stop unit inhibition was potent and delayed. This result solves the paradox: Interacting neural units can produce behavior that appears to be the outcome of a race between independent processes because the processes remain independent for most of their duration. The stop process interacts with the go process only for a very brief period of time. Thus, an implication of the interactive race model is that stopping is a multistage process consisting of an encoding stage during which no interaction with response preparation occurs and a brief interruption stage during which response preparation is inhibited.

The duration of SSRT is occupied by the following events. First, during stimulus encoding, the stop unit does not influence the go unit, satisfying the independence premise of the original race model. This interval, called D_{stop}, was between 51 and 67 ms, a range that corresponds to the typical latency of visual responses in frontal eye field and superior colliculus (e.g., Pouget, Emeric, Stuphorn, Reis, & Schall, 2005); the standard deviation of this latency is rather small, measured at 10–20 ms. Second, once activated, the stop unit potently and rapidly inhibits the go unit. This interval from activation of the stop unit (SSD + D_{stop}) until interruption of go unit accumulation (cancel time) can be called $stop_{interrupt}$. In the optimized interactive race model $stop_{interrupt}$ was 22 ms across SSDs and monkeys. Because this value approximates synaptic integration time, $stop_{interrupt}$ can be considered effectively instantaneous. Finally, the race model calculation of SSRT includes the ballistic interval preceding initiation of the movement (Logan & Cowan, 1984), which we refer to as $go_{ballistic}$. For saccade production, $go_{ballistic}$ can be considered to be the time at which inhibition of omnipause neurons on medium-lead burst neurons is released; this time has been measured at ~10 ms with 1-ms *SD* (e.g., Everling et al., 1998).

Thus, the SSRT measured using standard techniques is the sum of these intervals:

$$\mathrm{SSRT} = D_{stop} + stop_{interrupt} + go_{ballistic}. \quad (4)$$

$Go_{ballistic}$ is very brief, even approaching zero for some manual responses (e.g., De Jong et al., 1990). $Stop_{interrupt}$ is very brief, because the inhibition of the stop unit on the go unit is very potent. Thus, most of SSRT is occupied by D_{stop}, during which the go unit is uninfluenced by the stop unit. Thus, the go and stop processes are independent for most of their durations. Consequently, SSRT estimates from the independent race model provide valid measures of the process by which movement preparation is interrupted. This justifies the use of SSRT to measure inhibitory ability in studies of lifespan development, individual differences, and psychopathology.

Possible Neural Mechanisms of Stopping

We have identified stop and go units with fixation and movement neurons, respectively, in frontal eye fields and superior colliculus. This bridging proposition is a strong claim that must be defended against alternatives. On the computational side, Logan and Cowan (1984) distinguished between two ways to stop an action: by inhibiting it directly in the motor system or by removing its input, which could involve deleting the goal of moving or deleting the stimulus that drives the movement. Our interactive race model instantiates the first alternative; fixation cells inhibit movement cells in the motor system. The second alternative is possible computationally; the interactive race model does not rule it out. On the neural side, recent functional magnetic resonance imaging studies have described a broad network of structures in the brain that are involved in performing the stop-signal task with saccades (Curtis, Cole, Rao, & D'Esposito, 2005) as well as with manual responses (Aron & Poldrack, 2006; Li, Huang, Constable, & Sinha 2006). There is significant activation outside the frontal eye fields and superior colliculus, and it is logically possible that stop and go units are located in these other structures. It is logically possible that countermanding is accomplished by deleting or inhibiting the stimuli or the goals that drive the movement cells in frontal eye fields and superior colliculus.

Some neurophysiological data challenge the idea that the interaction between stop and go units could occur between cells other than movement and fixation neurons. The specific anatomical inputs to movement and fixation neurons are not known, but the most plausible possibility must be the visually responsive neurons in the frontal eye field and superior colliculus as well as from other areas in parietal and frontal cortex. The race between go and stop processes cannot be instantiated by these other neurons because most of the visual neurons in frontal eye field and superior colliculus do not modulate in stop-signal trials, and the few that do modulate after SSRT (Hanes et al., 1998; Paré & Hanes, 2003). Likewise, preliminary reports have shown that neurons in the supplementary eye field (J. W. Brown, Stuphorn, & Schall, 2001; Schall, Stuphorn, & Brown, 2002) and the lateral intraparietal area (Brunamonti & Paré, 2005) do not modulate in stop-signal trials either. Furthermore, the visual latencies of neurons in prefrontal

cortex approach or exceed the 100 ms of the SSRT (reviewed by Schall, 1991), so they cannot act in time to modulate the activity of movement neurons. Finally, physiological data show that the ballistic phase for saccadic eye movements is approximately 10 ms, which is the time required for movement processes in the brainstem to operate. If the interaction of stop and go units was prior to movement and fixation cells, the ballistic phase would have to be longer than this 10-ms interval (i.e., the time required to activate movement neurons would also contribute to the duration of the ballistic phase).

