Published in final edited form as:

Brain Res. 2010 March 8; 1318: 178–187. doi:10.1016/j.brainres.2009.12.018.

Optimal performance in a countermanding saccade task

KongFatt Wong-Lin^{a,b,*}, Philip Eckhoff^{a,b}, Philip Holmes^{a,b,c}, and Jonathan D. Cohen^{b,d}

^a Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA

^b Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA

^c Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

^d Department of Psychology, Princeton University, Princeton, NJ 08544, USA

Abstract

Countermanding an action is a fundamental form of cognitive control. In a saccade-countermanding task, subjects are instructed that, if a stop signal appears shortly after a target, they are to maintain fixation rather than to make a saccade to the target. In recent years, recordings in the frontal eye fields and superior colliculus of behaving non-human primates have found correlates of such countermanding behavior in movement and fixation neurons. In this work, we extend a previous neural network model of countermanding to account for the high pre-target activity of fixation neurons. We propose that this activity reflects the functioning of control mechanisms responsible for optimizing performance. We demonstrate, using computer simulations and mathematical analysis, that pre-target fixation neuronal activity supports countermanding behavior that maximizes reward rate as a function of the stop signal delay, fraction of stop signal trials, intertrial interval, duration of timeout, and relative reward value. We propose experiments to test these predictions regarding optimal behavior.

Keywords

Saccade-countermanding task; Frontal eye fields; Computational modeling; Cognitive control; Reward rate; Optimal performance

1. Introduction

Inhibiting an action is one of the most basic forms of cognitive control. In studies, subjects are usually externally cued to inhibit (i.e., countermand) an action (Verbruggen and Logan, 2008c). In the saccade-countermanding task, subjects indicate their response to a target by saccading in its direction. However, if a stop signal appears shortly after the target, the subject must maintain fixation. Neurophysiologists have recently begun to investigate such countermanding paradigms in behaving non-human primates (Munoz and Wurtz, 1992; Hanes and Schall, 1996; Hanes et al., 1998; Hanes and Carpenter, 1999; Everling et al., 1999; Paré and Hanes, 2003). A series of studies that combine behavioral and electrophysiological measurements using this task have found the neuronal correlates of countermanding behavior in the frontal eye fields and the superior colliculus. Some (movement) neurons from these areas exhibit firing rates that ramp up over time following target presentation before a saccadic eye movement is made, while other (fixation) neurons are activated when there is maintenance of

*Corresponding author. kfwong@princeton.edu (K.F. Wong-Lin).

gaze, e.g., when there is successful countermanding of a saccade (Hanes et al., 1998; Paré and Hanes, 2003, but see Hafeed et al., 2009).

For many years, the independent race model of Logan and Cowan (1984) has been regarded as the most successful in accounting for human data. To account for recent neurophysiological findings in which there is a suppression of movement neurons and reactivation of fixation neurons in cancelled trials, Boucher et al. (2007) employ the Usher and McClelland (2001) model, equivalent to an interactive race model. This model consists of two units: a “Go” unit (representing the movement neurons) that stochastically races against a “Stop” unit (representing the fixation neurons). A movement will be made to a target if the Go unit ramps up its activity over time and crosses a prescribed threshold. However, if the Stop unit can be activated by the stop signal in time to inhibit the activity of the Go unit, then fixation will persist. In this model, the Stop unit is assumed to have zero pre-target activity. However, studies have shown (Hanes et al., 1998; Munoz and Wurtz, 1993a, 1993b; Paré and Hanes, 2003) that fixation neuronal activities have characteristically high firing rates during the pre-target period, which may be relevant to adaptive control and the optimization of performance in this task. Models that account for high pre-target activity, with differing levels of complexity, have only recently been developed (Wilimzig et al., 2006; Cutsuridis et al., 2007; Heinzle et al., 2007; Lo and Wang, 2007; Wong et al., 2007; Lo et al., 2009).

There is experimental evidence that suggests that task demands can modulate pre-target activity of fixation neurons. For example, when comparing neuronal firing rates in the frontal eye fields or the superior colliculus between anti-saccadic and pro-saccadic tasks, pre-target activities of fixation and movement neurons are respectively higher and lower in the control-requiring anti-saccadic task than in the pro-saccadic task (Everling et al., 1999; Everling and Munoz, 2000; Munoz and Everling, 2004). Stuphorn and Schall (2006) have also shown that neurons in the supplementary eye field may play an important role in providing a source of executive control to countermanding saccades. Other experiments have shown that pre-frontal cortical neurons, a possible source of top-down control, have pre-target activities that differ between anti-saccadic and pro-saccadic tasks and are correlated with correct and error trials (Everling and DeSouza, 2005; Johnston et al., 2009).

In this theoretical study, we make use of the simplest model that can account for both neuronal and behavioral data: an extension of Boucher et al. (2007) that addresses pre-target fixation neuronal activity. Inspired by the above experimental results, we use computer simulations and mathematical analysis to explore this pre-target activity’s influence on reward rate optimization and describe how optimality can be distinctly revealed. Finally, we propose future experiments to test for optimal behavior.

