Review

Author for correspondence:
Imran Noorani
e-mail: in1@sanger.ac.uk

Towards a unifying mechanism for cancelling movements

Imran Noorani

Department of Physiology, Development and Neuroscience, University of Cambridge, Cambridge, CB2 3EG, UK

IN, 0000-0003-4289-2345

For precise motor control, we must be able to not only initiate movements with appropriate timing, but also stop them. The importance of stopping tended to be overlooked in research in favour of studying movement itself, so we are still only beginning to understand the neural basis of action cancellation. However, the development of models of behaviour in a wider range of tasks, and their relation to neural recordings has provided great insight into the underlying neurophysiology. Here we focus on developments of the linear approach to threshold with ergodic rate (LATER) model, relating these to complementary neurophysiological studies. It is tempting to consider that there may be a unifying mechanism for cancelling impending decisions in many contexts and how future efforts may clarify this possibility.

This article is part of the themed issue 'Movement suppression: brain mechanisms for stopping and stillness'.

1. Introduction

How we cancel our upcoming movements is a fundamental question in modern neuroscience. Cancelling, or stopping, actions is not just one process, but can occur in a number of different contexts [1,2]. To appreciate the importance of this behaviour, one only has to consider an everyday example of a driver about to press on the accelerator as the traffic light turns green before a pedestrian suddenly appears on the street—an immediate cancellation of the driver's foot movement is needed to avoid a disaster! Recent investigations into how the brain enables us to do this have been highly informative, and we focus our discussion on eye movement paradigms for stopping.

2. Measuring and modelling stopping behaviour

There are many kinds of stopping behaviour [1,2]. In order to study stopping behaviour in the laboratory, whether the subjects are humans or animals, a simple experimental task is needed. The best known and most established such task is called 'countermanding', also known as the 'Stop signal task'. In this paradigm, all trials start with the presentation of a 'Go' stimulus that a subject has been previously instructed to make a movement (such as a saccade) towards. On a subset of trials, a Stop signal is presented after the Go stimulus with a delay called the Stop signal delay. The subject is instructed to withhold the impending movement if the Stop signal is presented, but in some cases—failed Stop trials—they are unable to do so. An early model for countermanding proposed two neural units racing towards a threshold in Stop trials [3]. The idea of a neural decision unit accumulating activity towards threshold is well established in decision theory, with the accumulation itself representing sensory information providing evidence for a particular decision. In the case of countermanding, the first unit represents the 'Go' decision, and the second one (the Stop unit) starts accumulating after the Stop signal delay if the Stop signal is presented. Failed Stop trials can then be explained by the Go unit reaching threshold before the Stop unit and therefore escaping inhibition, whereas successfully cancelled trials are due to the Stop unit winning the race. A feature of this model is that the two units in this model move towards threshold independently from one another.
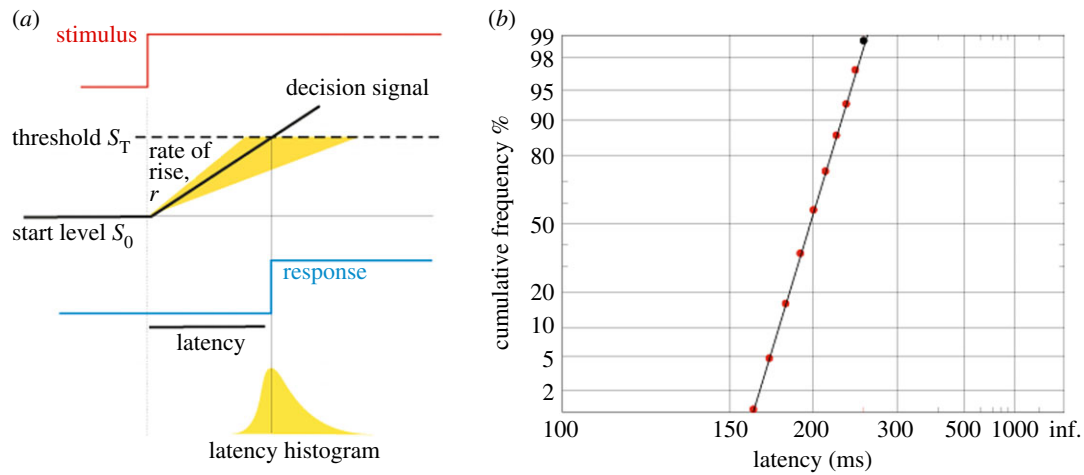
Figure 1. The LATER model. (a) A single LATER unit that accumulates a decision signal at a linear rate until threshold is reached, at which point the decision is made. (b) Simulated reciprobit plot showing the distribution of saccadic reaction times in a simple step task as predicted by a single LATER unit.

