# SCIENTIFIC REP☼RTS

# Saccadic inhibition interrupts ongoing oculomotor activity to enable the rapid deployment of alternate movement plans

Emilio Salinas & Terrence R. Stanford

Diverse psychophysical and neurophysiological results show that oculomotor networks are continuously active, such that plans for making the next eye movement are always ongoing. So, when new visual information arrives unexpectedly, how are those plans affected? At what point can the new information start guiding an eye movement, and how? Here, based on modeling and simulation results, we make two observations that are relevant to these questions. First, we note that many experiments, including those investigating the phenomenon known as "saccadic inhibition", are consistent with the idea that sudden-onset stimuli briefly interrupt the gradual rise in neural activity associated with the preparation of an impending saccade. And second, we show that this stimulus-driven interruption is functionally adaptive, but only if perception is fast. In that case, putting on hold an ongoing saccade plan toward location *A* allows the oculomotor system to initiate a concurrent, alternative plan toward location *B* (where a stimulus just appeared), deliberate (briefly) on the priority of each target, and determine which plan should continue. Based on physiological data, we estimate that the advantage of this strategy, relative to one in which any plan once initiated must be completed, is of several tens of milliseconds per saccade.

A fundamental function of oculomotor circuits is to determine where the eyes should look next and produce the appropriate eye movement. Neurally, each saccade is the culmination of a motor planning process whereby the firing rates of movement-related neurons rise gradually, monotonically, until the ramping activity of the population reaches a particular threshold level, at which point the plan becomes an uncancelable command and the saccade is triggered[1–6]. A timing conflict is likely to arise because this rise-to-threshold process unfolds over a sizable period of time, but new, potentially critical visual information may arrive at any moment. Furthermore, due to two prominent features of oculomotor circuits, this conflict must be extremely common.

On one hand, at least four lines of evidence indicate that saccades are programmed continuously, which is to say that there is always a subset of movement-related neurons (encoding a particular movement vector) that is steadily increasing its activity toward threshold. First, both humans and monkeys generate saccades at similar rates under a wide variety of viewing conditions[7–9], even in the dark[10,11]. Second, during free viewing, activity recorded in the frontal eye field (**FEF**) shows variations in firing rate, from one saccade to the next, occurring around a relatively high mean compared to that seen after prolonged fixation[9,12,13], as if the movement-related activity never fell far below threshold. Third, motor plans are not contingent on the completion of a target selection process[5,14–17] and may proceed in parallel[5,6,9,16,18–20]. And fourth, motor plans are strongly driven by current task contingencies and behavioral goals[21,22]; that is, by internal, already acquired information, and not necessarily by the visual stimuli currently in view.

On the other hand, saccades are rapidly and preferentially drawn to physically salient or highly relevant visual stimuli[8,23–30], such as a traffic light turning red, or a mouse suddenly scurrying across the kitchen floor. The draw of such stimuli is often described as "attentional capture" (refs[31,32]) or, when particularly potent, "oculomotor capture" (refs[25,26]). Stimuli that capture attention engage dedicated neural mechanisms that are highly sensitive

Department of Neurobiology and Anatomy, Wake Forest School of Medicine, 1 Medical Center Blvd., Winston-Salem, NC, 27157-1010, USA. Correspondence and requests for materials should be addressed to E.S. (email: esalinas@wakehealth.edu)
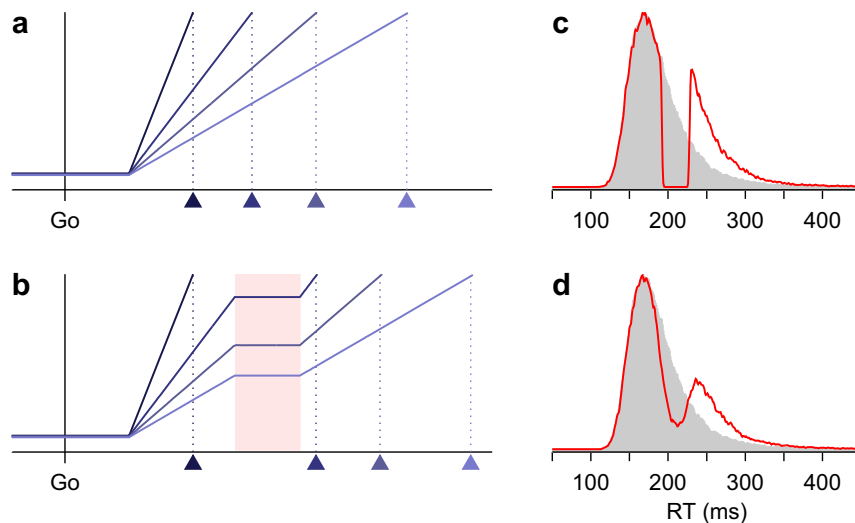
**Figure 1.** Interruptions in motor planning produce a dip in the distribution of saccadic RTs. (**a**) Four examples of the linear rise-to-threshold process. The y axis represents oculomotor activity as a function of time. Activity increases gradually, and when a critical threshold level is reached, a saccade is triggered. All four plans are identical except for their build-up rates. Triangles indicate saccade onset. (**b**) Four example motor plans that are briefly interrupted. The rise in activity halts during the interruption period (red shade). Build-up rates are the same as in the panel above. (**c**) Simulated RT distributions for motor plans that rise to threshold uninterrupted (gray shade) or that halt (red line) for 36 ms (between 192 and 228 ms after the go signal) but are otherwise identical. Distributions were obtained from 50,000 simulated trials using 1 ms RT bins and Gaussian smoothing with $\sigma = 1$ ms. Note the sharp discontinuity produced by the pause. (**d**) As in (**c**) except that the probability of interruption for any given trial was equal to 0.7 (rather than 1), and the onset and offset of the interruption interval varied normally with a standard deviation (SD) equal to 8 ms (instead of 0). Note the smooth dip in the distribution.

to bottom-up information and are closely related to the selective firing of visually-responsive neurons within oculomotor circuits[20,33–37].

Therefore, it must be typically the case that, when a salient stimulus appears abruptly, the corresponding visual information arrives at oculomotor planning areas while the activity associated with some movement vector is already rising toward threshold. What should the system do then? Should the ongoing plan be completed first, before the new information is acted upon, or should the plan be canceled and replaced with another one toward the recent stimulus? Clearly, the answer depends on the priorities of the targets. So, it may be best to put the ongoing saccade plan on hold, evaluate the new sensory information, and *then* determine whether to continue with the current plan or to cancel it in favor of an alternate one.

This intuition has been mentioned earlier[38–40]. For instance, Anderson and colleagues[40] articulated the problem this way: "involuntary attentional capture can rapidly orient the organism to unexpected changes that could signal danger or opportunity, but has the potential to cause distraction from intended acts of perception". That is, delaying a response to a visual stimulus that appears abruptly could be costly when it reveals an imminent threat or a fleeting opportunity, but always canceling the ongoing saccade plans would be quite inefficient, because not all salient stimuli have such importance. Here, we provide an initial framework for evaluating this intuitive tradeoff quantitatively. First we show that a simple descriptive model replicates a variety of experimental results which, in effect, indicate that ongoing saccade plans are transiently interrupted by salient, sudden visual stimuli. Then, based on this model, we use analytical and numerical calculations to determine the conditions under which such interruption is behaviorally advantageous, and estimate the actual temporal tradeoff that it entails.

## Results

**Behavioral manifestation of an interruption in motor planning.** First we make a simple observation, which is that, across multiple trials, a transient interruption in the rising activity that comprises the motor plan to make a saccade leaves a characteristic signature in the corresponding reaction time (**RT**) distribution – a dip.