2. Results

2.1. A firing rate model for countermanding saccades

To construct the simplest possible model that can capture both neuronal firing rates and psychophysical data, we adopt the leaky competing accumulator network of Usher and McClelland (2001):

$$dr_{FN} = \left(-r_{FN} + [-\beta_{FN} r_{MN} + I_{pre-target} + I_{stop}]_+ \right) \frac{dt}{\tau_E} + \sigma \sqrt{\frac{dt}{\tau_E}} \eta, \quad (1)$$

$$dr_{MN} = \left(-r_{MN} + [-\beta_{MN}r_{FN} + I_{\text{target}}]_+ \right) \frac{dt}{\tau_E} + \sigma \sqrt{\frac{dt}{\tau_E}} \eta; \quad (2)$$

See Fig. 1A. Here the threshold-linear function $[x]_+ = 0$ for $x < 0$ and $[x]_+ = x$ otherwise, r_{MN} and r_{FN} represent population firing rates of movement and fixation neurons respectively, β_{MN} and β_{FN} are the effective inhibitory synaptic strengths from fixation to movement neurons and from movement to fixation neurons, σ is the magnitude of additive noise and η is a random Gaussian variable with zero mean and unit variance. $I_{\text{pre-target}}$ and I_{target} are the input currents during pre-target fixation and target presentation, and I_{stop} is the input current due to the stop signal (Fig. 1B). The latency from target onset to stop signal onset (Fig. 1B) is defined as the stop-signal delay (SSD) and is controlled by the experimenter.

Augmenting the model of Boucher et al. (2007), we include an additional pre-target input $I_{\text{pre-target}}$ to account for high firing rates of fixation neurons. Such excitatory control has also been implemented in the more complex spiking network model of Lo et al. (2009). $I_{\text{pre-target}}$ is chosen such that the pre-target activity of r_{FN} is 80 Hz, matching that in Hanes et al. (1998). Furthermore, we necessarily include a leak term to allow fixation neurons to stabilize at a high firing rate prior to target onset. (This can be seen in Eq. (1) by setting the leak $-r_{FN}$ and $\beta_{FN}r_{MN}$ to be zero, in which case r_{FN} will continue to grow.) This model, which replaces inhibitory dynamics of the interneurons by direct mutual inhibition, can be justified if we assume that the excitatory timescale τ (e.g., mediated by NMDA receptors with decay time constant of ~ 50 – 150 ms) is much slower than its inhibitory counterpart (e.g., fast GABA_A receptors with decay time constant of ~ 5 – 10 ms; cf. Wong and Wang, 2006). Following the neuronal latencies in Hanes et al. (1998), movement neurons are assumed to respond to target onset after a delay of 100 ms and fixation neurons respond 80 ms after the stop signal appears. A saccade to the target is made if $r_{MN}(t)$ crosses a prescribed decision or movement threshold of 90 Hz. The reaction time (which includes the above latencies) is the time from target onset to threshold crossing time plus an additional duration of 10 ms (cf. Boucher et al., 2007) to represent the automatic and ballistic cascade of activity required for eye movement. Further details of the model parameters are given in Table 1.

We follow Hanes et al. (1998) in labeling the types of possible trials as follows. Trials with and without a stop signal are labeled as stop-signal trials and no-stop-signal trials, respectively. Stop-signal trials can be further categorized into cancelled and noncancelled trials (also known as “signal inhibit” and “signal respond,” respectively, in the literature), in which an impending saccade is cancelled (no crossing of threshold) and cancellation fails (crossing of threshold), respectively. In cancelled trials, inhibition from the fixation neurons is typically reactivated by the stop signal in time to suppress the movement neurons before they can reach threshold and fixation of gaze is maintained. In principle, in some no-stop-signal trials movement neurons may fail to reach threshold when inhibition due to high pre-target fixation activity is too strong. However, we do not observe such behavior over the range of parameters investigated here.

Figs. 1C–F and 2 show that this model can reasonably capture both the behavioral and neuronal data of Hanes et al. (1998). We did not attempt to reproduce the phasic burst of fixation neuronal activity (Fig. 2B) in stop-signal trials; in our model the activities approach a steady state (Fig. 2D). Trial-averaged activity of the movement neural unit in stop-signal trials also tend to deviate from no-stop-signal trials (Fig. 2C) early in the trial, unlike the data of Fig. 2A. This

is because the activity of the movement neural unit must ramp up sufficiently slowly in order to be suppressed in time by the fixation unit in stop-signal trials (Fig. 4).

2.2. Phase plane analysis

To further illuminate the model, we represent its dynamics in the phase space spanned by the firing rates of the two neural units. Fig. 3A shows that averaged orbits begin at stimulus onset on the lower right and traverse toward the upper left, crossing the movement threshold in no-stop-signal trials. In successfully-cancelled stop-signal trials, they turn back and travel to the right of the starting point. The fact that orbits traverse different regions of this phase space under the fixation, target and stop conditions implies that reduction to a one-dimensional drift-diffusion model, as in Bogacz et al. (2006), is impossible.

