NOTE: This unpublished manuscript had been submitted to Nature on 09/17/2001 (#C09401) and rejected on 10/17/2001.

Measuring the time to cancel a saccade

Hans Colonius & Adele Diederich

Institut für Kognitionsforschung, Universität Oldenburg, P.O. Box 2503, D-26111 Oldenburg, Germany.

Adaptive behaviour requires not only the ability to initiate and maintain action, but also to withhold a planned movement. In the eye movement countermanding task subjects are instructed to withhold a saccade in response to an imperative stop signal. The decision to initiate or to withhold an eye movement arises out of a balance of activity in neurons representing gaze-holding and gaze-shifting mechanisms in the frontal eye fields and superior colliculus[1]. Analysis and interpretation of neurophysiological and behavioural data have been related to a race model that claims statistical independence between these mechanisms[2-4]. Given recent behavioural evidence against this assumption[5-6], here we suggest an alternative, diffusion-type model. It leads to estimates of the time required to cancel a saccade notably different from those based on the race model, and it accounts for behavioural data including possible subject strategies in performing the task. At present, it seems premature to identify either model with the structure and functions of a single specific brain structure.

In the eye movement countermanding task the subject is instructed to make a saccade as quickly as possible from the fixation point to a visual target suddenly appearing somewhere in the periphery. In about 25 % of the trials a stop signal (often, the reappearance of the fixation point) is presented after a variable time delay that requires the subject to inhibit the saccade. The probability of a successful inhibition is a decreasing function of the time delay between the onset of the target and the onset of

the stop signal (*stop signal delay*). This observation is consistent with the notion that the generation of a saccade is inherently stochastic with saccadic latency varying across trials[7]. A quick saccade has a high chance of being executed even if a stop signal was presented, especially if the stop signal is presented relatively late, because the saccade would be initiated before the stop signal could influence the system. On the other hand, if the stop signal is presented very shortly after the onset of the go signal and/or the saccade latency is rather long, it is very likely that the stop signal has enough time to become effective and to cause an inhibition of the movement. Observable data in this paradigm are the frequencies of successful inhibition as a function of stop signal delay and the distribution of saccadic response times both in the control condition (no stop signal) and in the presence of stop signals (non-cancelled saccades).

A common interpretation of these data is in terms of a race going on between a go process and a stop process[2]. If the stop process is faster than the go process, the saccade is inhibited. Otherwise, the response escapes inhibition and the saccade is executed in spite of a stop signal. In a particular instantiation of the race model[4] the go and the stop signal are assumed to generate neural activities linearly rising towards a fixed threshold with the rate of rise varying randomly from trial to trial. Whichever process crosses the threshold first will determine whether a saccade is executed or the response is inhibited. Note that stop signal reaction time (SSRT), defined as the time needed to cancel the planned movement, cannot be observed directly in behavioural experiments due to the nature of this paradigm: when the saccade is successfully inhibited, no response time can be measured. However, assuming that the go process activity grows independently of the stop process it is possible to estimate average SSRT from the behavioural data[2]. Estimates range from 70 ms to 150 ms depending on species (monkey or human) and on the experimental setup (e.g., stop signal modality).

For behavioural data, a test of this independence assumption is provided by comparing the cumulative distribution of the response times when no stop signal was presented with the cumulative distributions of response times under the various stop signal delays: the latter distributions should lie entirely above the former[5]. Recent evidence suggests that severe inhibition between go and stop processes, rather than independence, may occur[5-6,8]. This, however, implies that estimates of SSRT, being based on independence, may be severely biased.

Recent studies have shown that cells with movement- and fixation-related activity within the frontal eye fields (FEF) generate signals sufficient to control eye movement[1,9]. Since the race model allows to estimate the time to cancel a saccade (SSRT), special efforts have been made to validate the assumptions of the model through cell recording data. In particular, no difference was found in average neural activity associated with movements executed without or in spite of the stop signal[1]. While this finding is consistent with the independence assumption of the race model, a direct test of stochastic independence between the growth of go signal and stop signal related activities would require simultaneous recordings from fixation and buildup cells[1] in FEF. Note, however, that such independence is at odds with current models of the superior colliculus (SC) that assume that buildup movement neurons are inhibited directly or indirectly by fixation neurons within the SC[10].