The neural events that precede the onset of a saccade are well understood, particularly when the eye movement is in response to the presentation of a visual target. After the target is shown, activity in oculomotor areas, most prominently the FEF and superior colliculus (**SC**), starts building up, with the activated neurons encoding the movement vector required to fixate the target. The ramping activity keeps increasing until it reaches a point of no return – the threshold – and shortly thereafter (~10–20 ms) the saccadic movement is initiated[1–3]. Although the baseline and threshold levels of this rise-to-threshold process may vary and contribute substantially to its dynamics[6], in general, RTs are most strongly correlated with variations in the build-up rate[1,4–6,16,41]. Thus, a reasonable simplification is that short versus long RTs largely correspond to steeper versus shallower excursions in activity between baseline and threshold (Fig. 1a).

In agreement with this account, it is well established that the RT distributions of simple reactive saccades, which have a characteristic skew, are accurately replicated by a linear rise-to-threshold process in which the build-up rate, $r_{BU}$, is constant within each trial but is drawn from a Gaussian distribution from one trial to the next[42,43]. Under such conditions, the RT can be expressed as

$$RT = T_A + \frac{\Theta - B}{r_{BU}} \tag{1}$$

where $T_A$, $B$, and $\Theta$ are constants (the afferent delay, baseline activity, and threshold), and for each trial the build-up rate, $r_{BU}$, is a different sample from the same Gaussian distribution. Simulated RT distributions based on this expression (Fig. 1c,d, gray shades) closely approximate distributions obtained experimentally based on saccades to single targets. This is the mechanism that we use for relating oculomotor activity to saccadic RT (we comment on its relationship to other modeling frameworks at the end of this section).

With this rise-to-threshold mechanism in place, we can now introduce the main focus of our study, which is the characteristic effect produced when the build-up process is interrupted during a consistent time period. If the build up of activity is momentarily halted, such that the firing level is maintained constant during the pause (Fig. 1b), then the RT distribution is essentially split into two pieces separated by an empty interval (Fig. 1c, red trace). The location of the void within the distribution depends on the temporal offset between the go signal (or more precisely, the start of the rise to threshold) and the onset of the interruption, and the length of the void is equal to the duration of the pause. Whatever the offset, if the two parts of the distribution were brought together, the original distribution without interruption would be recovered. Thus, ideally, the behavioral manifestation of a consistent interruption of the saccadic motor plans is a discontinuous split of the RT distribution.

Notably, though, any variability in the interruption mechanism will turn the discontinuous split into a smooth dip (Fig. 1d). (The sharpness of the discontinuity will also depend on the bin width used to generate the RT histograms, but this effect should be minor.) In particular, there are three ways in which such smoothing would be likely to occur. (1) Rather than halting, the motor plan could keep increasing at a low build-up rate. A weak but nonzero build-up rate would partially fill in the void in the RT distribution, and if that low rate fluctuated randomly across trials, the edges of the void would be softened. (2) Rather than always halting, the motor plan could halt on some trials but not on others. That is, the probability of halting could be less than 1. (3) Rather than being constant, the duration of the interruption (and/or its onset and offset) could also fluctuate randomly across trials. As with point 1, the larger the fluctuations, the heavier the smoothing. These factors can easily turn a sharp, discontinuous split in the RT distribution into a visible but much more subtle, smooth dip (compare red traces in Fig. 1c,d), with the width of the dip depending on both the mean and the variability of the interruption (Methods).

The effect of the interruption can also be appreciated analytically, based on how much each individual RT is delayed. If on a given trial the interruption halts the motor plan starting at a time $I_{on}$ and lasts $q$ ms, then the observed RT is simply

$$RT_I = \begin{cases} RT_0 & \text{if } RT_0 < I_{on} \\ RT_0 + q & \text{if } RT_0 \geq I_{on} \end{cases} \tag{2}$$

where $RT_0$ is the value that would have been measured had the same saccade plan (with the same build-up rate) not been interrupted. In other words, as can be seen in Fig. 1b, if the interruption starts at $I_{on}$ (left edge of red shade), saccades that are programmed to reach threshold before $I_{on}$ are executed normally, whereas saccades that are programmed to reach threshold after $I_{on}$ are delayed by $q$ ms, which is the duration of the pause. This is true for any $RT_0$. So, assuming that the activity level does not change during the interruption and that $I_{on}$ and $q$ are constant, equation (2) leads to a sharp break in the RT distribution that starts at $I_{on}$ and is $q$ ms wide. However, when $I_{on}$ and $q$ are not constant, the fluctuations across trials simply smooth the edges of the otherwise discontinuous, empty interval. This result, i.e., equation (2), is important because it means that the sharp break (or smooth dip) in the RT distribution occurs independently of the distribution's original shape (without the interruption).

Note that equivalent effects would be obtained if the neural activity during the interruption decreased instead of remaining constant (Supplementary Fig. 1). This would result in an apparently longer interruption, because after the suppression interval the firing rate would require additional time to recover the ground lost in its excursion toward threshold. Although these alternative responses (constant versus decreasing) would be neurophysiologically distinct, they would be impossible to disambiguate based on RT measurements alone. For simplicity, then, in what follows we only consider interruptions with constant or weakly increasing activity, but the possibility of active suppression during the interruption is left open.

In the rest of the paper we review the conditions under which such interruption is likely to occur, show that it is consistent with a wide variety of experimental studies, and explore what functional role it may play within perceptuo-motor interactions. It will be useful to keep in mind that the models we consider have two elements, the rise-to-threshold process and the interruption itself, which are substantially independent. For the former, we use a linear rise, as in equation (1), because this representation captures an essential neurophysiological finding – that the variance in saccadic RT is largely due to the variance in build-up rate – accurately and with minimal complexity[1,5,6,16]. This mechanism is similar in many ways to other RT-generation models that make slightly different assumptions (e.g., refs[43,44]), and could be augmented to take into account, say, the strength of the sensory stimulus that the saccade is directed to (as in refs[45,46]), or specific aspects of the local circuit dynamics (as in ref.[47]). However, our focus here is not on the details of the saccadic choice process, but rather on the effects of the interruption, which should be qualitatively the same for any saccade generation model based on activity that rises to a bound (for evidence of this, see Supplementary Fig. 2).

**Saccadic inhibition as an interruption.** Many psychophysical studies have characterized the phenomenon best known as "saccadic inhibition" or the "remote distracter effect," which is typically produced by presenting a brief distracting stimulus while a participant is just about to make a saccade to a target. Over multiple trials, the influence of the distracter manifests as a dip in the distribution of saccadic RTs (for correct responses), very much as in Fig. 1d (red trace). In this section we describe experimental and accompanying simulation results that support our functional interpretation of the dip – that it is the signature of a stimulus-driven pause in the rise-to-threshold process that comprises a saccadic motor plan.

Before analyzing one specific experiment in detail, we wish to underscore the robustness of the phenomenon. In saccadic inhibition experiments, the characteristic dip is observed whether the distracter stimulus is discrete and localizable (e.g., refs[39,47]) or widely distributed and lacking a well-defined spatial location (e.g., wide bars flashing above and below the intended saccade target[38]). The strength of the effect depends on the saliency of the distracting event; for instance, it is stronger for larger[48–50] and higher contrast stimuli[47]. Manipulation of other stimulus dimensions, such as spatial frequency, produce only small variations (~10 ms) in the leading phase of the dip[51]. And notably, the effect is stronger and more prolonged when the stimulus appears in an attended or task-relevant location as opposed to an unattended, task-irrelevant one[48,52]. Thus, all manner of abrupt visual stimuli produce saccadic inhibition, but their saliency matters.