The dynamics are explained by the three vector fields of the noise-free system plotted in Figs. 3B–D, respectively showing pre-target fixation (B), target presentation (C), and target plus stop signal presentation (D). These show nullclines on which the derivatives dr_{FN}/dt and dr_{MN}/dt change sign, stable fixed points at their intersections, and vectors indicating the local direction and speed of orbits. In the fixation period orbits approach the stable state $(r_{FN}, r_{MN}) = (I_{pre-target}, 0)$ (B). When $I_{pre-target}$ turns off and I_{target} turns on, the vector field changes and the stable fixed point jumps to $(r_{FN}, r_{MN}) = (0, I_{target})$, attracting the orbit and inducing threshold crossing (C). When I_{stop} is applied, the orbit turns and descends to approach a third fixed point $(r_{FN}, r_{MN}) = (I_{stop}, 0)$, remaining below threshold if the SSD is short enough (D); for longer SSDs it may still continue to rise, thus crossing threshold in a noncancelled trial. The more complex dynamics of Lo et al. (2009, Fig. 8), with coexisting attractors, is not necessary to capture these behaviors.

2.3. Reward rate: A measure of behavioral performance

The overall performance of the model is measured by the average reward rate received over a block of trials, calculated by dividing the total number of correct responses over the total time taken for the block. Correct responses are counted by summing the total number of successfully cancelled trials N_{canc} plus the number of completed no-stop trials $N_{no-stop}$. Denoting the total time spent in a block of trials as T_{total} , the average reward rate R is therefore

$$R = \frac{\gamma N_{canc} + N_{no-stop}}{T_{total}} = \frac{\gamma N_{canc} + N_{no-stop}}{N_{canc} T_{canc} + N_{nonc} (mRT_{nonc} + T_{nonc}) + N_{no-stop} (mRT_{no-stop} + T_{no-stop})} \quad (3)$$

where N_{nonc} is the total number of non-cancelled trials, mRT_{nonc} and $mRT_{no-stop}$ are the mean reaction times for noncancelled and no-stop signal trials, and γ is the reward value of cancelled trials relative to that of responding correctly in no-stop-signal trials. Rewards are given only when there is a correct response in no-stop-signal trials or when there is successful cancellation of movement in stop-signal trials, and $\gamma = 1$ unless otherwise specified. The experimentally controlled nondesired time delays for cancelled (T_{canc}), no-stop ($T_{no-stop}$), and noncancelled trials (T_{nonc}) are scheduled to match those of Hanes et al. (1998). Specifically, a successfully cancelled trial (or a no-stop trial) consists of withholding gaze at fixation point (or choice target) for a duration of $T_{fix} = 700$ ms (or $T_{target} = 700$ ms), while there is a timeout of $T_{timeout} = 500$ ms following a noncancelled trial.

All trial types have inserted an additional constant intertrial interval (ITI) of duration $T_{ITI} = 500$ ms, unless otherwise specified. This particular value is selected so that the change in reaction time can have a significant effect on reward rates. We will later study the effects of longer T_{ITI} and $T_{timeout}$. Thus, $T_{canc} = T_{fix} + T_{ITI}$, $T_{no-stop} = T_{target} + T_{ITI}$ and $T_{nonc} = T_{timeout} + T_{ITI}$.

As there is no explicit mathematical solution to the model's reward rate performance, we will first use computer simulations to search for parameters that produce optimal reward rates. Then we provide further understanding of the model behavior through mathematical analysis, with some approximations guided by results from the simulations.

2.4. Reward rate: Simulations

To understand the basic effects of varying the pre-target firing rate of fixation neurons (FN), we first examine only stop-signal trials. Fixing SSD at 169 ms, we find that increasing pre-target activity delays the rise in activity of the movement neural unit after target onset (Figs. 4A and B). This generally slows the ramping toward movement threshold, resulting in longer reaction times in noncancelled trials (Fig. 4C) but producing more successful cancellations (Fig. 4D). Accordingly, there may exist an optimal pre-target FN activation at which speed is appropriately balanced with accuracy of responses (and nonresponses) to optimize reward rate (R).

Next, we search for optimal reward rates under various SSDs, fixing the fraction of stop signal trials in a block at 0.25 (cf. Hanes et al., 1998). For very short SSDs (< 50 ms), reward rates decrease monotonically with pre-target FN activity, while for long SSDs (> 150 ms) they increase monotonically (Fig. 5A). For intermediate SSDs, reward rate is maximized at a nonzero FN activity, which grows for longer SSDs. Thus, varying SSD alone can reveal optimal countermanding behavior as a function of FN activity. We then increase the fraction of stop signal trials with fixed SSD=100 ms. This shifts the optimal FN activity rightward (Fig. 5B), significantly increasing the range of reward rates produced by different SSDs (Fig. 5C). More importantly, this wider range allows reward rate slopes to change sign, producing maxima within the physiological range of FN activities.