The stop signal diffusion model presented here assumes a variable growth to a fixed threshold, as suggested by measurements of movement-related activity recorded in FEF[9] and SC[11]. Rather than claiming separate growths of go signal and stop signal related activities, however, the neural balance between gaze-holding and gaze-shifting mechanisms at any instance is represented by a single process that unfolds over time in a random walk-like[12] manner between two fixed criterion thresholds. Crossing the go-criterion (θ_{go}) results in the execution of the saccade, whereas crossing the stop-

criterion (θ_{stop}) results in a permanent cancellation of the planned movement to the go signal. Fig. 1 illustrates this mechanism. For each trial the onset of the go signal triggers a growth process represented by a stochastic trajectory drifting towards the upper boundary θ_{go} . In the absence of a stop signal the average trajectory (indicated by the line) has a positive slope resulting in a mean saccadic response time determined by the time point corresponding to the crossing of the go criterion (Fig. 1a). Presentation of a stop signal at a certain point after the go signal shifts the slope of the linear drift to a negative value (Fig. 1b). Those trajectories which have not yet crossed the upper boundary will then tend in the direction of the stop criterion θ_{stop} . Due to stochastic variability, however, individual trajectories may still cross the upper boundary resulting in a response in spite of the stop signal. Note that, like the race model, the diffusion model does not predict different rates of rise in activity for responses in non-cancelled trials and in latency-matched no-stop-signal trials.

Obviously, the percentage of saccades reaching the go criterion increases with stop signal delay. Thus, the diffusion model predicts an important qualitative feature of typical stop signal data: the later the stop signal is presented the less likely a successful inhibition of the saccade becomes (Fig. 1 b). The diffusion model also provides a quantitative fit to behavioural data. In a recent experiment in our lab[8] subjects were instructed to make a saccade to a visual target. In stop trials, an auditory stop signal was presented after one of three possible stop signal delays. Fig. 2 presents observed and predicted probability of inhibition and mean saccadic response time as a function of stop signal delay for one of the subjects. Note that the diffusion model is able to account for the non-monotonic relation between mean response time and stop signal delay. In contrast, the race model always predicts mean response time to increase with stop signal delay[2].

The race model´s estimate of about 100 ms for the SSRT means that the movement is actually cancelled in just 50 ms if one takes into account that the latency of the visual response to the foveal stop signal is no less than 50 ms. This amounts to five action potentials in a neuron firing 100 spikes/s. It has been suggested[13] that the potency of the stop signal is due to its foveal presentation directly activating the gaze fixation system. However, recent experiments with auditory[5] and peripheral visual[6] stop signals gave similar low estimates for SSRT under the race model.

The diffusion model´s prediction for the speed of the stopping process differ sharply from estimates based on the race model. For the diffusion model, *inhibition time* (IT) is defined as the interval from the presentation of the stop signal until the trajectory reaches the stop criterion. Thus, IT depends on the momentary level of activity towards the go threshold (represented by the trajectory location) at the time the stop signal is presented. It implies that the average time to cancel a saccade increases with stop signal delay. For example, estimates of IT for the subject presented are 530, 570, and 580 ms for stop signal delays of 70, 100, and 130 ms, respectively. Even if one subtracts some 30 ms for the latency of the response to the auditory stop signal, the resulting estimates are one order of magnitude larger than the estimates for SSRT under the race model. This discrepancy reflects an important difference between the diffusion and the race model: While both IT and SSRT are initiated by the presentation of a stop signal, the termination of IT in the diffusion model indicates that inhibition of the saccade has become certain, whereas termination of SSRT in the race model means that stop signal processing is finished, but actual inhibition of the saccade occurs only if the go signal processing has not been terminated earlier. Estimates for IT of about 500 ms in the diffusion model may appear implausible, but it should be noted that this includes the time to suppress the go signal activity completely. If, for example, go and stop signal are presented nearly simultaneously, resulting in a very high probability of

successful inhibition, estimates for IT can go down by an order of magnitude, depending on the value of the drift parameters.