The independent race model does not posit definitively what the underlying neural mechanism for stopping might be, and later neurophysiological studies of neural correlates of Stop and Go processes in the frontal eye fields and superior colliculus (SC) of monkeys have suggested that a network of interacting rather than independent neurons may generate the countermanding behaviour [4,5]. Boucher *et al.* [6,7] made an attempt to reconcile these neurophysiological findings with race model analysis by proposing the 'interactive race model' for countermanding, in which stopping occurs through a network of mutually inhibitory Go and Stop neurons [6,7]. As with the independent model from Logan, Cowan and Davis [3], the Go unit is activated following a short processing delay after stimulus presentation, and the Stop unit is later activated after another processing delay if the Stop signal is also presented. The key difference here is that the two neural units laterally inhibit one another. Model simulations implied that the Stop unit would strongly inhibit Go unit very late, therefore producing reaction times (RTs) almost identical to those of the independent race model.

Salinas and Stanford have recently proposed a similar model that explains the behavioural data and neural recordings, the difference being that their model puts emphasis on the time needed to detect the Stop signal; in their model, a saccadic decision signal begins to rise towards a threshold, but presentation of the Stop signal leads to detection with a particular speed and reliability; the Go signal then declines towards zero [8,9]. The similarity between these various models of countermanding highlights an important problem in the field: 'model mimicry'; it is difficult to distinguish reliably between them by simple behavioural and neuronal measurements, without further complex analysis [10,11]. Further work is also needed to determine whether similar models can explain cancelling of limb in addition to eye movements [12].

3. The LATER model

A type of decision-making model that has become increasingly popular is the 'linear approach to threshold with ergodic rate' (LATER) model [13–16]. In this model, a potential decision is represented by a unit that starts from a baseline point (S_0) and then rises linearly towards threshold (S_T) with a mean rate of rise (μ) and variance (σ) (figure 1). The mean rate of rise is determined by the rate at which information for a particular decision is supplied, rather like accumulating evidence for a given hypothesis; this decision signal represents the log likelihood of the hypothesis that a particular decision should be made (such as looking towards a visual stimulus if it is present). An important feature is that although the mean rate of rise is fixed for a decision unit, the rate of rise varies randomly between trials, allowing the model to account for the random variability in reaction times. The popularity of this model stems largely from its conceptual simplicity and ability to simulate reaction time (RT) data for decision tasks with relatively few free parameters, as well as being grounded in Bayesian decision theory. Early evidence for the accuracy of LATER in representing neural processes came from a behavioural study in which subjects made saccades to visual targets with varying prior probabilities of the target appearing. Altering the prior probability S_0 should result in a specific change, a 'swivel', in the subjects' RT distribution; and indeed this is what was observed [13]. The LATER model itself has been recently reviewed elsewhere [1,2,17].

In a study aimed at modelling countermanding with LATER, Hanes & Carpenter [18] measured saccadic latencies in a group of human subjects performing this task. They found that the Stop signal reaction times of these human subjects were between 125 and 145 ms. Using a race model with two LATER units that independently rise towards threshold linearly with a rate that varies randomly between trials, they were able to accurately capture not only the error rates and mean reaction times, but also the entire statistical distributions of these RTs in quantitative detail. The Go unit is initiated by the Go stimulus following a short delay, and the Stop unit would similarly be activated by the Stop signal, with the winner of the race between these two units determining whether the impending saccade would successfully be cancelled or not. The entire RT distributions were simulated with only four parameters—μ_{go}, σ_{go}, μ_{stop} and σ_{stop}. The accuracy with which this model predicted the data is a strong testimony to the power of this type of model in simulating human decision-making behaviour. Although modelling data based on mean reaction times or a small number of 'quintiled' reaction times can be perfectly acceptable in many cases, this approach can often be too
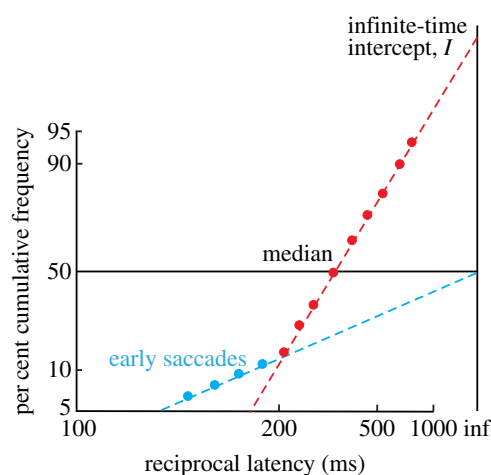
Figure 2. Early saccades. There are two separate latency distributions: a main distribution, and an early one; the two are modelled by separate, competing LATER units, with the parameters μ and σ for the early unit being very large.

simplistic and lead to oversights in important points in the data that could point towards alternative underlying neural processes. In particular, plotting full RT distributions, such as with 10 ms bins, allows one to discern subtle differences at particularly short or long latencies where populations of responses may not follow the main distribution, and which tend to be hidden in conventional linear plots. For example, this approach allows 'early' saccades to be seen more clearly (figure 2). It is important to note these extra populations of responses because a thorough model of the data should also be able to account for these [19].