Optimal countermanding behavior can also be brought into sharper contrast by increasing intertrial interval (ITI) or timeout duration (compare Fig. 6A with Figs. 6B and C), although the former yields much lower reward rates with less clear maxima. Finally, if the reward for canceling a response in a stop-signal trial is four times that for responding to a no-stop trial, the range of reward rates is again wider (Fig. 6D). Thus, a judicious choice of experimental conditions can help to distinguish optimum parameter values.

2.5. Reward rate: Analysis

To further illuminate optimal countermanding behavior, we examine the mathematical relationships between the rate of change of reward rate and the pre-target FN activity and study how this relationship is affected by the following experimental parameters: SSD, fraction of stop-signal trials $(N_{\text{canc}} + N_{\text{nonc}})/(N_{\text{canc}} + N_{\text{nonc}} + N_{\text{no-stop}})$, intertrial interval T_{ITI} , timeout duration T_{timeout} , and relative reward value of a cancelled trial γ .

Optimal reward rates can be found by finding the condition $dR/dp=0$, where p denotes pre-target FN activity. We can then write

$$\frac{dR}{dp} = \frac{dN_{\text{canc}}/dp}{T_{\text{total}}} - \frac{(\gamma N_{\text{canc}} + N_{\text{no-stop}})}{T_{\text{total}}^2} \frac{dT_{\text{total}}}{dp} \quad (4)$$

where

$$\frac{dT_{\text{total}}}{dp} = \frac{dN_{\text{canc}}}{dp} T_{\text{canc}} + \frac{dN_{\text{nonc}}}{dp} (mRT_{\text{nonc}} + T_{\text{nonc}}) + N_{\text{nonc}} \frac{d(mRT_{\text{nonc}})}{dp} + N_{\text{no stop}} \frac{d(mRT_{\text{no-stop}})}{dp} \quad (5)$$

Obvious inequalities in Eq. (5) are

$$dN_{\text{canc}}/dp > 0 \quad (6)$$

$$d(mRT_{\text{no-stop}})/dp > 0 \quad (7)$$

$$dN_{\text{nonc}}/dp = d(N_{\text{stop}} - N_{\text{canc}})/dp = -dN_{\text{canc}}/dp < 0 \quad (8)$$

The first two inequalities above are due to the fact that higher pre-target FN activity can result in more cancelled trials and slower reaction times in no-stop signal trials. The third inequality is due to the definition that $N_{\text{stop}} = N_{\text{canc}} + N_{\text{nonc}}$. Figs. 2A and C show that higher inhibition from fixation neurons results in slower ramping of movement neurons and thus longer noncancelled reaction times $d(mRT_{\text{nonc}})/dp > 0$. In fact, Fig. 2C shows that $d(mRT_{\text{nonc}})/dp \approx d(mRT_{\text{no stop}})/dp$, and we can then approximate Eq. (5) as

$$\frac{dT_{\text{total}}}{dp} \approx \frac{dN_{\text{canc}}}{dp} (T_{\text{canc}} - mRT_{\text{nonc}} - T_{\text{nonc}}) + (N_{\text{nonc}} + N_{\text{no-stop}}) \frac{d(mRT_{\text{no-stop}})}{dp} \quad (9)$$

Substituting Eq. (9) into Eq. (4), we have

$$T_{\text{total}}^2 \frac{dR}{dp} \approx \frac{dN_{\text{canc}}}{dp} T_{\text{total}} - \frac{dN_{\text{canc}}}{dp} (\gamma N_{\text{canc}} + N_{\text{no-stop}}) (T_{\text{canc}} - mRT_{\text{nonc}} - T_{\text{nonc}}) - (\gamma N_{\text{canc}} + N_{\text{no-stop}}) (N_{\text{nonc}} + N_{\text{no-stop}}) \frac{d(mRT_{\text{no-stop}})}{dp} \quad (10)$$

According to inequalities (7) and (8), the first term on the right of Eq. (10) is always positive while the third term is always negative. The second term can be reduced to $-(dN_{\text{canc}}/dp) (\gamma N_{\text{canc}} + N_{\text{no-stop}}) (T_{\text{fix}} - mRT_{\text{nonc}} - T_{\text{timeout}})$, such that its sign is strictly dependent on the sign of the factor $(T_{\text{fix}} - mRT_{\text{nonc}} - T_{\text{timeout}})$. If we use the values of $T_{\text{fix}} = 700$ ms and $T_{\text{timeout}} = 500$ ms as in Hanes et al. (1998), then this term can either be positive or negative, depending on the noncancelled mean reaction time mRT_{nonc} . We shall now examine dR/dp under different experimental conditions.