Subjects have considerable leeway in performing the countermanding task. Explicit speed instructions and feedback (e.g., when response speed falls below a level achieved in a preceding experimental condition without any stop signals) do influence the subjects' performance[8]. In the diffusion model, a bias in favouring either stopping performance or response speed is easily accounted for by letting a trajectory in the activation space start from a level closer to the stop criterion or to the go criterion, respectively. Introducing this bias parameter allows to predict sequential effects like a higher probability of cancelling a saccade if the movement had failed to be cancelled on the previous trial[14] and thus relates this parameter to a hypothesized[15] monitoring and control system in the supplementary eye field (SEF).

Both the race and the diffusion model have been developed primarily to account for behavioural data. The increasing availability of cell recordings from structures involved in producing these behaviours opens up the possibility to identify the neural substrate of these models. At present, however, it seems premature to equate these models with the structure or function of a single specific brain structure, like FEF or SC. The diffusion model seems better suited than the race model to describe the behavioural data and to interpret the complex network of brain structures underlying the control of saccadic behaviour.

Methods

Stop signal experiment.

White dots (0.1°, 19.8 cd/m^2) presented on a dark background at a distance of 57 cm served as go signals, auditory noise signals (500-14,000 Hz, 72 dB SPL) presented over headphones as stop signals. Go signals were presented randomly 15° left or right from fixation. After each block of 72 trials, subjects were informed about their mean reaction times and proportion of successfully inhibited reactions. They were asked to maintain

response speed at a level previously assessed in two sessions without stop signals occurring. Three subjects served in 40 blocks of trials over 8 to 10 experimental sessions. Eye movements were measured with an infrared light reflecting system (IRIS, Skalar Medicals). Saccades larger than 1° into the direction of the target carried out during the first 500 ms after target presentation were counted as responses. Saccadic reaction time was determined as time difference between target onset and saccade onset.

Diffusion model.

The growth process in the diffusion model is represented by a standard Brownian motion (or Wiener) process[16] A(t) with drift rate $\mu(t)$ and two absorbing barriers θ_{GO} and θ_{STOP} . The process is time-inhomogeneous because the drift rate changes with the occurrence of the stop signal at $t = t_{SSD}$:

$$\mu(t) = \mu_{GO} \text{ for } t \leq t_{SSD} \text{ and } \mu(t) = \mu_{STOP} \text{ for } t > t_{SSD}.$$

The first-passage times are defined as

$$T_{GO} = \inf\{t: A(t) \geq \theta_{GO}; A(\tau) \geq \theta_{STOP} \text{ for all } \tau < t \}$$

$$T_{STOP} = \inf\{t: A(t) \leq \theta_{STOP}; A(\tau) \leq \theta_{GO} \text{ for all } \tau < t \},$$

with $\theta_{STOP} < A(0) < \theta_{GO}$. A finite-state Markov chain approximating the diffusion process was taken to compute these distributions using spectral methods[16-18]. Observable saccadic reaction time (SRT) was taken as $SRT = T_{GO} + c$ (c a motor constant) and observable inhibition probability as $Prob(T_{GO} < T_{STOP})$. For these data, A(0) was set to zero, assuming no bias. Estimated model parameters are the drift rates for the go signal and the stop signal (μ_{GO} and μ_{STOP}), the distance between the go and the stop threshold ($\theta_{GO} - \theta_{STOP}$), plus a (pre)motor constant (c). Given 7 data points the 4 parameter estimates were obtained by the least-squares method using the FMINSEARCH routine from the OPTIMUM MatLab® package.

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

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

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

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

5. Colonius, H., Özyurt, J. & Arndt, P. Countermanding saccades with auditory stop signals: Testing the race model model. *Vis. Res.* **41**, 1951–1968 (2001).

6. Asrress, K. N. & Carpenter, R. H. S. Saccadic countermanding: a comparison of central and peripheral stop signals. *Vis. Res.* **41**, 2645–2651.