In addition, the phenomenon is surprisingly independent of volitional control and of the way in which saccades are triggered. For instance, similar effects on RT are produced whether participants are instructed to make a saccade to the target or an antisaccade away from it[38]. Consistent with this, robust saccadic inhibition (i.e., a clear dip) is observed in the context of more cognitively demanding behaviors as well, such as search tasks[51], double-step saccades[39], and reading[48]. At the other extreme, saccadic inhibition is equally robust for microsaccades[17] and for the reflexive saccades observed during the quick phase of nystagmus, both during its normal operation[53] (optokinetic nystagmus) and when it is pathological[54] (infantile nystagmus).

These findings indicate that, when a saccade is being prepared and a salient stimulus appears suddenly, the result is almost invariably a visible dip in the saccadic RT distribution – an effect that occurs regardless of how the saccade plans are initiated and that cannot be prevented. Our hypothesis is that the dips observed experimentally are the consequence of a consistent interruption in motor planning, as depicted in Fig. 1b. If this is correct, then the proposed interruption mechanism should be able to account quantitatively for such experimental data.

To investigate this, we model a saccadic inhibition experiment by Bompas and Sumner[47] that displays a rich phenomenology (see Methods). We consider data from one of their human participants, but the results were qualitatively similar for three others[47]. In the experiment (Fig. 2a), participants gazed at a central fixation point for 700 ms, after which a target appeared randomly to the left or right of fixation (±8°). Participants were instructed to saccade rapidly to the target, ignoring any other stimuli. In 83% of the trials, a distracter was briefly flashed at a location diametrically opposite to the target, with the temporal interval between target and distracter (the stimulus onset asynchrony, or **SOA**) varying randomly across trials (SOA = 0, 20, 40, 60, or 80 ms). The RT was measured as the interval from target onset to saccade onset.

Even though the distracter appeared in most trials, it had a marked impact on performance. In trials with no distracter, the distribution of saccadic RTs was unimodal, with the usual long tail on the right side (Fig. 2b). In trials in which a distracter was shown, the resulting RT distributions demonstrated a clear dip, with the minimum consistently occurring about 100 ms after distracter onset. That is, as the SOA increased, the dip in the RT distribution shifted further to the right, consistent with an interruption in motor planning that is time-locked to the onset of the distracter (Fig. 2c–g, black traces). Interestingly, for an SOA of 0 ms, the effect looks less like a split and more like a pure rightward shift of the whole distribution – which is exactly what one would expect from an interruption that occurs right at the onset of the rise-to-threshold process. The resulting progression goes from the original, unimodal RT distribution without any distracters, to the shifted distribution at the shortest SOA, to distributions with a pronounced dip that gradually shifts to the right.

This progression was accurately reproduced by a simple rise-to-threshold model similar to that of the idealized case discussed in the previous section (Fig. 1). As before, the parameters of the interruption (probability of occurrence, mean and variability of onset and offset) were fixed, except that, in this case, the onset of the interruption was time-locked to the onset of the distracter (on average, the interruption started 85 ms after distracter onset; Methods). In this way, the simulated interruption produced a dip in the RT distribution that shifted as a function of the SOA (Fig. 2b–g, red traces) – precisely as observed in the experimental data (black traces).

Note the parsimony of the model. Both the build-up process and the interruption used to generate the synthetic data were statistically the same for all the SOAs. Furthermore, although the simulation was based on a rise-to-threshold mechanism, it made no assumptions about the underlying distribution of build-up rates; instead, the build-up rates were effectively inferred from the experimental data (specifically, from the distribution of RTs measured[47] in the no-distracter condition; Fig. 2b; Methods). The results show that an interruption in motor planning that is triggered by the appearance of the distracter can account for the variations in the RT distribution observed experimentally across SOAs.

Having set the interruption parameters to fit the RT distributions (Fig. 2c–g), the simulated data were then aligned to the onset of the distracter and combined across SOAs to generate a single response histogram. For each simulated trial, we calculated the interval between distracter onset and saccade onset, which we refer to as the processing time, or **PT** (Fig. 2a). This is the maximum amount of time that the system has to process the distracter signal in each trial. The resulting PT distribution (Fig. 2h, red trace) revealed a more prominent dip, and the dip observed in the experimental distribution closely matched the expectation (Fig. 2h, black trace). The fact that the experimental PT histogram, which combines data across all SOAs, indeed reveals a pronounced dip is further evidence that the interruption is synchronized to the presentation of the distracter.

Curves like those in Fig. 2h were similar for the other participants in the study by Bompas and Sumner[47] (see their Fig. 4 and our Supplementary Fig. 3), and have been reported in many studies[38,39,48–54]. The point here is
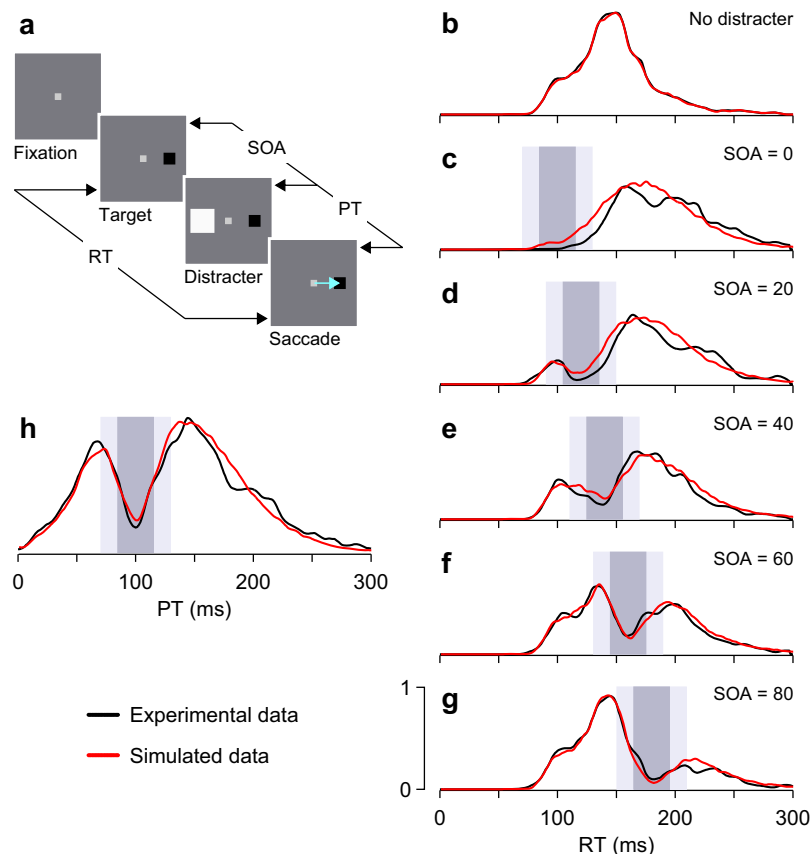
**Figure 2.** Interruptions in motor planning account quantitatively for observed dips in RT distributions. (**a**) The paradigm used by Bompas and Sumner[47] to study saccadic inhibition. Participants were instructed to make a saccade (blue arrow) to the target (black square) as soon as it appeared, and to ignore any distracters. In most trials, a distracter (white square) was briefly flashed (50 ms) opposite to the target. The delay between target onset and distracter onset (SOA) varied randomly across trials. For each trial, the processing time (PT) is equal to RT − SOA, and corresponds to the time interval between distracter onset and saccade onset. (**b**–**g**) RT distributions reported by Bompas and Sumner[47] (observer 1, black traces, 1200 trials per trace) along with our model results (red traces). SOAs are indicated, in ms. In the rise-to-threshold model, the mean interruption occurred from 85 to 115 ms after distracter onset (dark gray shades), with the onset and offset times varying across trials (light gray shades show 1 SD in each direction). (**h**) PT distributions for the simulated (red trace) and experimental data from Bompas and Sumner[47] (observer 1, black traces). PT histograms include aggregated data from all SOAs. All experimental data were redrawn from Bompas and Sumner[47] (Methods).

simply that the transient interruption in motor planning provides a compact and accurate phenomenological description of the empirical data – with minimal assumptions about the saccade generation mechanism

Direct neurophysiological evidence is also available indicating that a sudden-onset distracter briefly inhibits the progress of an ongoing saccade plan. In the SC, presenting an irrelevant distracter while a monkey is preparing to make an eye movement to a known location produces a brief decrease in the activity of visuomotor neurons precisely at the time expected based on the behavioral studies, about 90 ms after distracter onset[55,56].