4. Complex stopping experimental tasks

Although countermanding is the simplest and most studied experimental task requiring cancellation of impending actions, other tasks have also more recently emerged as likely to require stopping. These tasks differ, not only in terms of the actual choices involved in that there are different potential responses, but also in terms of the model architecture and therefore the underlying neural decision processes. One such task is the Go/No-Go task. A study aimed at devising a race model for the behaviour in this task employed two different coloured visual stimuli: the human subjects were instructed to make a saccade towards the target if it was red but not to do so if the target was blue [20,21]. Subjects made a variable number of errors, saccades towards the target even if it was the wrong colour. Observing the distributions of reaction times of correct and error responses, it is evident that at lower latencies the two distributions are essentially identical, suggesting indiscriminate responding to the target irrespective of whether or not it is the correct colour. With longer latencies, the number of errors tends to saturate, whereas the correct responses become more frequent. This suggests that the subjects are using the colour information of the target to decide whether to make a saccade or not only at longer latencies, most likely because there is insufficient time for cortical processing of this colour information at very short reaction times. An important feature noticed from the RT distributions that would have otherwise been missed had only mean RTs been analysed is the bimodality in correct responses: there were early and late responses,

often with a short period in between when very few responses occurred (see figure 3, for an example). This feature made it clear that to model the decision-making behaviour, a Stop unit is needed, accounting for the complete suppression of error responses after the delay for cortical processing of colour and for the bimodality in the RTs for correct responses. The Go/No-Go task model is therefore very similar to that of countermanding, both having a Go and a Stop unit, except in this task there is also a further decision Go unit for the correct colour target that is activated after a cortical colour processing delay of 60 ms [22,23]. This is a satisfying demonstration of how sensory information (colour) may be used by the brain during the decision-making process in order to modify its outcome.

Another more complex eye movement task that involves cancelling an unwanted action in favour of a more considered one is the anti-saccade task. This requires the subject to look in the opposite direction to a given visual stimulus: if, for example, a target appears on the left, the subject must resist the tendency to look left and instead look right. As before, subjects tend to make errors (saccade towards the target instead of away from it) probably because the tendency to look towards a novel stimulus is prepotent. Early neural decision models for anti-saccades involved a simple race between two Go units—one for the pro-saccade towards the target, and one for the anti-saccade, with the winner in a trial determining whether an anti-saccade or error would occur [24–27]. However, a more recent study investigating the complete RT distributions of both errors and anti-saccades has attempted to simulate the data in quantitative detail [28]. Each subject performed both a control step (pro-)saccadic task as well as the anti-saccade task, allowing the best fit parameters for μ and σ of the main decision unit to be estimated. These parameters for each individual were then used to model their corresponding anti-saccade data for two key units—the anti-saccade and pro-saccade unit. In this way, the model was highly constrained, with much of the behaviour already predicted by the RT distributions in control trials. A feature of this model was that the anti-saccade unit was delayed for an extra 50 ms before beginning to rise, corresponding to the time needed for the cerebral cortex to invert the spatial location of the target (needed for knowledge of the anti-saccade target) [29]. A simple two-unit model, however, failed to capture RT distributions accurately but did so with the inclusion of a Stop unit that would suppress the pro-saccade unit. The key observation in the latency distributions here is that there was a significant bimodality in responses, with many early errors and then mostly late anti-saccades, suggesting there is a distinct process suppressing errors before anti-saccades arise. This provided strong justification for the importance of stopping in the complex decision task and further highlights the need for detailed quantitative analysis for supporting neural models (figure 3).