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

8. Özyurt, J., Colonius, H. & Arndt, P. A. Countermanding saccades with auditory stop signals: effects of position and speed instruction. *Percept. Psychophys.* (submitted).

9. Hanes, D. P. & Schall, J. D. Countermanding saccades in macaque. *Vis. Neurosci*. **12**, 929-937 (1995).

10. Munoz, D. P. & Wurtz, R. H. Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cell. *J. Neurophysiol.* **73**, 2313-2333 (1995).

11. Hanes, D. P. & Paré, M. Neural control of saccade production studied with the countermanding paradigm: superior colliculus. *Soc. Neurosci. Abstr.* **24**, 418 (1998).

12. Luce, R. D. *Response time models : their role in inferring elementary mental organisation*. (Oxford Univ. Press, New York, 1986).

13. Schall, J. D. & Thompson, K. G. Neural selection and control of visually guided eye movements. *Ann. Rev. Neurosci.* **22,** 241-259 (1999).

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

15. Stuphorn, V., Taylor, T. L. & Schall, D. J. Performance monitoring by the supplementary eye field. *Nature* **408**, 857-860 (2000).

16. Bhattacharya, R. N. & Waymire, E. C. *Stochastic processes with applications*. (Wiley, New York, 1990).

17. Busemeyer, J. R. Townsend, J. T. Fundamental derivations from decision field theory. *Math. Soc. Sci.* **23**, 255-282.

18. Diederich, A. Dynamic stochastic models for decision making under time constraints. *J. Math. Psychol.* **41**, 260-274 (1997).

Acknowledgements

This work was supported DFG-SFB Neurokognition .

Correspondence and requests for materials should be addressed to H. C. (e-mail: hans.colonius@uni-oldenburg.de).

Figure 1 Three hypothetical trajectories in the activation space simulated by the diffusion model. **a**, In go trials, the drift rate is constant. A saccade is initiated when the trajectory crosses the upper threshold θ_{GO} for the first time. The line presents the average trajectory. **b**, In stop signal trials, the drift rate switches at presentation of a stop signal (SSD). The saccade is inhibited with certainty once the lower threshold θ_{STOP} has been crossed the first time. The average trajectory switches slope from positive to negative at SSD.

Figure 2 Observed data (squares) and predictions (diamonds) from diffusion model for one subject. **a**, Stop failure probability as a function stop signal delay. Rightmost point refers to no-stop-signal condition. **b**, Mean saccadic response time as function of stop signal delay.