In conclusion, plentiful evidence indicates that the oculomotor system responds reflexively and rapidly to sudden changes in visual input in a manner that is consistent with ongoing saccade plans making a brief pause in their rise toward threshold.

**Temporal advantage of saccadic inhibition.** In this section we explore the conditions under which it is advantageous to halt an ongoing saccade plan, and estimate the actual temporal benefit of doing so. We do this by means of theoretical (and complementary numerical) calculations that compare different strategies for updating and prioritizing saccade plans toward two targets. Specifically, suppose that a stimulus appears at location $B$ while a saccade plan toward location $A$ is already underway. Then, the key question is this: if the stimulus (at $B$) is indeed relevant, i.e., worth examining right away, how quickly can a saccade to it be generated?

First, note that if the arriving visual signal is spatially congruent with the currently planned saccade endpoint (i.e, $B \approx A$), then there is no conflict and the ongoing plan should simply continue. In that case canceling or halting the current plan would be unnecessary and waste time. And indeed, experimental evidence from saccadic inhibition experiments[49,57,58] is consistent with this intuition (see Discussion). Therefore, we only consider the
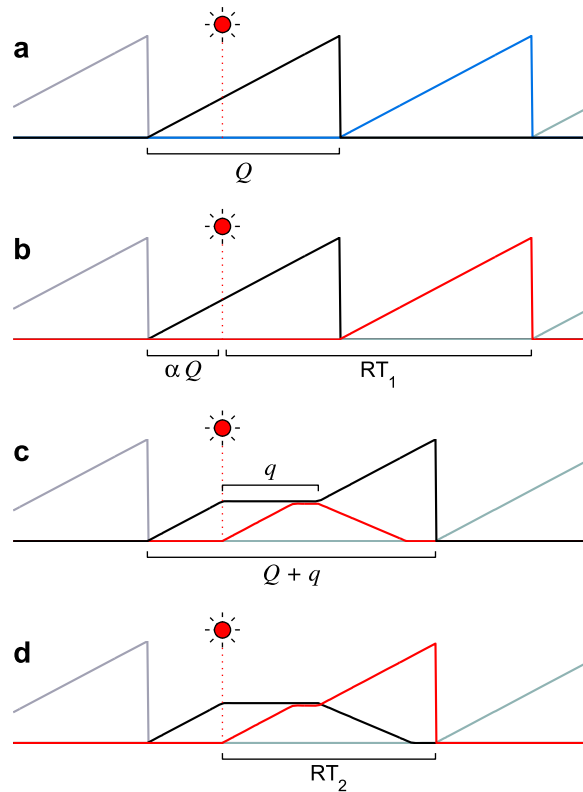
**Figure 3.** Two scenarios whereby saccade plans may incorporate new visual information. Traces represent the activity of saccade-related neuronal populations. Each color represents activity associated with a different saccade direction. Activity rises progressively, and when it reaches threshold, a saccade is triggered and a new motor plan is initiated. Here, all plans have the same build-up rate. The circuit detects a new stimulus (red flash) at a time $\alpha Q$ relative to the onset of the ongoing plan (black trace). (**a,b**) A serial programming scenario. Saccade plans are produced every $Q$ ms and cannot be interrupted. After stimulus detection, the next saccade may be to an unrelated location ((**a**) blue trace) or toward the flash ((**b**) red trace), in which case its latency is $RT_1$ ms from the time of detection. (**c,d**) A concurrent-programming scenario. After stimulus detection, the ongoing plan (black trace) is halted for $q$ ms and a second plan, toward the stimulus (red trace), starts rising. After the interruption, either the second plan is canceled (**c**), or the first one is (**d**), in which case a saccade to the flash is made after $RT_2$ ms from the time of detection.

situation in which the new visual signal and the ongoing plan are spatially incongruent, in which case the best course of action will depend on the priorities of the two potential visual targets.

Assuming that saccades are generated by a linear rise-to-threshold mechanism, we compare two simplified saccade programming scenarios: one (serial) in which, once started, the build-up process cannot be interrupted, and another (concurrent) in which it can (Fig. 3). In both scenarios, movement-related activity starts at zero and rises toward a threshold level $\Theta$ with a build-up rate $r_{BU}$, so that the time between saccades is normally $Q = \Theta/r_{BU}$. For simplicity, we first consider $r_{BU}$ to be constant; later we will relax this constraint (and others) and show that additional, intrinsic randomness in the build-up rate does not affect the argument or the conclusions.

In the serial scenario (Fig. 3a,b), saccades are generated every $Q$ ms, one after another. Critically, when a new visual stimulus is detected at location $B$, the ongoing plan toward location $A$ (black trace) keeps advancing. After the detection, while the activity builds up, the system evaluates the priority of the stimulus. Then, once the saccade to $A$ is made, a new saccade plan is initiated, with its direction reflecting the outcome of the deliberation: if the stimulus was considered of low priority, the new plan is toward an unrelated target (Fig. 3a, blue trace), whereas if the stimulus was considered important, the new plan is toward it (Fig. 3b, red trace). In the latter case, assuming that the visual signal is detected by the circuit at a time $\alpha Q$ from the start of the original plan, the latency for making a saccade to the stimulus at location $B$ is

$$RT_1 = Q(2 - \alpha) \tag{3}$$

counting from the moment of detection. Note that $\alpha$ varies between 0 and 1, and indicates when the stimulus is detected by the circuit relative to how advanced the ongoing motor plan is.

In the concurrent scenario (Fig. 3c,d), the movement-related activity also rises with a build-up rate $r_{BU}$, so in the absence of new visual information saccades are still produced every $Q$ ms. However, when a new visual stimulus is detected at location $B$, again at time $\alpha Q$, two things happen: first, the ongoing plan toward $A$ is halted for $q$ ms, and second, a new, competing plan toward the stimulus starts rising immediately, with the same build-up
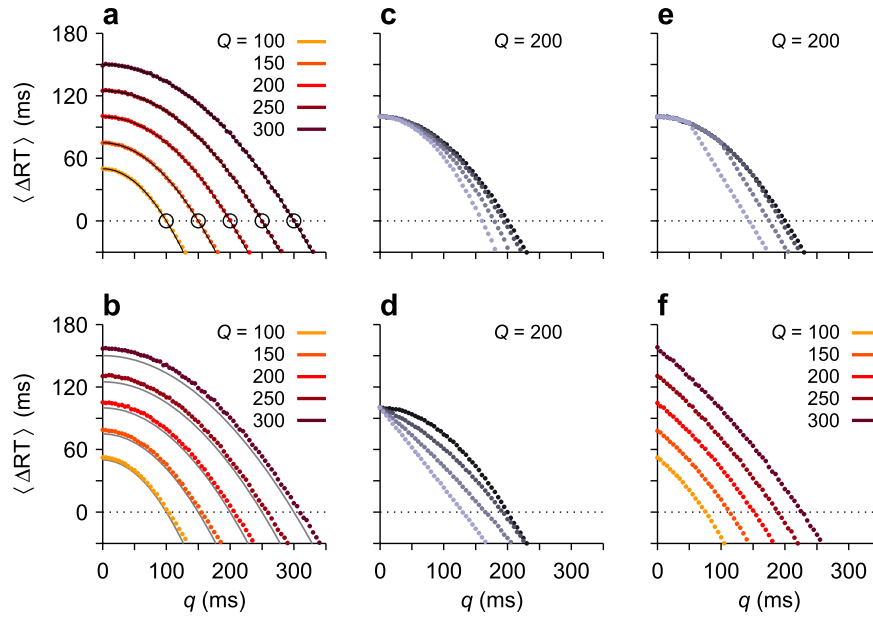
**Figure 4.** Temporal advantage of an interruption in motor planning. Each panel shows the mean difference in RT between a serial strategy, in which plans are not interrupted, and a parallel strategy in which saccade plans are interrupted and deployed concurrently. The x axes show the deliberation time, $q$. The mean saccade latency, $Q$, is indicated. (**a**) Expected difference in RT from theoretical calculations. Colored dots are results from computer simulations; gray lines are from equation (6). Circles mark the points beyond which concurrent programming ceases to be advantageous. (**b**) As in (**a**) but with Gaussian variability in the simulated build-up rates. The SD of $Q$ was approximately equal to $Q/4$. (**c**) Results when the ongoing motor plan is not fully arrested but keeps rising at a lower build-up rate. Going from dark to light points, the build-up rate, $r_{BU}$, of the ongoing plan was equal to 0, 0.25, 0.5 and 0.75 times the initial value. (**d**) Results when the concurrent motor plan toward the stimulus rises at a lower build-up rate. Going from dark to light points, the build-up rate, $r_{BU}$, of the second plan was equal to 1, 0.75, 0.5 and 0.25 times the maximum value. (**e**) Results when the concurrent motor plan toward the stimulus can rise only up to a point below threshold. Going from dark to light points, the maximum rise is equal to 1, 0.75, 0.5 and 0.25 times the maximum value, $\Theta$. (**f**) As in (**b**) but with all the effects in panels (**c**–**e**) with coefficients of 0.5 applied simultaneously.