In trials of the anti-saccade task where the subject fails to cancel the impending pro-saccade instead of making an anti-saccade, they almost always make a correction saccade after the error. Analysis of the RT distributions for these corrections reveals that they are faster than making anti-saccades afresh. The original LATER anti-saccade model successfully predicted these with just a simple modification to the model requiring no extra free parameters: the anti-saccade unit resets and starts racing towards threshold after the pro-saccade unit wins the initial race because the Stop unit failed to reach threshold
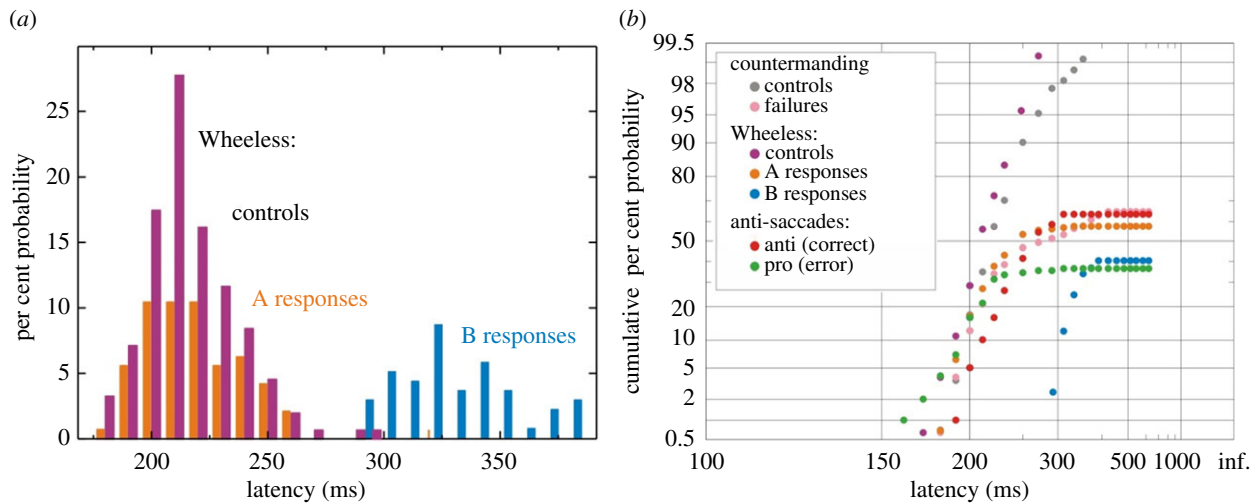
Figure 3. RT distributions of Stop tasks. (*a*) Histogram of latencies of type A and B responses in the Wheeless task; note the bimodality in the distribution, with a gap between the end of most A responses and the beginning of most B responses, suggestive of a Stop unit acting on the A responses. This bimodality is also common to the Go/No-Go task and anti-saccades. (*b*) Reciprobits plots of countermanding, Wheeless and anti-saccade responses demonstrating that error responses typically occur earlier than correct ones. Note: the data presented here were from one individual subject, previously unpublished.
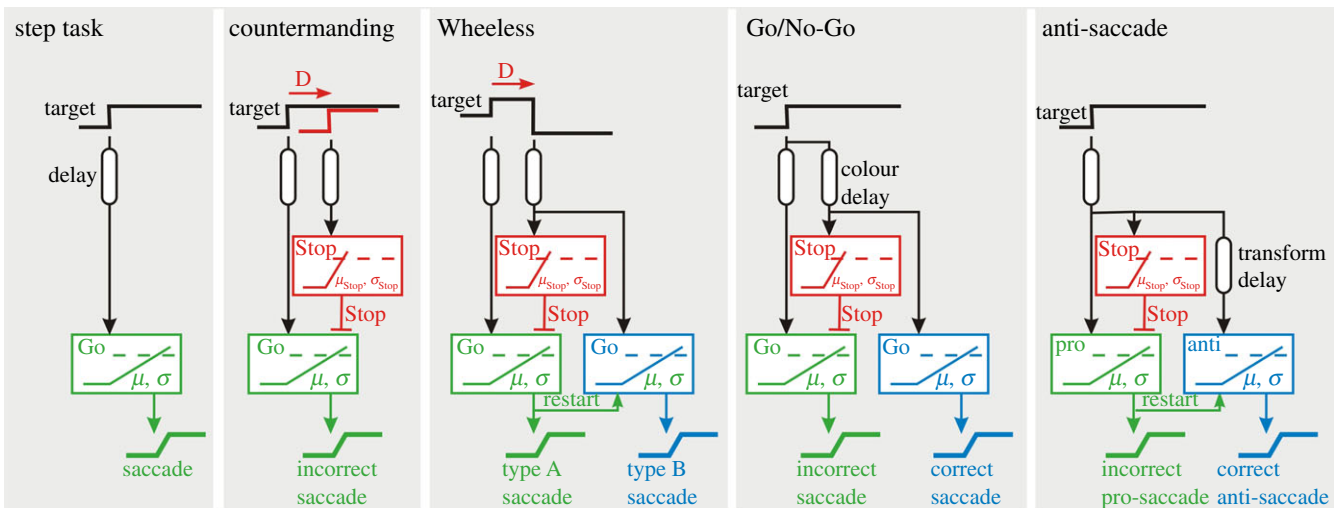

Figure 4. LATER-based decision models for the step task, and four types of Stop task of which countermanding has the simplest model. All essentially have a race between a Go and a Stop unit, and it can be seen that models for Wheeless, Go/No-Go and anti-saccades are very similar with a key difference being in the delay before the Stop unit starts rising.

quickly enough. Since there is then no need for further cortical processing for inversion of spatial location, it is clear why the correction anti-saccade is faster than normal anti-saccades [30,31], and is an example of a seemingly complex stopping task that can be accurately conceptualized by just a few simple racing decision processes.