Although we believe the neurophysiological data are consistent with our identification of stop and go units with fixation and movement neurons, respectively, we tried to address the issue computationally by developing models in which the interaction between stop and go units operated on the inputs to fixation and movement neurons. We assumed the go and stop processes rose to a threshold, and when one of them reached it, the input to both units was turned off. In one version of the model, noise turned off with the input. In the other version of the model, noise remained on continuously, as it had in our other models, even after the input turned off. These models did not fit the behavioral data very well. They accounted for the inhibition function and the go RTs on no-stop-signal trials, but they produced signal-respond RTs that were longer than the observed signal-respond RTs and even longer than the predicted no-stop-signal RTs. Because there was no direct inhibition between stop and go units to produce the modulation of go unit activity on successful stop trials, a leakage parameter greater than zero was included that served to bring the go unit activation back to zero after the input was cut off. This produced a modulation in go unit activation on successful stop trials, but the modulation occurred after SSRT in both models (i.e., cancel time was after SSRT). Thus, neither the computational analysis nor the physiological data are consistent with the alternative hypothesis that stop and go units can be identified with neurons that provide inputs to fixation and movement cells in frontal eye fields and superior colliculus.

Asymmetric Inhibition

As reviewed above, the critical neural events responsible for saccade preparation occur in a distributed circuit including frontal eye field, superior colliculus, basal ganglia, cerebellum, thalamus, and brainstem. Saccade initiation is inhibited at several levels including the omnipause neuron inhibition on the excitatory burst neurons (Scudder, Fuchs, & Langer, 1988), inhibition within the superior colliculus (Meredith & Ramoa, 1998; Munoz & Istvan, 1998; Munoz & Wurtz, 1993a, 1995), and tonic inhibition on the superior colliculus from the substantia nigra pars reticulata, which is itself inhibited before saccades by the oculomotor region of the caudate nucleus (Hikosaka et al., 2000) that is innervated by frontal eye field (Parthasarathy, Schall, & Graybiel, 1992). The architecture of the interactive race model that fit behavior required asymmetric inhibition between the go unit and the stop unit. We believe this implies that during the period described by this model, gaze-holding neurons inhibit gaze-shifting neurons much more than vice versa. To our knowledge, this asymmetric inhibition has not been described before or anticipated by previous physiological or anatomical studies, so it is a novel prediction from our interactive race model. It is certainly not an explicit feature of current models of saccade generation (for a review, see Girard & Berthoz, 2005). Further neurophysiological and neuroanatomical studies are necessary to identify how such an asymmetry can arise.

Stages of Stopping

A strong implication of the interactive race model is that SSRT consists of an initial period during which no influence is exerted on the go process followed by a subsequent, brief period of potent interaction—if and only if the movement is inhibited. We refer to the two stages as *encoding* and *interruption*. Although the RT preceding overt movements has been regarded as consisting of stages (e.g., Sternberg, 2001), most previous descriptions of SSRT have not emphasized this characteristic (but see van den Wildenberg & van der Molen, 2004b).

Selective influence is one of the strongest lines of evidence for stages of processing, and several studies have shown that certain variables may selectively influence different stages of SSRT. Several researchers manipulated the discriminability of the stop signal and found variations in SSRT, suggesting selective influence on the encoding stage (Asrress & Carpenter, 2001; Cabel et al., 2000; Cavina-Pratessi, Bricolo, Prior, & Marzi, 2001; van den Wildenberg & van der Molen, 2004a). Other researchers found prolonged SSRT in response-incompatible conditions of flanker and Stroop tasks that require inhibition of prepotent responses, suggesting selective influence on the interruption stage (Kramer et al., 1994; Ridderinkhof et al., 1999; Verbruggen, Liefooghe, & Vandierendonck, 2004). Further research with our interactive race model is necessary to assess the validity of these suggestions.

Stopping Eyes and Stopping Hands

We believe the results we have obtained generalize to stop-signal conditions using other effectors, such as key presses[4] or vocalizations, and other stop-signal modalities, such as acoustic. Several lines of evidence lead to this conclusion.

First, the quality of performance on the saccade stop-signal task corresponds precisely to that on manual stop-signal tasks (e.g., Hanes & Schall, 1995; Logan & Irwin, 2000). We have investigated performance of combined eye and hand stopping (Boucher et al., in press). In a stop-signal task, subjects were instructed to initiate but occasionally inhibit eye, hand, or eye + hand movements in response to a color-coded foveal stop signal. Performance on this task could be accounted for very well by the race model because uncertainty in the mapping of stop signal onto effector did not produce performance data qualitatively different from any previous stop-signal study. Although SSRT was shorter for eye movements than for hand movements, SSRT did not vary with knowledge about which movement to cancel.

Second, the high-level control of saccades appears indistinguishable from the high-level control of manual movements. For example, when subjects are asked to generate a sequence of

[4] Manual stop-signal studies have provided evidence that these movements can be produced with lower velocity or force or even interrupted during execution (De Jong et al., 1990). In contrast, saccades produced in the stop-signal task are not slower or hypometric (Hanes & Schall, 1995). Such differences in effector will need to be recognized in generalizing this model.

saccades, the latency of the first saccade increases with the number of movements in the sequence (Inhoff, 1986; Zingale & Kowler, 1987), following the same pattern observed for speech and typing (Sternberg, Monselli, Knoll, & Wright, 1978). In addition, the influence of foreperiod on RT is the same for movements of the eyes (Findlay, 1981; Hanes, Tu, & Schall, 1992) as it is for movements of the limbs (e.g., Niemi & Näätänen, 1981).