Case 1 (mostly no-stop signal trials)—When there are more no-stop-signal than stop-signal trials, $N_{\text{stop}} \ll N_{\text{no-stop}} \approx N_{\text{total}}$, and Eq. (10) becomes

$$T_{\text{total}}^2 \frac{dR}{dp} \approx N_{\text{no-stop}} \frac{dN_{\text{canc}}}{dp} \times (\text{mRT}_{\text{no-stop}} + T_{\text{no-stop}} - T_{\text{fix}} + \text{mRT}_{\text{nonc}} + T_{\text{timeout}}) - N_{\text{no-stop}}^2 \frac{d(\text{mRT}_{\text{no-stop}})}{dp} \quad (11)$$

For very short SSD, most of the stop-signal trials are cancelled trials. This implies $N_{\text{nonc}} \ll N_{\text{canc}} \approx N_{\text{stop}}$. More importantly, Fig. 2D shows that $dP_{\text{canc}}/dp = (1/N_{\text{total}})(dN_{\text{canc}}/dp) > 0$ can be very small due to a ceiling effect. Then the first term in Eq. (11) is small and dR/dp can be negative, as found in our simulation studies. This will not occur with longer SSDs.

Case 2 (mostly stop-signal trials)—If most of the trials are stop-signal trials, then $N_{\text{no-stop}} \ll N_{\text{stop}} \approx N_{\text{total}}$, and Eq. (10) becomes

$$T_{\text{total}}^2 \frac{dR}{dp} \approx \frac{dN_{\text{canc}}}{dp} (N_{\text{nonc}} + N_{\text{canc}}) (\text{mRT}_{\text{nonc}} + T_{\text{nonc}}) - N_{\text{canc}} N_{\text{nonc}} \frac{d(\text{mRT}_{\text{no-stop}})}{dp} \quad (12)$$

with relative reward rate $\gamma = 1$. As mentioned in Case 1, for very short SSD, $N_{\text{nonc}} \ll N_{\text{canc}} \approx N_{\text{stop}}$, and both terms on the right of Eq. (12) can be small since both dN_{canc}/dp and N_{nonc} are small. On the other hand, for very long SSD, most stop signal trials are actually non-cancelled trials, i.e., $N_{\text{canc}} \ll N_{\text{nonc}} \approx N_{\text{stop}}$. Furthermore, dN_{canc}/dp is no longer small. So, we have

$$T_{\text{total}}^2 \frac{dR}{dp} \approx \frac{dN_{\text{canc}}}{dp} N_{\text{nonc}} (\text{mRT}_{\text{nonc}} + T_{\text{nonc}}) > 0 \quad (13)$$

and this result of positive slope is confirmed by our simulations.

Case 3 (longer T_{ITI} , T_{timeout} or larger γ)—Up to this point, we have fixed the ITI and timeout to be 500 ms, and relative reward rate γ to be 1. However, these parameters are set by experimentalists (e.g., duration of timeout is specifically set to discourage impulsive inaccurate responses from subjects). In Eq. (10), T_{ITI} and T_{timeout} are implicit in the first and second terms on the right through T_{total} and T_{nonc} , respectively, which contribute positively to dR/dp . Thus, increasing ITI or timeout can allow optimal behavior to exist even for small SSDs or very few stop signal trials. Finally, if the reward value for a cancelled trial is higher than that of a no-stop signal trial ($\gamma > 1$), then it generally contributes to a more negative reward rate slope (mainly through the third right term in Eq. (10)). The range of reward rates can then be larger and optimal countermanding will become more distinctive. These results are supported by the simulations described above.

3. Discussion

3.1. Summary of results

We have extended the modeling work of Boucher et al. (2007) to include the pre-target activity of fixation neurons. Constrained by the data of Hanes et al. (1998), using computer simulations and mathematical analysis, we searched for conditions under which optimal countermanding behavior can occur as a function of pre-target FN activity. Our results show that this is most clearly observed when a small fraction of the trials are stop-signal trials, there are low SSDs, and higher rewards are offered for cancelled than for no-stop trials. In principle, increasing ITI

has a similar effect, but it can cause overall reward rates and the range of achievable reward rates to decrease, so that optima become difficult to distinguish. Although our parameter values were based mostly on Hanes et al. (1998), we expect our general results to hold for other parameter regimes (e.g., with human data), as shown in our analysis.

A key assumption in this study is that the pre-target FN activities can vary depending on task demands. Due to recurrent inhibition, the dynamics of movement neurons can therefore be indirectly controlled. This means that pre-target activity of fixation neurons may provide a measure of the level of cognitive control. During the epoch between target and stop signal onset, we allow the activity of fixation neurons in our model to decay at a constant rate toward baseline, irrespective of their initial firing rate level. This implies that the higher this activity is, the more the ramping up activity of movement neurons will be temporally shifted and delayed, as shown in Figs. 4A–C. In fact, the size of this temporal shift is $\sim \Delta \exp(-t/\tau_E)$, where Δ is the change in the pre-target FN activity.

Although cognitive control in our model is implemented somewhat similarly to the independent modeling study of Lo et al. (2009), our simple model does not require temporal jitter in the offset of the fixation neuronal activities and additional neuronal/synaptic nonlinearities to account for both the neuronal and behavioral data. Furthermore, with considerably fewer parameters, our model is perhaps also more amendable to Bayesian approaches in which the priors can be directly linked to the pre-target activity of fixation neurons.