A further advanced decision task requiring stopping is the Wheeless task [32]. This is closely related to the anti-saccade task in that one must stop an impending response in favour of another, but is slightly less artificial because here the desired response is triggered by an external stimulus (rather than an internal one as for anti-saccades). This paradigm (and closely related variants such as the redirect task) has been used in clinical studies to identify deficits in Stop signal reaction times in patients with Parkinson's disease and schizophrenia [33,34]. The subject is presented with a visual target either on the left or right and has to make a saccade directly to the target; for some trials, this is the end, but for others a second visual stimulus is presented in the opposite direction following a delay (D). No specific instructions are given to the subjects beforehand,

who are only told to follow the visual targets with their eyes. Once again, their behaviour is stochastic: in some experimental trials, they direct their eyes to the first target and then to the second (a type A response), whereas in others they make a saccade straight to the second target (B). As might be expected, the bigger D is, the greater the probability of making a type A response. The importance of including a Stop unit in general two-step paradigms of this kind was highlighted in earlier models [35,36]. Initial LATER model attempts with a Stop unit and two Go units (one for each possible response) failed to account for the unexpectedly high rates of cancellation in Wheeless; even if the rate of rise of the Stop unit was infinite, it would not predict this high stopping efficiency in certain circumstances. The implication is that stopping is given a certain advantage in speed over the other two decision units here; it was found that if the Stop unit started rising towards threshold 10–20 ms before the Go units, then the model accurately predicts both error rates and full RT distributions (figures 3 and 4). It is quite possible, therefore, that the neural Stop processes can be activated particularly fast in

conditions when it is important to cancel impending errors quickly [37]. Indeed, recent neural recordings have emphasized that onset of neural activity is also crucial in determining RT and therefore should be carefully considered in designing accurate neural race models of behaviour [8,38].

A critical question arising from these recent modelling efforts is whether each of these tasks employs its own separate, neurally distinct Stop decision process, or whether the Stop unit is actually identical in the various paradigms. Comparing the tasks, there is a fundamental difference in how the Stop process is triggered: in countermanding, Go/No-Go and Wheeless tasks, there is an external Stop signal of some sort presented in a given trial that instructs the subject that they need to cancel the impending eye movement, whereas with anti-saccades this Stop signal is generated internally because the subject knows before a given trial that they need to suppress their tendency to look at the visual stimulus. Could this difference in how stopping is triggered mean that the Stop decision process itself is different, at least in anti-saccades compared with the other tasks? The observation that, in the Wheeless task, the Stop unit can be activated faster than expected in certain conditions again raises the possibility that there is another Stop decision unit acting in these circumstances. The alternative explanation is that the difference lies purely in the amount of processing needed to detect and evaluate the Stop signal in each case, but that the neural Stop unit is itself identical. Although with current evidence we cannot answer these questions, it is tempting to consider that all these tasks embody a single Stop mechanism racing against Go units, and further that the values of its parameters (μ and σ) are rather similar between the tasks.

5. Neural evidence

Race models of decisions in simple as well as complex tasks have generated much interest in their possible implementation by the underlying neural processes of Stop and Go. Early neuronal recordings in the macaque frontal eye fields revealed movement neurons whose activity rises essentially linearly before the saccade, and fixation cells whose activity increased following the Stop signal and before the Stop signal RT, suggesting potential neuronal correlates of the decision units in typical models [4]. It is likely that these signals are passed downstream to the SC, where neurons fire less when saccades are successfully countermanded with the implication that their firing must reach a threshold level to trigger a saccade, although of course the SC may simply be reflecting the activity of the cortical regions [5].

The neural origin of the Stop process has proved to be more elusive than for Go. Some light has been shed on this by Stuphorn & Schall [39], who have performed an elegant experimental study in macaque monkeys to investigate the role of the supplementary eye field (SEF) in countermanding. The SEF has close anatomical connections with the oculomotor system (specifically with the frontal eye field, basal ganglia, SC and brainstem), potentially allowing modification of eye movements, although these connections are unlikely to be strong enough to directly initiate saccades. However, previous work had already established that the SEF demonstrates error- and reward-related signals. Stuphorn and Schall weakly stimulated different regions of the SEF while the monkeys performed a saccadic countermanding task; for many regions, microstimulation resulted in fewer failed Stop trials, as well as increasing the delay in mean saccade latency for Go trials. This effect was specific for countermanding, since the same stimulation did not increase saccade latency for monkeys performing a simple step (Go) task. These effects on error rates and latencies in countermanding for SEF stimulation can potentially be explained by a more active Stop unit in the independent race model, which would suppress faster errors and therefore lengthen mean reaction times (due to a reduction in faster saccades that are now cancelled). This study therefore demonstrates a role for the SEF in saccade inhibition in the saccade unit, rather akin to the Stop unit of race models. A related study in humans using fMRI showed that the frontal eye fields were more active on Stop signal trials irrespective of successful cancellation, whereas the SEF was more active specifically in Stop signal trials that were successfully cancelled [40].