Third, forelimb movements are initiated when movement-related activity in primary motor cortex reaches a threshold, and variability in RT comes from variability in growth (Lecas, Requin, Anger, & Vitton, 1986), as observed for saccades (Hanes & Schall, 1996). Similar evidence has been presented measuring the amplitude of the lateralized readiness potential, a scalp potential concomitant of movement preparation (Gratton, Coles, Sirevaag, & Eriksen, 1988).

Fourth, it is well known that the initiation of forelimb and other skeletomotor movements is under inhibitory control primarily through the internal segment of the globus pallidus of the basal ganglia (e.g., Alexander, DeLong, & Strick, 1986). Unfortunately to date, no single-unit recordings have been conducted in the skeletal motor system in monkeys performing a manual stopping task.

Neurally Constrained Cognitive Modeling

Our work exemplifies an emerging approach of mapping cognitive processes onto neural processes through converging evidence from formal models and neurophysiology (see also Bundesen, Habekost, & Kyllingsbaek, 2005; Mazurek, Roitman, Ditterich, & Shadlen, 2003; Ratcliff, Cherian, & Segraves, 2003). This effort is distinct in several respects from other kinds or levels of models. Some models are built from a reasonably complete understanding of the details of the neural circuitry in a particular brain region, such as models of the brainstem saccade generator (e.g., Cannon & Robinson, 1985; Chun & Robinson, 1978; Scudder, 1988). Although these models provide detailed accounts of neural properties, their scope is rather limited. Other models simulate circuits involving multiple brain areas engaged in more complex behavior (e.g., J. W. Brown et al., 2004; Hamker, 2005; Mitchell & Zipser, 2001; Optican, 1995; Quaia, Lefevre, & Optican, 1999; Zipser & Andersen, 1988). Although models such as this respect many neural details, they lack stochastic elements capable of accounting for natural variability in performance. Finally, cognitive process models account for a broad range of complex behaviors with no commitment to neural architecture (e.g., S. Brown & Heathcote, 2005; Bundesen, 1990; Busemeyer & Townsend, 1993; Carpenter & Williams, 1995; Hanes & Carpenter, 1999; Link, 1975; Logan & Cowan, 1984; Logan & Gordon, 2001; Nosofsky & Palmeri, 1997; Ratcliff, 1978; Usher & McClelland, 2001).

Establishing useful mapping between cognitive models and neurophysiology is an essential next step for both fields. First, cognitive models can tell neurophysiologists what neural events are relevant to measure. For example, we mapped the activation of model units onto specific neuron firing rates changing over time. This is an explicit commitment to rate coding. We could have attempted to map the model unit activation onto the temporal pattern of action potentials or the information content of the spike train, but the race model can be instantiated most simply in the accumulator framework. Furthermore, the mapping between the race model and neural processes was established most directly through the timing of particular moments of strong neural modulation. This focus on specific events in the firing-rate records prevented overfitting the model to idiosyncrasies in the discharge patterns of neurons. However, it is conceivable that the mapping of neural events onto cognitive processes requires knowledge of the membrane properties of neurons and the intrinsic dynamics of complex neural circuits. In fact, natural bridges between models of hundreds of integrate-and-fire neurons with realistic receptor mechanisms and models of the form we used are being formulated (e.g., Wong & Wang, 2006). Ultimately, if proper mappings can be elucidated, then by articulating the computations performed, the cognitive models provide explanations of what neurons do.

Second, neurophysiology can help distinguish between alternative cognitive models. Cognitive models are often faced with the conundrum of *model identifiability* or *model mimicry* (e.g., Logan, 2004; Massaro, 1993). Several approaches to deal with this issue have been proposed (e.g., Anderson, 1990; Myung, Pitt, & Kim, 2005; Oaksford & Chater, 1998; Shepard, 1994; Townsend & Ashby, 1983), but we suggest that beyond these approaches, the characteristics of neural activity can contribute decisively to resolving model mimicry (Hanes & Schall, 1996). Our work demonstrates a case in which two models fit the behavioral data equally well. Therefore, appealing to parsimony would have demanded exclusion of the more complex interactive race model in favor of the simpler independent race model. However, mapping the pattern of neural modulation onto the model activation required the more complex model.

Conclusions

The goal of our research was to resolve a paradox in linking a cognitive model of countermanding performance to the neurophysiological events that produce the behavior: How can an independent race model account so well for behavior that is produced by systems of interacting neurons? Our resolution of this paradox was to propose an interactive race model, in which stop and go units are independent for much of their duration, but the stop unit can interrupt the go unit potently and briefly. The interactive race model presented here is based on a linking proposition that identifies the formal go and stop processes for saccade preparation with activation of movement and fixation neurons in the circuit that includes the frontal eye field and superior colliculus that implements the operations required to perform the stop-signal task. Through this linking proposition, the interactive race model explains what the neurons do and justifies the use of SSRT as a measure of inhibitory control in cognitive psychology, lifespan development, individual differences, and psychopathology. We anticipate that a coordinated analysis of behavioral and neural data with formal cognitive models to converge on underlying mechanisms will be useful in other domains such as visual attention, categorization, and executive control.