3.2. Experiments for exploring optimal countermanding

Our optimal reward rate curves are not as striking as those for the two-alternative forced-choice task paradigm studied in Bogacz et al. (2006). Their optimal reward rate distinctly drops as the decision threshold increases whereas in our study, it does not decrease steeply as pre-target FN activity increases. The main reason for this in the countermanding paradigm is that accuracy changes more than reaction time, since the latter does not vary significantly. Our results suggest that increasing the timeout duration or the reward value of cancelled trials shape the conditions under which optimal countermanding behavior occurs. Perhaps making the countermanding task more difficult (Logan, 1981, 1983, 1985; Logan and Burkell, 1986; Verbruggen and Logan, 2009) would help increase overall reaction times and allow more distinctive optimal behavior to be observed. A more interesting experiment may be to implement countermanding of two-alternative forced-choice tasks in non-human primates.

3.3. Future work

Alternative controlling mechanisms were not considered in this work. One such mechanism may be active suppression of MN activation (via FN) after target onset. Another is adaptation of movement or decision thresholds (Bogacz et al., 2006). Although these neural mechanisms differ from our proposals, we expect their behavioral effects to be similar. In future work, it may be worthwhile to compare our simple model to more complex models (Frank, 2006; Heinzle et al., 2007; Lo et al., 2009).

Important across-trial effects in countermanding tasks (Emeric et al., 2007; Verbruggen et al., 2008; Verbruggen and Logan, 2008b,a) have been ignored in this work. For example, learning may be involved to optimize countermanding behavior (Liddle et al., 2009). In Wong et al. (2007), a simple leaky linear integrator added to our model suffices to reproduce the covariation of response time with the fraction of stop trials (Schall et al., 2002, Fig. 3). It will thus be interesting to implement the full performance-monitoring-adjustment mechanisms similar to Schall and Boucher (2007).

4. Experimental procedures

The Euler–Marayama method (Higham, 2001) with an integration time step of 0.1 ms was used in computer simulations when integrating the stochastic differential Eqs. (1–2); smaller time steps did not change the results. When plotting the activity time course, 5,400 independently simulated realizations (trials) were averaged to produce the data. When plotting the reward rate figures, a data point for a set of model parameters was obtained by first simulating separately no-stop trials and stop-signal trials and then combining these to calculate the reward rate using Eq. (3). Simulation code was written in Matlab and run on a Linux workstation. Phase plane analysis was performed using the software XPPAUT.

Acknowledgments

We thank Gordon D. Logan and Jeffrey D. Schall for careful readings of an earlier version of this manuscript and suggestions of useful references in the literature. We would also like to thank two anonymous reviewers for helping improve the manuscript. P.E. was supported by the Fannie and John Hertz Foundation and National Science Foundation, and K.F.W.-L., P.H. and J.D.C. by Public Health Service Grants MH 62196 and Air Force Office of Scientific Research Grant FA9550-07-1-0537. The United States government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the United States government.

References

- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol Rev* 2006;113:700–765. [PubMed: 17014301]
- Boucher L, Palmeri TJ, Gordan GD, Schall JD. Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psych Rev* 2007;114:376–397.
- Cutsuridis V, Smyrnis N, Evdokimidis I, Perantonis S. A neural model of decision-making by the superior colliculus in an antisaccade task. *Neural Netw* 2007;20:690–704. [PubMed: 17446043]
- Emeric EE, Brown JW, Boucher L, Carpenter RH, Hanes DP, Harris R, Logan GD, Mashru RN, Par' e M, Pouget P, Stuphorn V, Taylor TL, Schall JD. Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vis Res* 2007;47:35–49. [PubMed: 17081584]
- Everling S, DeSouza JF. Rule-dependent activity for prosaccades and antisaccades in the primate prefrontal cortex. *J Cogn Neurosci* 2005;17:1483–1496. [PubMed: 16197701]
- Everling S, Klein MCDRM, Munoz DP. Role of primate superior colliculus in preparation and execution of anti-saccades and prosaccades. *J Neurosci* 1999;19:2740–2754. [PubMed: 10087086]
- Everling S, Munoz DP. Neuronal correlates for preparatory set associated with pro-saccades and anti-saccades in the primate frontal eye field. *J Neurosci* 2000;20:387–400. [PubMed: 10627615]
- Frank MJ. Hold your horses: a dynamics computational role for the subthalamic nucleus in decision making. *Neural Netw* 2006;19:1120–1136. [PubMed: 16945502]
- Hafed ZM, Goffart L, Krauzlis RJ. A neural mechanisms for microsaccade generation in the primate superior colliculus. *Science* 2009;323:940–943. [PubMed: 19213919]
- Hanes DP, Carpenter RH. Countermanding saccades in humans. *Vis Res* 1999;39:2791.
- Hanes DP, Patterson WF, Schall JD. Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J Neurophysiol* 1998;79:817–834. [PubMed: 9463444]
- Hanes DP, Schall JD. Neural control of voluntary movement initiation. *Science* 1996;274:427–430. [PubMed: 8832893]
- Heinzle J, Hepp K, Martin KA. A microcircuit model of the frontal eye fields. *J Neurosci* 2007;27:9341–9353. [PubMed: 17728448]
- Hestrin S, Sah P, Nicoll RA. Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices. *Neuron* 1990;5:247–253. [PubMed: 1976014]
- Higham DJ. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Rev* 2001;43:535–546.