The basal ganglia, which of course also contain oculomotor as well as more general motor machinery, represent another region where neural representations of both Stop and Go decision units have been implicated. In a countermanding study in rats that used lateral head movements to indicate decisions, Schmidt et al. [41] recorded from neurons of various basal ganglia nuclei. Subthalamic nucleus (STN) neurons responded to the Stop cue irrespective of whether the rat successfully cancelled the impending head movement. In contrast, neurons in the substantia nigra (SNr) downstream from the STN responded to the Stop signal only in trials where stopping was successful, and it is significant that this neuronal response occurred before the Stop signal RT, suggesting it is early enough to influence behaviour and potentially cause the stopping itself. The proposal is therefore that the STN–SNr pathway represents the Stop unit of countermanding. Further recordings have demonstrated that striatal input to the SNr caused pauses in SNr firing correlating with Go trials, suggesting the striatal–SNr pathway is related to the Go unit. In this way, groups of neurons in these two regions could be seen to race against one another, much akin to the two decision units of race models, with the motor outcome of a given trial depending on which pathway wins the race in that particular trial. This is an excellent example of how early race models have in many cases been vindicated by careful neural recordings in subjects performing decision tasks. One important question arising from this study is whether these basal ganglia pathways are the decision-makers themselves or alternatively (and more likely) they are a reflection of decisions being made in higher regions, such as the SEF for the case of stopping [4,30,31,39]. Moreover, it is unclear whether this race between Stop and Go processes in the basal ganglia also generates behaviour in other types of stopping task, such as anti-saccades and the Wheeless paradigm. As discussed earlier, it is not known whether the Stop unit in these various tasks is the same or whether each task has its own neurally distinct Stop process, and careful neural recording studies in these other Stop paradigms are needed to clarify whether this race in the basal ganglia can also explain them in a similar fashion.

6. Conclusion and future directions

The importance of cancelling impending undesired actions cannot be overestimated: just imagine how many 'wrong'

actions we would perform if there were no neural mechanisms to suppress them. Through simple laboratory paradigms such as countermanding, we have been able to conceptualize neural Stop processes with race decision models. These have progressed from those that characterized very basic features of behaviour such as mean reaction times and error rates, to more recent ones that are now able to fully predict the quantitative details of behaviour with surprising accuracy. Neurophysiological recordings have been crucial in justifying the processes posited in these models as well as attempting to localize where in the brain these are occurring. It is still unclear where in the brain these processes originate, and it may well be that the decision units in race models are represented by functionally linked groups of neurons in distributed regions of the brain that must act together in order to generate a decision. Moreover, the recent work identifying distinct race pathways in the basal ganglia for countermanding raises questions as to whether there is just one Stop process that can also apply to more complex behaviours like anti-saccades, or whether there are multiple Stop processes represented in the brain for cancelling unwanted actions in different contexts, although this important question has undoubtedly been considered before

[42–46]. In my opinion, the fact that recent work has demonstrated that these seemingly disparate Stop tasks can all be conceptualized in a relatively simple manner, with similar race models, suggests that there may well be a unifying mechanism for stopping impending movements. Of course, the tasks discussed in this paper all involve eye movements, which provide useful paradigms for experimental study in a laboratory, but more work would be needed to determine whether similar results and conclusions are drawn from limb movements. Further insights will undoubtedly be gleaned from detailed modelling of Stop behaviours, correlation with neural recordings in these more advanced tasks and comparison of data between these tasks to identify potentially unifying decision processes. In particular, studies with larger numbers of subjects that each perform several of these Stop tasks and have their behaviour simulated with race models will allow unprecedented valid comparisons of the neural processes underlying these paradigms.

References

1. Noorani I, Carpenter RH. 2016 The LATER model of reaction time and decision. *Neurosci. Biobehav. Rev.* **64**, 229–251. (doi:10.1016/j.neubiorev.2016.02.018)

2. Noorani I, Carpenter RHS. 2017 Not moving: the fundamental but neglected motor function. *Phil. Trans. R. Soc. B* **372**, 20160190. (doi:10.1098/rstb.2016.0190)

3. Logan GD, Cowan WB, Davis KA. 1984 On the ability to inhibit simple and choice reaction time responses: a model and a method. *J. Exp. Psychol. Hum. Percept. Perform.* **10**, 276–291. (doi:10.1037/0096-1523.10.2.276)

4. Hanes DP, Patterson 2nd WF, Schall JD. 1998 Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J. Neurophysiol.* **79**, 817–834.