References

Akerfelt, A., Colonius, H., & Diederich, A. (2005). Visual-tactile saccadic inhibition. *Experimental Brain Research, 169,* 554–563.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organi-

zation of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience, 9,* 357–381.

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Arai, K., Keller, E. L., & Edelman, J. A. (1994). Two-dimensional neural network model of the primate saccadic system. *Neural Networks, 7,* 1115–1135.

Armstrong, I. T., & Munoz, D. P. (2003). Inhibitory control of eye movements during oculomotor countermanding in adults with attention-deficit hyperactivity disorder. *Experimental Brain Research, 152,* 444–452.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience, 26,* 2424–2433.

Asrress, K. N., & Carpenter, R. H. (2001). Saccadic countermanding: A comparison of central and peripheral stop signals. *Vision Research, 41,* 2645–2651.

Badcock, J. C., Michie, P. T., Johnson, L., & Combrinck, J. (2002). Acts of control in schizophrenia: Dissociating the components of inhibition. *Psychological Medicine, 32,* 287–297.

Band, G. P. H., van der Molen, M. W., & Logan, G. D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologia (Amst), 112,* 105–142.

Barkley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. *Psychological Bulletin, 121,* 65–94.

Becker, W., & Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vision Research, 19,* 976–983.

Bedard, A. C., Nichols, S., Barbosa, J. A., Schachar, R., Logan, G. D., & Tannock, R. (2002). The development of selective inhibitory control across the lifespan. *Developmental Neuropsychology, 21,* 93–111.

Bogacz, R., & Cohen, J. D. (2002). *Parameterization of connectionist models* (Tech. Rep. No. 1). Princeton, NJ: Princeton University, Center for the Study of Brain, Mind, and Behavior.

Boucher, L., Stuphorn, V., Logan, G. D., Schall, J. D., & Palmeri, T.(in press). Dissecting the stop process: Eye-hand coordination in a stop task. *Perception & Psychophysics.*

Bozis, A., & Moschovakis, A. (1998). Neural network simulations of the primate oculomotor system III. A one-dimensional, one-directional model of the superior colliculus. *Biological Cybernetics, 79,* 215–230.

Brindley, G. S. (1970). *Physiology of retina and visual pathways* (2nd ed.). Baltimore: Williams & Wilkins.

Brown, J. W., Bullock, D., & Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Networks, 17,* 471–510.

Brown, J. W., Stuphorn, V., & Schall, J. D. (2001). Reliability of macaque FEF but not SEF movement neurons predicting saccade initiation. [Abstract]. *Society for Neuroscience Abstracts Online, 575.9.* Retrieved from http://sfn.scholarone.com/itin2001/.

Brown, S., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review, 112,* 117–128.

Bruce, C. J., & Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *Journal of Neurophysiology, 53,* 603–635.

Brunamonti, E., & Paré M. (2005). Neural control of saccade production studied with the countermanding paradigm: Parietal cortex area LIP. [Abstract]. *Society for Neuroscience Abstracts Online, 166.18.* Retrieved from http://sfn.scholarone.com/itin2005/

Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97,* 523–547.

Bundesen, C., Habekost, T., & Kyllingsbaek, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review, 112,* 291–328.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459.

Buttner-Ennever, J. A., Horn, A. K., Henn, V., & Cohen, B. (1999). Projections from the superior colliculus motor map to omnipause neurons in monkey. *Journal of Comparative Neurology, 413,* 55–67.

Cabel, D. W., Armstrong, I. T., Reingold, E., & Munoz, D. P. (2000). Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Experimental Brain Research, 133,* 431–441.

Cannon, S. C., & Robinson, D. A. (1985). An improved neural-network model for the neural integrator of the oculomotor system: More realistic neuron behavior. *Biological Cybernetics, 53,* 93–108.

Carpenter, R. H. S. (1991). The visual origins of ocular motility. In J. R. Cronly-Dillon (Ed.), *Vision and visual function. Eye movements* (Vol. 8, pp. 1–10). London: Macmillan.

Carpenter, R. H. S., & Williams, M. L. (1995, September 7). Neural computation of log likelihood in control of saccadic eye movements. *Nature, 377,* 59–62.

Carter, J. D., Farrow, M., Silberstein, R. B., Stough, C., Tucker, A., & Pipingas, A. (2003). Assessing inhibitory control: A revised approach to the stop signal task. *Journal of Attention Disorders, 6,* 153–161.

Cavina-Pratessi, C., Bricolo, E., Prior, M., & Marzi, C. A. (2001). Redundancy gain in the stop-signal paradigm: Implications for the locus of coactivation in simple reaction time. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 932–941.

Choi, W. Y., & Guitton, D. (2006). Responses of collicular fixation neurons to gaze shift perturbations in head-unrestrained monkey reveal gaze feedback control. *Neuron, 50,* 491–505.