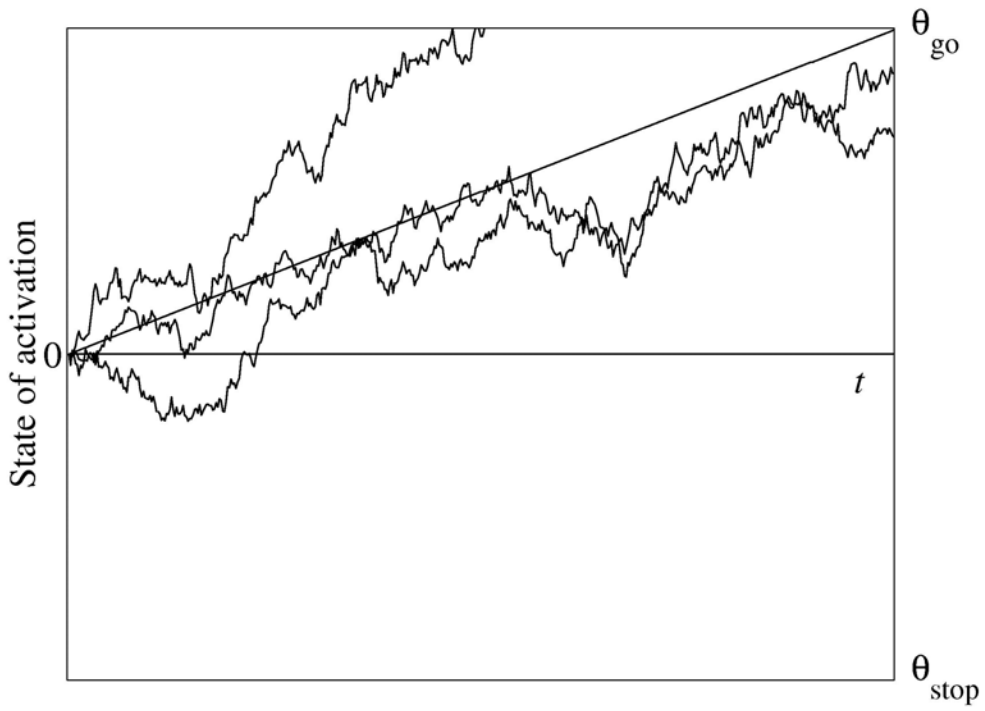
θgo
State of activation
0
t
θstop

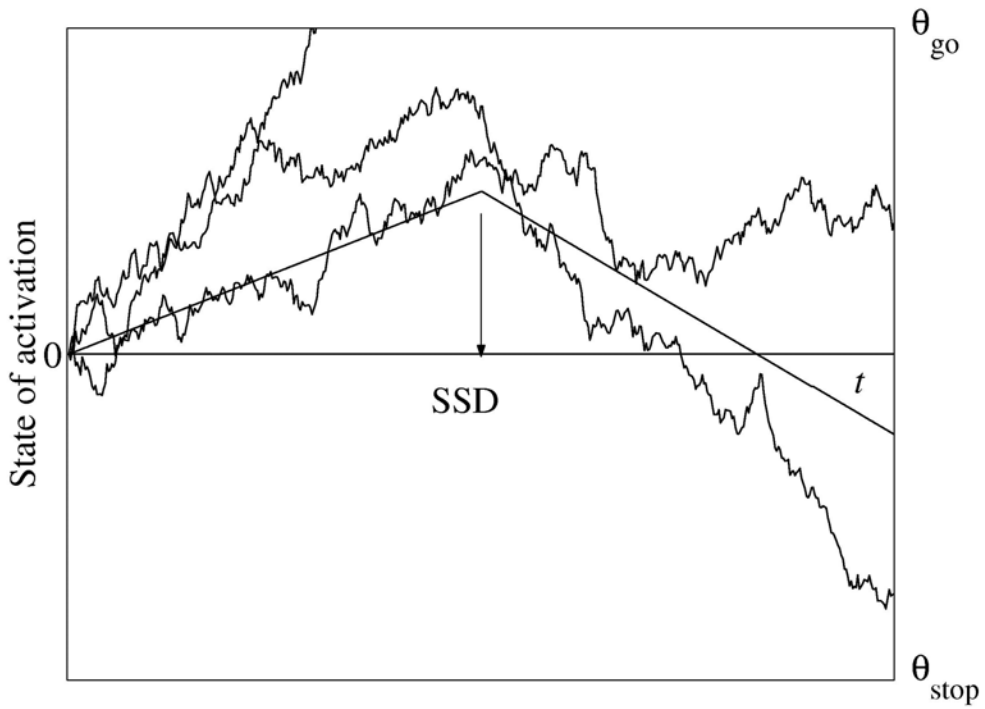
θgo
State of activation
0
SSD
t
θstop

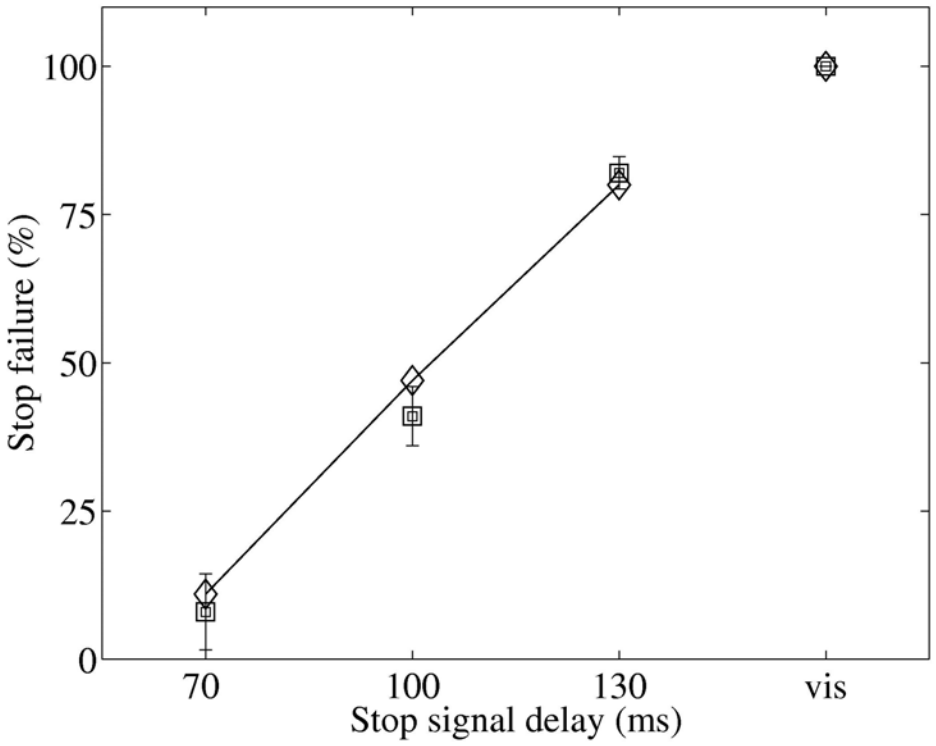