rate, $r_{BU}$. During the $q$ ms that the pause lasts, this second, parallel plan can only rise up to the level of the first plan, not further. Thus, the second plan can either catch up with the first one or stay below it – but cannot overtake it. In this scenario the deliberation occurs during the interruption, so it lasts $q$ ms. Finally, after the interruption, again two outcomes are possible: either the initial plan toward $A$ continues and the budding plan toward $B$ is canceled (Fig. 3c), or the reverse, the initial plan is canceled and the later one, toward the stimulus at location $B$, keeps advancing (Fig. 3d). In the former case, when the stimulus is deemed of low priority, the eyes still land at the original target location, $A$, but the RT associated with that saccade expands from $Q$ to $Q + q$ ms. In the latter case, when the stimulus turns out to be of high priority, the saccade is toward location $B$ and its latency is

$$\text{RT}_2 = Q + q - \min\{\alpha Q, q\} \tag{4}$$

again relative to the moment of detection, where $\min\{a, b\}$ is the minimum of $a$ and $b$. The question now is, which scenario is more efficient?

For a saccade to the stimulus, the difference in latency between the two scenarios is

$$\Delta\text{RT} = \text{RT}_1 - \text{RT}_2 = \begin{cases} Q - q & \text{if } \alpha < q/Q \\ Q - \alpha Q & \text{if } \alpha \geq q/Q \end{cases} \tag{5}$$

where a positive value indicates a faster response in the concurrent case. This expression shows that the advantage in RT depends on when the stimulus is detected relative to how advanced the ongoing plan is, which is what $\alpha$ represents. If the first plan is just starting to rise ($\alpha \approx 0$), the difference is potentially large; whereas if the planned saccade toward $A$ is just about to be executed ($\alpha \approx 1$), the difference is minimal. The mean difference, averaging over detection times (i.e., integrating over $\alpha$, assuming it is uniformly distributed), is

$$\langle\Delta\text{RT}\rangle = \frac{Q}{2}\left[1 - \left(\frac{q}{Q}\right)^2\right] \tag{6}$$

This is the main result. According to this expression, re-prioritizing on the fly pays off when the perceptual evaluation time (equal to the interruption duration), $q$, is shorter than the intersaccadic interval, $Q$. It is in that case that a positive average difference in RT is obtained. Coupled to concurrent planning, the motor interruption is advantageous – but only if perception is sufficiently fast.

Note that the above result is conservative, in that it (slightly) favors the purely serial strategy. This is because, in the serial scenario, no additional deliberation time is ever required, even in the extreme case when the stimulus is detected just before the saccade to $A$ is triggered. The perceptual evaluation of the stimulus is always completed while the first motor plan is ongoing, however short that interval may be. By contrast, in the concurrent scenario the interruption time is equal to the deliberation time, i.e., $q$ is precisely the amount of time needed to resolve whether the stimulus at $B$ is of low or high priority.

What is the actual advantage of concurrent planning, in milliseconds? According to the theory, this depends on the duration of the interruption relative to the typical RT (Fig. 4a), but a rough estimate can be made based on physiologically plausible values for $Q$ and $q$. First consider an intersaccadic interval of 200–250 ms, which is most common in humans and monkeys[7–10], and an interruption time of approximately 70 ms, as estimated from saccadic inhibition experiments[38,47,48,50]. Based on equation (6), those numbers give a mean difference $\langle\Delta RT\rangle$ between 88 and 115 ms favoring the parallel programming strategy, which is quite large as a proportion of the median intersaccadic interval (~45%).

There is a cost associated with the interruption of the ongoing plan, of course, which is that when the stimulus turns out to be irrelevant, the original RT is lengthened by $q$ (70) ms (Fig. 3c). Also, the previous calculation might overestimate the true advantage, because saccades *can* be produced considerably more quickly (see Discussion), and because the interruption could conceivably be longer, depending on the location and saliency of the stimulus, as well as on the robustness of the motor plans. Assuming that conditions are optimized for producing short saccadic latencies ($Q \approx 150$ ms), as happens when target locations are predictable and subjects are highly motivated[6,28,35,59], and that the interruption is as long ($q = 115$ ms) as it could be under extreme circumstances[6,39], the average time saved according to equation (6) is now 31 ms. This is considerably smaller than above – but still sizable as a proportion of the time between fixations (~21%).

More generally, the dependence of $\langle\Delta RT\rangle$ on $q$ and $Q$ indicates that interrupting the ongoing saccade plan and initiating a parallel plan while deliberating on the priority of the new information is clearly advantageous over a wide range of physiologically relevant values (Fig. 4a) – again, as long as the perceptual processes underlying the deliberation (e.g., stimulus detection and identification) are at least as fast as the mean intersaccadic interval. Indeed, perceptual processing speeds typically satisfy this requirement, as elaborated in the Discussion.

One could worry, though, that the analytical result involves strong simplifications, so we made similar comparisons based on computer simulations in which various assumptionss were relaxed. As a check, we first simulated the same conditions assumed by the theory, and indeed, the analytical and numerical results were in agreement (Fig. 4a, lines vs. dots). Then we allowed the build-up rate to be different for each saccade plan ($r_{BU}$ was drawn from a Gaussian distribution). Introducing a large amount of variability in the build-up rates had a very small effect that tended to increase the advantage of the concurrent programming strategy (Fig. 4b, compare dots vs. gray lines; see caption). Next, we relaxed three key assumptions of the concurrent scenario in ways that tended to lessen its advantage. (1) The interruption did not fully arrest the ongoing plan, but simply diminished its build-up rate, and if the ongoing plan reached threshold before the deliberation was over, the plan toward the stimulus had to restart from the baseline. (2) The plan toward the stimulus could rise during the interruption interval but slowly, at a fraction of its nominal build-up rate. (3) The plan toward the stimulus could only rise so far during the interruption; that is, as before, the second plan could not overtake the first one, but in addition, it was not allowed to increase beyond an absolute level below the saccade threshold. All of these manipulations curtailed the amount of progress that the concurrent plan toward the stimulus could make during the pause, but the effects were gradual and required large changes in the parameters, on the order of 50%, to be substantial (Fig. 4c–f). The results demonstrate a frank advantage of the interruption mechanism over strict serial programming over a wide range of physiological conditions.