- Johnston K, DeSouza JF, Everling S. Monkey prefrontal cortical pyramidal and putative interneurons exhibit differential patterns of activity between prosaccade and antisaccade tasks. *J Neurosci* 2009;29:5516–5524. [PubMed: 19403819]
- Liddle EB, Scerif G, Hollis CP, Batty MJ, Groom MJ, Liotti M, Liddle PF. Looking before you leap: a theory of motivated control of action. *Cognition* 2009;112:141–158. [PubMed: 19409540]
- Lo CC, Boucher L, Paré M, Schall JD, Wang XJ. Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. *J Neurosci* 2009;29:9059–9071. [PubMed: 19605643]
- Lo C-C, Wang X-J. A spiking neural network model of control of action in the countermanding task. *Soc Neurosci Abstract No.* 636.10. Nov;2007
- Logan, GD. Attention, automaticity, and the ability to stop a speeded choice response. In: Long, J.; Baddeley, AD., editors. *Attention and Performance*. New Jersey: Erlbaum; 1981. IX
- Logan GD. On the ability to inhibit simple thoughts and actions: I. Stop signal studies of decision and memory. *J Exp Psychol Learn Mem Cogn* 1983;9:585–606.
- Logan GD. On the ability to inhibit simple thoughts and actions: II. Stop-signal studies of repetition priming. *J Exp Psychol Learn Mem Cogn* 1985;11:675–691.
- Logan GD, Burkell J. Dependence and independence in responding to double stimulation: a comparison of stop, change, and dual-task paradigms. *J Exp Psychol Hum Percept Perform* 1986;12:549–563.
- Logan GD, Cowan WB. On the ability to inhibit thought and action: a theory of an act of control. *Psych Rev* 1984;91:295–327.
- Munoz DP, Everling S. Look away: the anti-saccade task and the voluntary control of eye movement. *Nat Rev Neurosci* 2004;5:218–228. [PubMed: 14976521]
- Munoz DP, Wurtz RH. Role of the rostral superior colliculus in active visual fixation and execution of express saccades. *J Neurophysiol* 1992;67:1000–1002. [PubMed: 1588382]
- Munoz DP, Wurtz RH. Fixation cells in superior colliculus. I Characteristics of cell discharge. *J Neurophysiol* 1993a;70:559–575. [PubMed: 8410157]
- Munoz DP, Wurtz RH. Fixation cells in superior colliculus. II Reversible activation and deactivation. *J Neurophysiol* 1993b;70:576–589. [PubMed: 8410158]
- Paré M, Hanes DP. Controlled movement processing: superior colliculus activity associated with countermanded saccades. *J Neurosci* 2003;23:6480–6489.
- Schall JD, Boucher L. Executive control of gaze by the frontal lobes. *Cogn Affect Behav Neurosci* 2007;7:396–412. [PubMed: 18189013]
- Schall JD, Stuphorn V, Brown JW. Monitoring and control of action by the frontal lobes. *Neuron* 2002;36:309–322. [PubMed: 12383784]
- Spruston N, Jonas P, Sakmann B. Dendritic glutamate receptor channel in rat hippocampal CA3 and CA1 pyramidal neurons. *J Physiol* 1995;482:325–352. [PubMed: 7536248]
- Stuphorn V, Schall JD. Executive control of countermanding saccades by the supplementary eye field. *Nat Neurosci* 2006;9:925–931. [PubMed: 16732274]
- Usher M, McClelland JL. On the time course of perceptual choice: the leaky competing accumulator model. *Psychol Rev* 2001;108:550–592. [PubMed: 11488378]
- Verbruggen F, Logan GD. Automatic and controlled response inhibition: associative learning in the go/no-go and stop-signal paradigms. *J Exp Psychol Gen* 2008a;137:649–672. [PubMed: 18999358]
- Verbruggen F, Logan GD. Long-term aftereffects of response inhibition: memory retrieval task goals and cognitive control. *J Exp Psychol Hum Percept Perform* 2008b;34:1229–1235. [PubMed: 18823207]
- Verbruggen F, Logan GD. Response inhibition in the stop-signal paradigm. *Trends Cogn Sci* 2008c;12:418–424. [PubMed: 18799345]
- Verbruggen F, Logan GD. Proactive adjustments of response strategies in the stop-signal paradigm. *J Exp Psychol Hum Percept Perform* 2009;35:835–854. [PubMed: 19485695]
- Verbruggen F, Logan GD, Liefoghe B, Vandierendonck A. Short-term aftereffects of response inhibition: repetition priming or between-trial control adjustments? *J Exp Psychol Hum Percept Perform* 2008;34:413–426.