5. Pare M, Hanes DP. 2003 Controlled movement processing: superior colliculus activity associated with countermanded saccades. *J. Neurosci.* **23**, 6480–6489.

6. Boucher L, Palmeri TJ, Logan GD, Schall JD. 2007 Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychol. Rev.* **114**, 376–397. (doi:10.1037/0033-295X.114.2.376)

7. Boucher L, Stuphorn V, Logan GD, Schall JD, Palmeri TJ. 2007 Stopping eye and hand movements: are the processes independent? *Percept. Psychophys.* **69**, 785–801. (doi:10.3758/BF03193779)

8. Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E. 2010 Perceptual decision making in less than 30 milliseconds. *Nat. Neurosci.* **13**, 379–385. (doi:10.1038/nn.2485)

9. Salinas E, Stanford TR. 2013 The countermanding task revisited: fast stimulus detection is a key determinant of psychophysical performance.

J. Neurosci. **33**, 5668–5685. (doi:10.1523/JNEUROSCI.3977-12.2013)

10. Cutsuridis V. 2017 Behavioural and computational varieties of response inhibition in eye movements. *Phil. Trans. R. Soc. B* **372**, 20160196. (doi:10.1098/rstb.2016.0196)

11. Schall JD, Palmeri TJ, Logan GD. 2017 Models of inhibitory control. *Phil. Trans. R. Soc. B* **372**, 20160193. (doi:10.1098/rstb.2016.0193)

12. Pouget P, Murthy A, Stuphorn V. 2017 Cortical control and performance monitoring of interrupting and redirecting movements. *Phil. Trans. R. Soc. B* **372**, 20160201. (doi:10.1098/rstb.2016.0201)

13. Carpenter RH, Williams ML. 1995 Neural computation of log likelihood in control of saccadic eye movements. *Nature* **377**, 59–62. (doi:10.1038/377059a0)

14. Reddi BA, Carpenter RH. 2000 The influence of urgency on decision time. *Nat. Neurosci.* **3**, 827–830. (doi:10.1038/77739)

15. Reddi BA, Asrress KN, Carpenter RH. 2003 Accuracy, information, and response time in a saccadic decision task. *J. Neurophysiol.* **90**, 3538–3546. (doi:10.1152/jn.00689.2002)

16. Carpenter RH, Reddi BA, Anderson AJ. 2009 A simple two-stage model predicts response time distributions. *J. Physiol.* **587**, 4051–4062. (doi:10.1113/jphysiol.2009.173955)

17. Noorani I. 2014 LATER models of neural decision behavior in choice tasks. *Front. Integr. Neurosci.* **8**, 67. (doi:10.3389/fnint.2014.00067)

18. Hanes DP, Carpenter RH. 1999 Countermanding saccades in humans. *Vision Res.* **39**, 2777–2791. (doi:10.1016/S0042-6989(99)00011-5)

19. Noorani I, Carpenter RH. 2011 Full reaction time distributions reveal the complexity of neural

decision-making. *Eur. J. Neurosci.* **33**, 1948–1951. (doi:10.1111/j.1460-9568.2011.07727.x)

20. Sperling G, Dosher B. 1986 Strategy and optimization in human information processing. In *Handbook of human perception & performance*, vol. 1 (eds K Boff, L Kaufman, J Thomas), pp. 2.1–2.65. New York, NY: Wiley.

21. Noorani I, Gao MJ, Pearson BC, Carpenter RH. 2011 Predicting the timing of wrong decisions with LATER. *Exp. Brain Res.* **209**, 587–598. (doi:10.1007/s00221-011-2587-1)

22. Thompson KG, Hanes DP, Bichot NP, Schall JD. 1996 Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J. Neurophysiol.* **76**, 4040–4055.

23. Schall JD, Bichot NP. 1998 Neural correlates of visual and motor decision processes. *Curr. Opin. Neurobiol.* **8**, 211–217. (doi:10.1016/S0959-4388(98)80142-6)

24. Mokler A, Fischer B. 1999 The recognition and correction of involuntary prosaccades in an antisaccade task. *Exp. Brain Res.* **125**, 511–516. (doi:10.1007/s002210050709)

25. Kristjansson A, Chen Y, Nakayama K. 2001 Less attention is more in the preparation of antisaccades, but not prosaccades. *Nat. Neurosci.* **4**, 1037–1042. (doi:10.1038/nn723)

26. Trappenberg TP, Dorris MC, Munoz DP, Klein RM. 2001 A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *J. Cogn. Neurosci.* **13**, 256–271. (doi:10.1162/089892901564306)

27. Cutsuridis V, Smyrnis N, Evdokimidis I, Perantonis S. 2007 A neural model of decision-making by the superior colicullus in an antisaccade task.