These results show that psychophysical measurements of saccadic latencies and interruption times are largely consistent with the proposed theoretical mechanism whereby ongoing saccade plans are put on hold and re-prioritized on the fly, enabling significantly faster eye movements toward recently detected stimuli, with the benefit being on the order of several tens of milliseconds per saccade.

**Pronounced interruptions in the double-step task.**    Our functional interpretation of the interruption is that it exists because, when a salient visual stimulus is detected, the oculomotor system pauses to deliberate on its priority. If this is correct, then the interruption should depend only on the stimulus itself, regardless of the outcome of the perceptual evaluation process. This leads to a concrete prediction: given a novel stimulus, the interruption should be the same whether such stimulus is deemed relevant (i.e., a target) or irrelevant (i.e., a distracter).

In this section we consider an elegant experiment[39] for which this prediction is shown to be correct. This example is also important because it specifically illustrates how the interruption and concurrent motor programming can jointly determine a subject's performance in a well-known oculomotor paradigm, the double-step task.

We discuss the version of the double-step task implemented by Buonocore and colleagues[39] (Fig. 5a; for further task details, see Methods). All trials start in the same way, with the participant fixating on a central point flanked by two empty circles, or placeholders, one on the left and another on the right. Target presentation corresponds to one of the circles being filled. In one half of the trials (control trials), a target appears and stays on for the rest of the trial (700 ms). In the other half of the trials (actual double-step trials; Fig. 5a), a target appears but then steps to the opposite location after a short delay (SOA). The participant is instructed to look at the target,
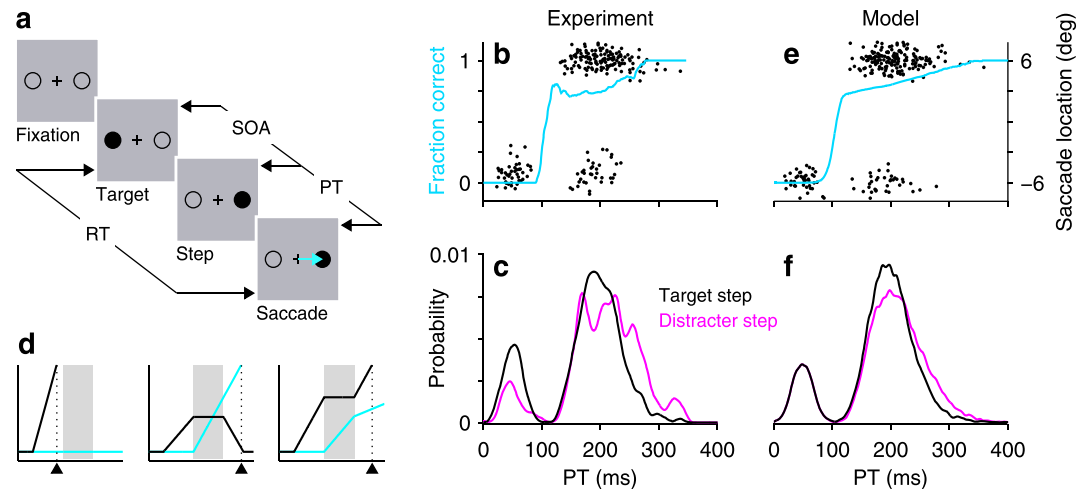
**Figure 5.** Interruptions in motor planning due to targets and distracters. (**a**) The double-step paradigm used by Buonocore *et al.*[39]. Participants were instructed to make a saccade (blue arrow) to the target (filled circle) as soon as it appeared. In one half of the trials the target did not change (not shown); in the other half it stepped to the opposite location after a delay (SOA). The PT in each trial is equal to RT − SOA. (**b**) Performance of one participant in the double-step task. Saccade landing points are arranged by PT (black dots, 240 trials, right y axis; data redrawn from ref.[39], experiment 1, participant 3). Targets were located at ±6°. The direction transition function (blue trace, left y axis) is a running histogram (bin size = 81 ms) of the proportion of correct saccades toward the stepped target as a function of PT. We computed it based on the shown saccade landing points. (**c**) PT distributions in double-step (black trace) and distracter-step (pink trace) blocks performed by the same participant (data redrawn from ref.[39], experiment 1, participant 3). In distracter-step blocks, participants were instructed to always make a saccade to the first target, ignoring the step. (**d**) Motor plans toward initial (black traces) and stepped (blue traces) target locations in 3 simulated double-step trials. A fast error (left), a correct saccade (middle), and a slow error or lapse (right) are shown. Shades indicate mean interruption interval. Triangles mark saccade onset. (**e**) Simulated double-step responses (black dots; similar number of trials as in (**b**)) and direction transition function (blue trace; based on 20,000 trials). Correct and incorrect simulated saccades were assigned to ±6° landing points with additional random scatter. (**f**) PT distributions in double-step (black trace) and distracter-step (pink trace) simulated trials. For all data the SOA was constant at 120 ms.

and to move his or her eyes to the second target if it appears. Importantly, control and double-step trials are randomly interleaved, and so are left and right target locations. Thus, in step trials, depending on the SOA and the participant's readiness to respond, the resulting eye movement may be toward the first (incorrect) or the second (correct) target location.

Performance in the task is quantified via the direction transition function (Fig. 5b, blue trace). This curve describes the probability of making a correct saccade to the stepped target as a function of PT (Fig. 5a), which in this case is the interval between step onset and saccade onset (i.e., the amount of time available in any particular trial to view the target after it has stepped, where PT = RT − SOA). As can be seen for one of the six participants that performed this experiment[39] (Fig. 5b, blue trace), when the step is viewed for less than 100 ms or so, the resulting saccade is to the first, incorrect location, but as the PT increases beyond that mark, the likelihood of making a correct saccade to the stepped target rises steeply and then levels off. The transition from (fast) incorrect responses to correct ones occurs extremely rapidly, within 20 ms of PT or so.

Crucially, Buonocore and colleagues[39] showed that the timing of the saccades in the double-step task is remarkably consistent with the target step producing robust saccadic inhibition of an initial motor plan toward the first target location. To see this, first consider the distribution of PTs (Fig. 5c, black trace), which includes both correct and incorrect responses (Fig. 5b, black points). The distribution is clearly bimodal, with a pronounced dip centered around PT ≈ 105 ms. This is, indeed, as if the target step had inhibited an ongoing saccade plan. To verify this interpretation quantitatively, Buonocore and colleagues[39] did the following manipulation. In separate blocks of trials, they presented the same displays as in the double-step task but instructed the participants to ignore the step and always make an eye movement to the initial target location (same sequence as in Fig. 5a, except with the correct saccade in the opposite direction). But note that this works as a classic saccadic inhibition paradigm, in which a salient distracter (the stepped target) must be ignored and, in addition, the actual target disappears. What they found was that the PT distribution from the distracter-step blocks (Fig. 5c, pink trace) was nearly identical to that in the double-step blocks; the saccades in both tasks produced the same unmistakable dip. Notably, the results were extremely similar across the six participants (their Fig. 2). These findings indicate that the interruption occurs identically in both cases and that it depends fundamentally on the stimulation conditions, regardless of whether the novel stimulus is considered a (high priority) target or a (low priority) distracter.