- Wilmzig C, Schneider S, Schöner G. The time course of saccadic decision making: dynamics field theory. *Neural Netw* 2006;19:1059–1074. [PubMed: 16942860]
- Wong K-F, Eckhoff P, Holmes P, Cohen JD. Analysis of a biologically realistic model for saccade-countermanding tasks. *Soc Neurosci Abstract* No. 719. Nov;2007
- Wong KF, Wang XJ. A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 2006;26:1314–1328. [PubMed: 16436619]

Fig. 1.

The model reproduces behavioral data. (A) Model architecture with asymmetrical mutual inhibition between movement (MN) and fixation (FN) neural units. (B) Timecourse of inputs to MN (green) and FN (red); SSD=stop signal delay. Cognitive control is assumed to vary pre-target FN input. Experimental data (C, D) and model simulations (E, F); SSDs in legends; axes for (C, D) are the same as (E, F), respectively. (C, E) Probability of noncancelled trials. (D, F) Cumulative distributions of reaction times for noncancelled and no-stop signal trials. (C, D) Adapted from Boucher et al. (2007), with permission; original data is from Hanes et al. (1998).

Fig. 2.

The model's neural activity timecourse is similar to experimental data. Time course of trial-averaged activities for movement (A, C) and fixation (B, D) neural neurons for no-stop (thin) and cancelled (bold) trials; (A, B) experimental data; (C, D) model simulations. Solid vertical lines in (A, B) denote onset of stop signal; arrows in (A, C) denote onsets of significant deviation of activities between no-stop and stop signal trials. (C, D) show SSDs of 69 ms (light grey), 117 ms (dark grey), and 169 ms (thick black). All data averaged over 200 trials. (A, B) Adapted from Boucher et al. (2007), with permission; original data is from Hanes et al. (1998).

Fig. 3.

Model dynamics in phase space. (A) Orbits averaged over many trials; no-stop signal trials shown dashed, cancelled trials for two SSDs solid. (B–D) Vector fields of the noise-free system (arrows) and nullclines for fixation and movement neurons (red and green respectively); nullclines intersect at stable fixed points (black filled circles). (B) Pre-target epoch: fixation neurons stabilize at high firing rate while movement neurons have low activity. (C) Target epoch: movement neuron activities ramp up toward a new high steady state above movement threshold (dashed blue). (D) In stop-signal trials orbits turn to approach a third steady state with high fixation neuron activity.

Fig. 4.

Effects of pre-target FN activity and stop-signal delay. High FN activities (red) delay ramping up of MN activities (green) in both noncancelled (A) and cancelled (B) trials; darker colors denote higher FN activity; SSD=169 ms. (C, D) Higher pre-target FN activity increases noncancelled mean reaction time (RT) (C) and probability of cancellation $P(\text{cancelled})$ (D); SSDs for (C, D) shown in legend on (C).

Fig. 5.

Optimal countermanding with various SSDs and fractions of stop-signal trials. (A, B) Increase in only SSD (A) or fraction of stop-signal trials (B) shifts optimal FN activity to higher values (filled circles), while overall reward rate decreases. (C) Range of reward rate over a range of SSD values increases with increasing fraction of stop signal trials, allowing reward rate slope to change from negative to positive, and hence reward rate maxima to occur. Bold, dashed, dashed-dot, dotted lines denote SSDs 0, 50, 100, and 150 ms, respectively. (A–C) ITI and timeout are 500 ms. Note that vertical scales differ.

Fig. 6. Increasing ITI, timeout or reward value of cancelled trials can reveal reward rate maxima. Fraction of stop signal trials in (A–D) are fixed at 10% such that a T_{ISI} and a $T_{timeout}$ of 500 ms does not show reward rate maxima. Increasing ITI (B), timeout (C), or reward value of a cancelled trial relative to a no-stop trial, here by FOUR times (D), can produce reward rate maxima (filled circles). Note vertical scales differ.

Table 1

Model parameters.

Parameter	Value	Notes
τ_E	50 ms	Synaptic decay time constant
β_{MN}	3.45	Fit behavioral data of H98
β_{FN}	0.4	Fit behavioral data of H98
$I_{\text{pre-target}}$	80 Hz	Fit neuronal data of H98
I_{target}	124 Hz	Fit behavioral, neuronal data of H98
I_{stop}	108 Hz	Fit behavioral, neuronal data of H98
σ	6.08 Hz	Fit behavioral data of H98
$T_{MN,\text{latency}}$	100 ms	Deduced from neuronal data of H98
$T_{FN,\text{latency}}$	80 ms	Deduced from neuronal data of H98
Motor/decision threshold	90 Hz	Fit behavioral data of H98

Parameters were selected either to fit behavioral data or to be consistent with neuronal data of Hanes et al. (1998), or both. Value of τ_E obtained from NMDA receptors (Spruston et al., 1995; Hestrin et al., 1990).