Neural Netw. **20**, 690–704. (doi:10.1016/j.neunet. 2007.01.004)

28. Noorani I, Carpenter RH. 2013 Antisaccades as decisions: LATER model predicts latency distributions and error responses. *Eur. J. Neurosci.* **37**, 330–338. (doi:10.1111/ejn.12025)

29. Zhang M, Barash S. 2000 Neuronal switching of sensorimotor transformations for antisaccades. *Nature* **408**, 971–975. (doi:10.1038/35048530)

30. Noorani I, Carpenter RH. 2014 Basal ganglia: racing to say no. *Trends Neurosci.* **37**, 467–469. (doi:10. 1016/j.tins.2014.07.003)

31. Noorani I, Carpenter RH. 2014 Re-starting a neural race: anti-saccade correction. *Eur. J. Neurosci.* **39**, 159–164. (doi:10.1111/ejn.12396)

32. Wheeless Jr LL, Boynton RM, Cohen GH. 1966 Eye-movement responses to step and pulse-step stimuli. *J. Opt. Soc. Am.* **56**, 956–960. (doi:10.1364/JOSA. 56.000956)

33. Joti P, Kulashekhar S, Behari M, Murthy A. 2007 Impaired inhibitory oculomotor control in patients with Parkinson's disease. *Exp. Brain Res.* **177**, 447–457. (doi:10.1007/s00221-006-0687-0)

34. Thakkar KN, Schall JD, Logan GD, Park S. 2015 Response inhibition and response monitoring in a saccadic double-step task in schizophrenia. *Brain Cogn.* **95**, 90–98. (doi:10.1016/j.bandc.2015.01.007)

35. Camalier CR, Gotler A, Murthy A, Thompson KG, Logan GD, Palmeri TJ, Schall JD. 2007 Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. *Vision Res.* **47**, 2187–2211. (doi:10.1016/j.visres. 2007.04.021)

36. Ramakrishnan A, Sureshbabu R, Murthy A. 2012 Understanding how the brain changes its mind: microstimulation in the macaque frontal eye field reveals how saccade plans are changed. *J. Neurosci.* **32**, 4457–4472. (doi:10.1523/ JNEUROSCI.3668-11.2012)

37. Noorani I, Carpenter RH. 2015 Ultrafast initiation of a neural race by impending errors. *J. Physiol.* **593**, 4471–4484. (doi:10.1113/JP270842)

38. Pouget P, Logan GD, Palmeri TJ, Boucher L, Pare M, Schall JD. 2011 Neural basis of adaptive response time adjustment during saccade countermanding. *J. Neurosci.* **31**, 12 604–12 612. (doi:10.1523/ JNEUROSCI.1868-11.2011)

39. Stuphorn V, Schall JD. 2006 Executive control of countermanding saccades by the supplementary eye field. *Nat. Neurosci.* **9**, 925–931. (doi:10.1038/ nn1714)

40. Curtis CE, Cole MW, Rao VY, D'Esposito M. 2005 Canceling planned action: an FMRI study of countermanding saccades. *Cereb. Cortex* **15**, 1281–1289. (doi:10.1093/cercor/bhi011)

41. Schmidt R, Leventhal DK, Mallet N, Chen F, Berke JD. 2013 Canceling actions involves a race between basal ganglia pathways. *Nat. Neurosci.* **16**, 1118–1124. (doi:10.1038/nn.3456)

42. Rubia K *et al.* 2001 Mapping motor inhibition: conjunctive brain activations across different versions of go/no-go and stop tasks. *Neuroimage* **13**, 250–261. (doi:10.1006/nimg.2000.0685)

43. Verbruggen F, Logan GD. 2009 Models of response inhibition in the stop-signal and stop-change paradigms. *Neurosci. Biobehav. Rev.* **33**, 647–661. (doi:10.1016/j.neubiorev.2008.08.014)

44. Aron AR. 2011 From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol. Psychiatry* **69**, e55–e68. (doi:10.1016/j.biopsych.2010.07.024)

45. Levy BJ, Wagner AD. 2011 Cognitive control and right ventrolateral prefrontal cortex: reflexive reorienting, motor inhibition, and action updating. *Ann. NY Acad. Sci.* **1224**, 40–62. (doi:10.1111/j. 1749-6632.2011.05958.x)

46. Wiecki TV, Frank MJ. 2013 A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychol. Rev.* **120**, 329–355. (doi:10.1037/ a0031542)