We performed computer simulations of these experiments (Methods) to verify this conclusion, i.e., that the initial motor plan is interrupted in the same way in both conditions, except that in one case (distracter-step trials) that first plan is meant to continue to threshold, whereas in the other (target-step trials) it is meant to be canceled

and replaced by an alternate one. Motor plans toward the initial target location (Fig. 5d, black traces) were generated as in Fig. 1, by drawing from a Gaussian distribution a different build-up rate for each trial. These initial plans were always interrupted, unless they reached threshold before the onset of the pause, of course (Fig. 5d, left). The onset and offset of the interruption varied across trials, to produce smooth RT distributions. In the distracter-step blocks, after the interruption the initial motor plan simply continued to threshold (the concurrent plan played no role in this case, because it was never meant to overtake the first one and, indeed, participants made few errors). In contrast, in the target-step blocks, a concurrent plan was launched at the beginning of the pause and the initial plan was canceled at the end of the pause (Fig. 5d, middle; this was true for correct responses; see below). In this way, the simulations reproduced not only the direction transition function (Fig. 5e, blue trace), but also the PT distributions for the two conditions (Fig. 5f). These had slightly different tails on the right side – a difference that was, in fact, consistent with the data of most participants in the study by Buonocore and colleagues[39] (as with Fig. 5c). Most importantly, the two distributions displayed identical dips.

In the double-step task, the example participant made some errors at long PTs (Fig. 5b). These slower errors can be considered lapses, incorrect responses that cannot be attributed to insufficient viewing time. In the simulations, directional errors corresponding to lapses were generated by assuming that the cancelation of the initial plan is not 100% certain (Methods). That is, most often the initial plan is successfully interrupted and canceled, and a correct saccade to the stepped location is triggered (Fig. 5d, middle); however, sometimes that initial plan is interrupted but not canceled, resulting in an incorrect saccade (Fig. 5d, right). This mechanism produced lapses with the appropriate timing (Fig. 5e, lower group of dots with PT > 100 ms). In this way, a simple model that allows the interruption of an ongoing saccade plan and concurrent programming of an alternate plan accounted for all the main features of this oculomotor experiment.

## Discussion

Because saccadic eye movements are planned continuously, the question arises of what to do when new visual information is detected. As a first approximation, there are three ways to proceed: (1) always let the ongoing plan continue, and deal with the new stimulus later, (2) always cancel the ongoing plan in favor of a new one congruent with the new stimulus, or (3) pause the ongoing plan, initiate a second, concurrent plan toward the stimulus, and deliberate to determine which one has a higher priority and should be executed next. Our review of the literature indicates that the latter, more flexible strategy is ubiquitous and robust, and our modeling results reveal two important aspects of it. First, the experimental manifestations of such strategy are accurately reproduced by a simple descriptive model in which the rise in oculomotor activity associated with a saccade plan is temporarily halted by a sudden-onset stimulus. This same mechanism applies to diverse experimental conditions, with only quantitative variations. And second, the combined pause/concurrent-planning strategy is not necessarily optimal just because it is more flexible. For it to be advantageous, perception must be fast. Specifically, the advantage depends on how fast the new information can be interpreted relative to how fast the saccades can be programmed (equation (6)). When the deliberation is quick, on average it pays off to briefly pause the ongoing plan and initiate an alternative one right away *every time*. In this way, although ongoing plans to targets detected earlier are delayed, responses to novel, high-priority ones are rushed.

The results suggest a compromise whereby the RTs to unexpected, high priority targets are shortened by a few tens of milliseconds per saccade while the responses to previously identified targets are delayed by a comparable amount; that delay is the price to be paid for being ready to respond when novel information mandates an immediate action. The price may seem steep, given that many saccades are postponed without any apparent necessity, but even a modest temporal benefit could be critical sometimes, e.g., when a cat detects a mouse springing out of its hiding place. Indeed, a remarkable aspect of many saccadic inhibition experiments, including those we discussed in detail[39,47], is that the interruption cannot be overcome at all, even when the distracter is predictable. Such a seemingly hardwired, reflexive effect, suggests strong evolutionary pressure at work. In a competitive world the currency of survival is time itself, and mechanisms that enable slightly faster reactions – when they matter most – may confer a vital advantage.

**Deliberation and perceptual processing speed.** Our calculations highlight the importance of sensory evaluation for saccade planning, and indicate that their speeds must match, i.e., the perceptual deliberation required to determine the priority of a novel stimulus must, on average, be shorter than the typical time between fixations ($q < Q$, equation (6)), which is about 200 ms in monkeys and humans[7–9]. Although many laboratory tasks create conditions in which sensory stimuli are judged for several hundreds of milliseconds or more (e.g., refs[46,60,61]), the perceptual decisions that characterize sensing in natural environments are likely to be much faster[62,63]. Specifically, the amount of time needed to make an accurate color discrimination is about 25–50 ms under urgent conditions, when oculomotor choices occur within a typical 200–250 ms RT window[5,16,20]. This number corresponds to the minimum amount of time that it takes for performance to transition from chance to 75% correct or above, so it is approximately equal to the deliberation time considered here, which excludes any transmission delays. Deliberation/processing times very likely increase when stimulus discriminability goes down[64–66]. But, on the other hand, the deliberation time is likely to be even shorter, closer to 20–30 ms, for a stop signal[67], a single, salient stimulus that appears abruptly and is interpreted as a command to stop an action (e.g., a red traffic light). Thus, results based on urgent-choice paradigms reveal timescales for relatively easy, fast perceptual judgements of a few tens of milliseconds, consistent with the interruption durations inferred from saccadic inhibition experiments.

This timescale is also consistent with classic visual search experiments used to investigate attentional allocation[68,69]. In visual search tasks, the main measurement is the time needed to determine whether a target is present or absent in a display, and the key quantity characterizing the overall difficulty of the task is the search slope, i.e., the slope of the linear fit describing the RT as a function of the number of items in the display. For a given search

task, the search slope corresponds directly to the deliberation time needed to identify each additional item as either target or distracter. Notably, search slopes are typically smaller than 10 ms/item when targets are highly salient or highly discriminable from non-targets, but even for difficult versions of these tasks, such as searching for a blue *H* among green *H*s and blue *A*s, deliberation times exceed 50 ms/item only rarely[68,69]. These numbers indicate that perception is normally quite fast – fast enough to satisfy the key constraint identified here, which makes the interruption/concurrent-programming strategy advantageous.

Still, given this constraint ($q < Q$), one might worry about the fact that intersaccadic intervals occasionally happen to be very short[8,70] ($<100$ ms), such as when a "corrective" movement is produced immediately after an incorrect saccadic choice[18,19,25,71,72]. In all likelihood, however, such short-latency responses are fast precisely as the result of concurrent motor programming. For instance, the latency of the correctives decreases as a function of PT in the double-step task[18,73,74]. This, together with evidence from single-neuron recordings in awake monkeys[19,74], indicates that, during errors, the motor plan favoring the correct target keeps advancing up until the onset of the erroneous saccade (as depicted in Fig. 5d, right), and likely contributes to the corrective movement that follows shortly thereafter. Thus, although extremely short intersaccadic intervals do occur, rather than negating the functional advantage of the interruption mechanism, they likely are a consequence of the concurrent motor programming that the interruption is proposed to enable in the first place.

**Temporal and spatial specificity.** According to the proposed functional interpretation, the interruption should vary according to the spatial congruence between the ongoing motor plan and the abrupt-onset stimulus. When the novel stimulus appears close to or at the endpoint of the saccade that is currently being planned, there should be no interruption, as there is no competing motor plan. And indeed, the effect on RT has the expected spatial dependence. When a stimulus is precisely coincident with a saccade target, the eye movement is triggered sooner[57,58,70], so the stimulus actually reinforces the ongoing motor plan. The decrease in RT goes away within a few degrees, as the location of the abrupt-onset stimulus deviates from the endpoint of the planned saccade vector, until the dip in the RT distribution emerges, and as the separation between the corresponding vectors approaches 180°, it becomes most pronounced[49,58] (see also ref.[17]). The interruption is manifest to the degree to which the ongoing and potential motor plans are in conflict with each other.