Chun, K. S., & Robinson, D. A. (1978). A model of quick phase generation in the vestibuloocular reflex. *Biological Cybernetics, 28,* 209–221.

Colonius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review, 97,* 309–312.

Colonius, H., Ozyurt, J., & Arndt, P. A. (2001). Countermanding saccades with auditory stop signals: Testing the race model. *Vision Research, 41,* 1951–1968.

Corneil, B. D., & Elsley, J. K. (2005). Countermanding eye-head gaze shifts in humans: Marching orders are delivered to the head first. *Journal of Neurophysiology, 94,* 883–895.

Curtis, C. E., Cole, M. W., Rao, V. Y., & D'Esposito, M. (2005). Canceling planned action: An FMRI study of countermanding saccades. *Cerebral Cortex, 15,* 1281–1289.

Das, S., Gandhi, N. J., & Keller, E. L. (1995). Open-loop simulations of the primate saccadic system using burst cell discharge from the superior colliculus. *Biological Cybernetics, 73,* 509–518.

De Jong, R., Coles, M. G. H., & Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 498–511.

De Jong, R., Coles, M. G. H., Logan, G. D., & Gratton, G. (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 164–182.

Dimitrov, M., Nakic, M., Elpern-Waxman, J., Granetz, J., O'Grady, J., Phipps, M., et al. (2003). Inhibitory attentional control in patients with frontal lobe damage. *Brain and Cognition, 52,* 258–270.

Dominey, P., & Arbib, M. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex, 2,* 153–175.

Dominey, P., Arbib, M., & Joseph, J. P. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience, 7,* 311–336.

Dorris, M. C., & Munoz, D. P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. *Journal of Neuroscience, 18,* 7015–7026.

Dorris, M. C., Paré, M., & Munoz, D. P. (1997). Neuronal activity in

monkey superior colliculus related to the initiation of saccadic eye movements. *Journal of Neuroscience, 17,* 8566–8579.

Eagle, D. M., & Robbins, T. W. (2003). Lesions of the medial prefrontal cortex or nucleus accumbens core do not impair inhibitory control in rats performing a stop-signal reaction time task. *Behavioural Brain Research, 146,* 131–144.

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: A dynamical model of saccade generation during reading. *Psychological Review, 112,* 777–813.

Everling, S., Paré, M., Dorris, M. C., & Munoz, D. P. (1998). Comparison of the discharge characteristics of brain stem omnipause neurons and superior colliculus fixation neurons in monkey: Implications for control of fixation and saccade behavior. *Journal of Neurophysiology, 79,* 511–528.

Feola, T. W., de Wit, H., & Richards, J. B. (2000). Effects of d-amphetamine and alcohol on a measure of behavioral inhibition in rats. *Behavioral Neuroscience, 114,* 838–848.

Findlay, J. M. (1981). Spatial and temporal factors in the predictive generation of saccadic eye movements. *Vision Research, 21,* 347–354.

Freedman, E. G., & Sparks, D. L. (1997). Activity of cells in the deeper layers of the superior colliculus of the rhesus monkey: Evidence for a gaze displacement command. *Journal of Neurophysiology, 78,* 1669–1690.

Gandhi, N. J., & Keller, E. L. (1999). Comparison of saccades perturbed by stimulation of the rostral superior colliculus, the caudal superior colliculus, and the omnipause neuron region. *Journal of Neurophysiology, 82,* 3236–3253.

Gauggel, S., Rieger, M., & Feghoff, T. A. (2004). Inhibition of ongoing responses in patients with Parkinson's disease. *Journal of Neurology, Neurosurgery & Psychiatry, 75,* 539–544.

Girard, B., & Berthoz, A. (2005). From brainstem to cortex: Computational models of saccade generation circuitry. *Progress in Neurobiology, 77,* 215–251.

Gratton, G., Coles, M. G. H., Sirevaag, E. J., & Eriksen, C. W. (1988). Pre- and poststimulus activation of response channels: A psychophysiological analysis. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 331–344.

Guitton, D., Bergeron, A., & Choi, W. Y. (2004). On the role of subcortical feedback mechanisms in the control of head-unrestrained gaze saccades. In W. C. Hall & A. Moschovakis (Eds.), *The superior colliculus: New approaches for studying sensorimotor integration* (pp. 241–275). Boca Raton, FL: CRC Press.

Hamker, F. H. (2005). The reentry hypothesis: The putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement. *Cerebral Cortex, 15,* 431–447.

Hanes, D. P., & Carpenter, R. H. S. (1999). Countermanding saccades in humans. *Vision Research, 39,* 2777–2791.

Hanes, D. P., Patterson, W. F., II, & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology, 79,* 817–834.

Hanes, D. P., & Schall, J. D. (1995). Countermanding saccades in macaques. *Visual Neuroscience, 12,* 929–937.

Hanes, D. P., & Schall, J. D. (1996, October 18). Neural control of voluntary movement initiation. *Science, 274,* 427–430.

Hanes, D. H., Tu, F. F., & Schall, J. D. (1992). Effects of temporal context (conditional probability) on saccade latency in macaque. *Society for Neuroscience Abstracts, 18,* 698.

Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiological Review, 80,* 953–978.

Hikosaka, O., & Wurtz, R. H. (1983a). Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. *Journal of Neurophysiology, 49,* 1230–1253.

Hikosaka, O., & Wurtz, R. H. (1983b). Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. *Journal of Neurophysiology, 49,* 1254–1267.

Hikosaka, O., & Wurtz, R. H. (1983c). Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *Journal of Neurophysiology, 49,* 1268–1284.

Inhoff, A. W. (1986). Preparing sequences of saccades under choice reaction conditions: Effects of sequence length and context. *Acta Psychologia, 61,* 211–228.

Kaneko, C. R. (2006). Saccade-related, long-lead burst neurons in the monkey rostral pons. *Journal of Neurophysiology, 95,* 979–994.

Kase, M., Miller, D. C., & Noda, H. (1980). Discharges of Purkinje cells and mossy fibers in the cerebellar vermis of the monkey during saccadic eye movements and fixation. *Journal of Physiology (London*), *300,* 539–555.

Kornylo, K., Dill, N., Saenz, M., & Krauzlis, R. J. (2003). Cancelling of pursuit and saccadic eye movements in humans and monkeys. *Journal of Neurophysiology, 89,* 2984–2999.

Kramer, A. F., Humphrey, D. G., Larish, J. F., Logan, G. D., & Strayer, D. L. (1994). Aging and inhibition: Beyond a unitary view of inhibitory processing in attention. *Psychology and Aging, 9,* 491–512.

Ladefoged, P., Silverstein, R., & Papcun, G. (1973). Interruptability of speech. *Journal of the Acoustical Society of America, 54,* 1105–1108.

Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology, 72,* 803–811.

Lauwereyns, J., Watanabe, K., Coe, B., & Hikosaka, O. (2002, July 25). A neural correlate of response bias in monkey caudate nucleus. *Nature, 418,* 413–417.

Lecas, J. C., Requin, J., Anger, C., & Vitton, N. (1986). Changes in neuronal activity of the monkey precentral cortex during preparation for movement. *Journal of Neurophysiology, 56,* 1680–1702.

Li, C. S., Huang, C., Constable, R. T., & Sinha, R. (2006). Imaging response inhibition in a stop-signal task: Neural correlates independent of signal monitoring and post-response processing. *Journal of Neuroscience, 26,* 186–192.

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

Lisberger, S. G., Fuchs, A. F., King, W. M., & Evinger, L. C. (1975). Effect of mean reaction time on saccadic responses to two-step stimuli with horizontal and vertical components. *Vision Research, 15,* 1021–1025.

Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A. Baddeley (Eds.), *Attention and performance IX* (pp. 205–222). Hillsdale, NJ: Erlbaum.

Logan, G. D. (1982). On the ability to inhibit complex movements: A stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 778–792.

Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 189–239). San Diego, CA: Academic Press.

Logan, G. D. (2004). Cumulative progress in formal theories of attention. *Annual Review of Psychology, 55,* 207–234.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91,* 295–327.

Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 276–291.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review, 108,* 393–434.

Logan, G. D., & Irwin, D. E. (2000). Don't look! Don't touch! Inhibitory

control of eye and hand movements. *Psychonomic Bulletin & Review, 7,* 107–112.

Logan, G. D., Schachar, R. J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science, 8,* 60–64.

Machens, C. K., Romo, R., & Brody, C. D. (2005, February 18). Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science, 307,* 1121–1124.

Marder, E., & Goaillard, J. M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience, 7,* 563–574.

Massaro, D. W. (1993). Information processing models: Microscopes of the mind. *Annual Review of Psychology, 44,* 383–425.

Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex, 13,* 257–269.

McDonald, S. A., Carpenter, R. H. S., & Shillcock, R. C. (2005). An anatomically constrained, stochastic model of eye movement control in reading. *Psychological Review, 112,* 814–840.

McGarry, T., Chua, R., & Franks, I. M. (2003). Stopping and restarting an unfolding action at various times. *Quarterly Journal of Experimental Psychology, 56,* 601–620.

Meredith, M. A., & Ramoa, A. S. (1998). Intrinsic circuitry of the superior colliculus: Pharmacophysiological identification of horizontally oriented inhibitory interneurons. *Journal of Neurophysiology, 79,* 1597–1602.

Mitchell, J., & Zipser, D. (2001). A model of visual-spatial memory across saccades. *Vision Research, 41,* 1575–1592.

Mitchell, J., & Zipser, D. (2003). Sequential memory-guided saccades and target selection: A neural model of the frontal eye fields. *Vision Research, 43,* 2669–2695.

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wagner, T. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology, 41,* 49–100.

Munoz, D. P., Dorris, M. C., Paré, M., & Everling, S. (2000). On your mark, get set: Brainstem circuitry underlying saccadic initiation. *Canadian Journal of Physiological Pharmacology, 78,* 934–944.

Munoz, D. P., & Istvan, P. J. (1998). Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. *Journal of Neurophysiology, 79,* 1193–1209.