100
75
50
25
0
Stop failure (%)
70
100
130
vis
Stop signal delay (ms)

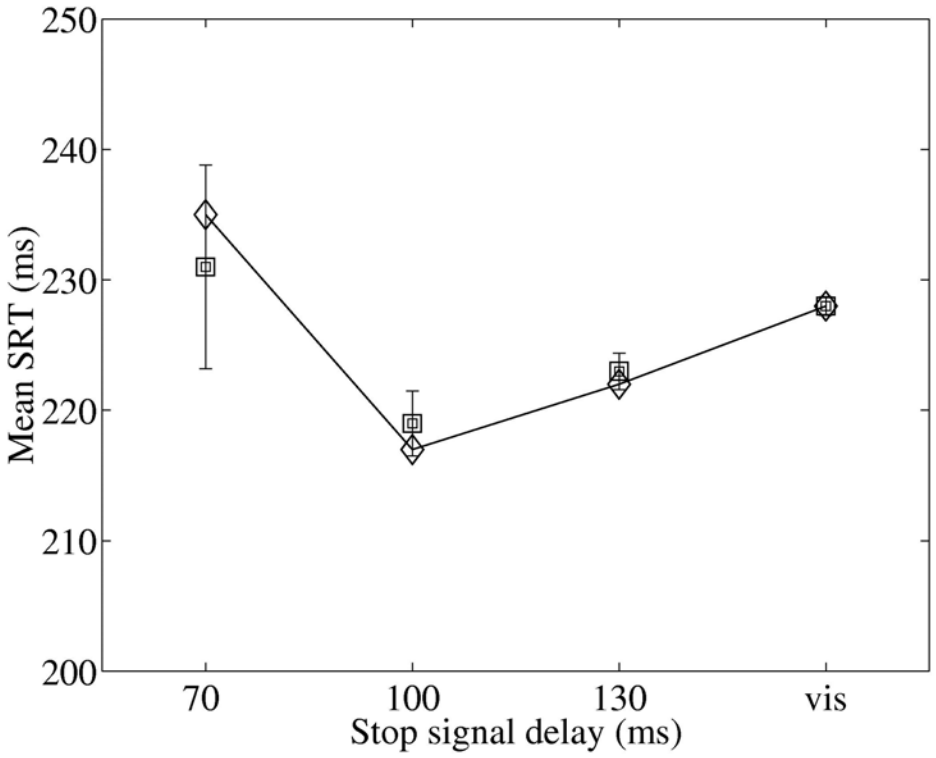
250
240
230
220
210
200
Mean SRT (ms)
70
100
130
vis
Stop signal delay (ms)