Munoz, D. P., & Schall, J. D. (2003). Concurrent, distributed control of saccade initiation in the frontal eye field and superior colliculus. In W. T. Hall & A. Moschovakis (Eds.), *The superior colliculus: New approaches for studying sensorimotor integration* (pp. 55–82). New York: CRC Press.

Munoz, D. P., & Wurtz, R. (1993a). Fixation cells in monkey superior colliculus. I. Characteristics of cell discharge. *Journal of Neurophysiology, 70,* 559–575.

Munoz, D. P., & Wurtz, R. H. (1993b). Fixation cells in monkey superior colliculus. II. Reversible activation and deactivation. *Journal of Neurophysiology, 70,* 576–589.

Munoz, D. P., & Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus. II. Spread of activity during saccades. *Journal of Neurophysiology, 73,* 2334–2348.

Myung, I. J., Pitt, M. A., & Kim, K. (2005). Model evaluation, testing and selection. In K. Lambert & R. Goldstone (Eds.), *Handbook of cognition* (pp. 422–436). Thousand Oaks, CA: Sage.

Nelder, J. A., & Mead, R. (1965). A simple method for function minimization. *Computer Journal, 7,* 308–313.

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin, 89,* 133–162.

Nigg, J. T. (2001). Is ADHD a disinhibitory disorder? *Psychological Bulletin, 127,* 571–598.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266–300.

Oaksford, M., & Chater, N. (1998). *Rational models of cognition.* Oxford, England: Oxford University Press.

Olman, R. T. (1973). Simple reactions with random countermanding of the "go" signal. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 571–581). New York: Academic Press.

Oosterlaan, J., Logan, G. D., & Sergeant, J. A. (1998). Response inhibition in AD/HD, CD, comorbid AD/HD + CD, anxious, and control children: A meta-analysis of studies with the stop task. *Journal of Child Psychology & Psychiatry, 39,* 411–425.

Optican, L. M. (1995). A field theory of saccade generation: Temporal-to-spatial transform in the superior colliculus. *Vision Research, 35,* 3313–3320.

Osman, A., Kornblum, S., & Meyer, D. E. (1986). The point of no return in choice reaction time: Controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 243–258.

Osman, A., Kornblum, S., & Meyer, D. E. (1990). Does motor programming necessitate response execution? *Journal of Experimental Psychology: Human Perception and Performance, 16,* 183–198.

Özyurt, J., Colonius, H., & Arndt, P. A. (2003). Countermanding saccades: Evidence against independent processing of go and stop signals. *Perception & Psychophysics, 65,* 420–428.

Paré, M. & Guitton, D. (1998). Brain stem omnipause neurons and the control of combined eye-head gaze saccades in the alert cat. *Journal of Neurophysiology, 79,* 3060–3076.

Paré, M., & Hanes, D. P. (2003). Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience, 23,* 6480–6489.

Parthasarathy, H. B., Schall, J. D., & Graybiel, A. M. (1992). Distributed but convergent ordering of corticostriatal projections: Analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *Journal of Neuroscience, 12,* 4468–4488.

Pouget, P., Emeric, E. E., Stuphorn, V., Reis, K., & Schall, J. D. (2005). Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *Journal of Neurophysiology, 94,* 2086–2092.

Quaia, C., Lefevre, P., & Optican, L. M. (1999). Model of the control of saccades by superior colliculus and cerebellum. *Journal of Neurophysiology, 82,* 999–1018.

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

Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology, 90,* 1392–1407.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review, 9,* 438–481.

Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The E-Z reader model of eye-movement control in reading: Comparisons to other models. *Behavioral Brain Science, 26,* 446–526.

Ridderinkhof, K. R., Band, G. P. H., & Logan, G. D. (1999). A study of adaptive behavior: Effects of age and irrelevant information on the ability to inhibit one's actions. *Acta Psychologica, 101,* 315–337.

Rowan, T. (1990). *Functional stability analysis of numerical algorithms.* Unpublished doctoral thesis, University of Texas at Austin.

Schachar, R. J., & Logan, G. D. (1990). Impulsivity and inhibitory control in normal development and childhood psychopathology. *Developmental Psychology, 26,* 710–720.

Schachar, R. J., Mota, V., Logan, G. D., Tannock, R., & Klim, P. (2000). Confirmation of an inhibitory control deficit in attention-deficit/hyperactivity disorder. *Journal of Abnormal Child Psychology, 28,* 227–235.

Schachar, R. J., Tannock, R., & Logan, G. D. (1993). Inhibitory control,

impulsiveness, and attention deficit hyperactivity disorder. *Clinical Psychology Review, 13,* 721–739.

Schall, J. D. (1991). Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: Comparison with supplementary eye fields. *Journal of Neurophysiology, 66,* 559–579.

Schall, J. D. (1997). Visuomotor areas of the frontal lobe. In K. Rockland & J. H. Kaas (Eds.), *Cerebral cortex* (pp. 527–638). New York: Plenum.

Schall, J. D. (2004). On building a bridge between brain and behavior. *Annual Review of Psychology, 55,* 23–50.

Schall, J. D., Stuphorn, V., & Brown, J. W. (2002). Monitoring and control of action by the frontal lobes. *Neuron, 36,* 309–322.

Schiller, P. H., & Körner, F. (1971). Discharge characteristics of single units in superior colliculus of the alert rhesus monkey. *Journal of Neurophysiology, 34,* 920–936.

Scudder, C. A. (1988). A new local feedback model of the saccadic burst generator. *Journal of Neurophysiology, 59,* 1455–1475.

Scudder, C. A., Fuchs, A. F., & Langer, T. P. (1988). Characteristics and functional identification of saccadic inhibitory burst neurons in the alert monkey. *Journal of Neurophysiology, 59,* 1430–1454.

Scudder, C. A., Kaneko, C. S., & Fuchs, A. F. (2002). The brainstem burst generator for saccadic eye movements: A modern synthesis. *Experimental Brain Research, 142,* 439–462.

Segraves, M. A., & Goldberg, M. E. (1987). Functional properties of corticotectal neurons in the monkey's frontal eye field. *Journal of Neurophysiology, 58,* 1387–1419.

Shepard, R. N. (1994). Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review, 1,* 2–28.

Slater-Hammel, A. T. (1960). Reliability, accuracy, and refractoriness of a transit reaction. *Research Quarterly, 31,* 217–228.

Sparks, D. L. (1975). Response properties of eye movement-related neurons in the monkey superior colliculus. *Brain Research, 6,* 147–152.

Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews Neuroscience, 3,* 952–964.

Sternberg, S. (2001). Separate modifiability, mental modules, and the use of pure and composite measures to reveal them. *Acta Psychologica, 106,* 147–246.

Sternberg, S., Monselli, S., Knoll, R. L., & Wright, C. E. (1978). The latency and duration of rapid movement sequences: Comparisons of speech and typewriting. In G. E. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 117–152). New York: Academic.

Takahashi, M., Sugiuchi, Y., Izawa, Y., & Shinoda, Y. (2005). Commissural excitation and inhibition by the superior colliculus in tectoreticular neurons projecting to omnipause neuron and inhibitory burst neuron regions. *Journal of Neurophysiology, 94,* 1707–1726.

Teller, D. Y. (1984). Linking propositions. *Vision Research, 24,* 1233–1246.

Teller, D. Y., & Pugh, E. (1983). Linking propositions in color vision. In J. D. Mollon & L. T. Sharpe (Eds.), *Colour vision: Physiology and psychophysics* (pp. 11–21). London: Academic Press.

Thier, P., Dicke, P. W., Haas, R., Thielert, C. D., & Catz, N. (2002). The role of the oculomotor vermis in the control of saccadic eye movements. *Annals of the New York Academy of Science, 978,* 50–62.

Townsend, J. T., & Ashby, F. G. (1983). *The stochastic modeling of elementary psychological processes.* Cambridge, England: Cambridge University Press.

Trappenberg, T. P., Dorris, M., Klein, R. M., & Munoz, D. P. (2001). A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *Journal of Cognitive Neuroscience, 13,* 256–271.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592.

van den Wildenberg, W. P. M., & van der Molen, M. W. (2004a). Additive factors analysis of inhibitory processing in the stop-signal paradigm. *Brain and Cognition, 56,* 253–266.

van den Wildenberg, W. P. M., & van der Molen, M. W. (2004b). Developmental trends in simple and selective inhibition of compatible and incompatible responses. *Journal of Experimental Child Psychology, 87,* 201–220.

van Opstal, A. J., & van Gisbergen, J. A. M. (1989). A model for collicular efferent mechanisms underlying the generation of saccades. *Brain, Behavior and Evolution, 33,* 90–94.

Verbruggen, F., Liefooghe, B., & Vandierendonck, A. (2004). The interaction between stop-signal inhibition and distractor interference in the flanker and Stroop task. *Acta Psychologica, 116,* 21–37.

Waitzman, D., Ma, T., Optican, L., & Wurtz, R. (1991). Superior colliculus neurons mediate the dynamic characteristics of saccades. *Journal of Neurophysiology, 66,* 1716–1737.

Williams, B. R., Ponesse, J. S., Schachar, R. J., Logan, G. D., & Tannock, R. (1999). Development of inhibitory control across the life span. *Developmental Psychology, 35,* 205–213.

Wong, K.-F. & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience, 26,* 1314–1328.

Wurtz, R. H., & Goldberg, M. E. (1972). Activity of superior colliculus in behaving monkey: III. Cells discharging before eye movements. *Journal of Neurophysiology, 35,* 575–586.

Wurtz, R. H., & Goldberg, M. E. (1989). *The neurobiology of saccadic eye movements.* Amsterdam: Elsevier.

Yoshida, K., Iwamoto, Y., Chimoto, S., & Shimazu, H. (2001). Disynaptic inhibition of omnipause neurons following electrical stimulation of the superior colliculus in alert cats, *Journal of Neurophysiology, 85,* 2639–2642.

Zingale, C. M., & Kowler, E. (1987). Planning sequences of saccades. *Vision Research, 27,* 1327–1341.

Zipser, D., & Andersen, R. A. (1988, February 25). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature, 331,* 679–684.

Received March 13, 2006
Revision received November 7, 2006
Accepted November 9, 2006 ■