Congruence in the temporal domain is also important. Motor plans are reliably interrupted when a distant stimulus appears abruptly just before saccade onset, i.e., within the RT interval of the ongoing eye movement plan. When the time interval between the distracter and saccade onset is much longer ($\geq 300$ ms), as typically happens when a saccade target is presented 100 ms or more after the distracter, then the effect may go away or even reverse[58,75].

According to our framework, a saccade plan should be interrupted when a novel stimulus is detected *and* it is in conflict, spatially and temporally, with the target of the ongoing plan. In the absence of conflict no deliberation is necessary. These experimental observations are highly consistent with such constraint.

**Ubiquity of interruptions.** Besides providing a functional interpretation for the phenomenology associated with saccadic inhibition and attentional capture, the results are also relevant for any oculomotor task in which visual stimuli appear abruptly, even if the task design is quite different from that of saccadic inhibition experiments. Such stimuli are expected to interrupt motor planning for several tens of milliseconds, which may introduce sizable behavioral and neurophysiological effects, as in the case of the double-step task (see also ref.[6]). This is particularly true for urgent-decision tasks, in which motor plans are intentionally initiated ahead of the sensory information required to determine the correct choice. The interruption must be taken into consideration when modeling such tasks – if the data are to be replicated in quantitative detail[5,16,76,77]. Based on that work, we predict that the interruption plays a major role in an urgent version of the antisaccade task, in which a single stimulus is revealed abruptly and the curve that characterizes performance as a function of PT exhibits a huge dip[78].

Beyond saccades, unexpected events of all sorts lead to the suppression of ongoing motor activity in general, for all effectors[79,80]. Thus, "pausing to deliberate" may be a much more widespread principle of perceptuo-motor coordination than previously appreciated.

**Change of mind or change of plan?** The term "change of mind" has become popular for describing situations in which the neural activity predicting choice *A* is initially strong, but shortly thereafter the neural activity favoring option *B*, which was initially weak, gains momentum and prevails, so the choice is toward *B* ultimately[60,81,82]. These swings within single choice trials are sometimes explicitly described as "vacillations," or signatures of indecision[81]. However, the current and previous results[5,6,16,76,77] indicate that, rather than rare anomalies, these events must be extremely common, occurring as a natural and functionally useful aspect of the dynamics of motor circuits. The work of Thura and Cisek[83] shows that, in premotor and motor cortex, motor plans favoring one or another arm movement are exquisitely sensitive to individual, quantized pieces of relevant sensory information, such that the prevalence of one plan over another can swing back and forth quite distinctly as pertinent evidence is presented over time. This is simply in keeping with the flexibility inherent to the definition of a "motor plan," which implies that plans may be promoted, canceled, or otherwise changed as necessary up until the point of commitment[84]. That oculomotor circuits demonstrate a similar capacity to switch motor plans is perhaps less evident, insofar as saccades involve shorter timescales, fewer degrees of freedom, and stronger competition between them, making changes in saccade plans somewhat difficult to resolve (but see refs[5,6,16,85]). Conceptually, however, they should not be surprising, as they simply reflect flexibility that is behaviorally advantageous. The transient interruption of ongoing motor plans is one specific gear in the oculomotor machinery that enables such flexibility.

## Methods

All simulations were performed using Matlab. The Matlab code used to analyze and/or generate the data in the current study is available from the corresponding author on reasonable request.

**Inclusion of experimental data.** We include experimental datasets from two studies, one by Bompas and Sumner[47] (Fig. 2) and another by Buonocore and colleagues[39] (Fig. 5). In both cases, the data were extracted from the published digital figures by converting them to jpeg format and using the Matlab function grabit. Once captured, the numerical values were used for further analysis and plotting.

**The rise-to-threshold process.** RTs were generated by simulating a linear rise-to-threshold process over multiple trials. The following procedure was the same for all simulations except those in Fig. 2, as explained below.

The firing rate variable $R$, which represents oculomotor activity, was integrated numerically by applying the update rule

$$R(t + \Delta t) = R(t) + v_R \Delta t \tag{7}$$

with a time step $\Delta t = 1\,\text{ms}$ and a rate of change $v_R$. The go signal was assumed to occur at $t = 0$, at which point $R$ was equal to a baseline value $B = 0$. The saccade was assumed to be triggered when $R$ reached a threshold $\Theta = 1000$ arbitrary units (**AU**), at which point the RT was computed (the efferent delay was ignored, as it simply contributes a constant that can be consolidated with the afferent delay). Over the course of the trial, the rate of change took the following values

$$v_R = \begin{cases} 0 & \text{if } t < T_A \\ r_I & \text{if } I_{on} < t < I_{off} \\ r_{BU} & \text{otherwise} \end{cases} \tag{8}$$

where $T_A$ is the afferent delay, $I_{on}$ and $I_{off}$ are the onset and offset of the interruption, respectively, $r_I$ is the build-up rate during the interruption, and $r_{BU}$ is the nominal build-up rate of the process. For each individual trial, the occurrence of an interruption was determined by drawing a random number according to the probability of interruption, $p_I$. For trials that were interrupted, $r_I = 0$, whereas for trials that were not interrupted $r_I = r_{BU}$. In non-interrupted trials, numerical integration gives the same RT as equation (1), given the same parameters.

**The idealized interruption.** In the simplest case (Fig. 1), RT distributions with and without interruptions in saccade planning were generated as described above, with the following parameters. The afferent delay, $T_A$, was drawn from a Gaussian distribution with a mean of 50 ms and a standard deviation (**SD**) of 5 ms. The nominal build-up rate was drawn from a Gaussian distribution with a mean of 7.7 and SD of 1.9 AU/ms. The onset and offset times were drawn from independent Gaussian distributions with means of 192 and 228 ms, respectively (for a mean interruption duration of 36 ms), and equal SDs of either 0 (Fig. 1c, red trace) or 8 ms (Fig. 1d, red trace). For simulation runs with no interruption (Fig. 1c,d, gray shade), the probability of interruption per trial was $p_I = 0$. For simulation runs with interruption, $p_I = 1$ (Fig. 1c, red trace) or $p_I = 0.7$ (Fig. 1d, red trace).

**Modeling experimentally measured dips in the RT distribution.** To simulate the experiment of Bompas and Sumner[47] (Fig. 2), equation (2) was applied directly. This method is better suited for quantitative comparison between experimental and model results than the integration method described above; while it also assumes a linear rise to threshold, it requires fewer parameters because it effectively takes the distribution of build-up rates ($r_{BU}$) directly from the data. To implement this method, the RT distribution in the no-distracter condition was digitally extracted from the published figure (Fig. 3 of ref.[47], observer 1, SOA = 0) and stored in the computer (Fig. 2b). The simulation then proceeded as follows.

In each trial, first, a nominal RT without interruption, $RT_0$, was sampled from the captured, empirical RT distribution with no distracter (Fig. 2b). This was done using the Matlab function randsample, which generates random samples from arbitrary distributions. Having drawn a sample $RT_0$, the interruption was then simulated for that trial just as in the previous case. Interruption onset and offset times, $I_{on}$ and $I_{off}$, were drawn from a bivariate Gaussian distribution with a correlation of $-0.8$, means equal to $\text{SOA} + 85$ and $\text{SOA} + 115$ ms, respectively, and identical SDs of 14.3 ms. The interruption duration in the trial was either $q = 0$, for non-interrupted trials, or else $q = I_{off} - I_{on}$, for interrupted trials, where any negative $q$ values were reset to zero. Finally, the simulated RT in the trial, $RT_I$, was obtained by inserting $RT_0$, $I_{on}$, and $q$ into equation (2). For the results in Fig. 2, all distracter-present trials were interrupted (i.e., $p_I = 1$). In total, this simulation required 5 parameters. The values just listed were determined by the optimization procedure described in the next subsection.

Note that the correlation parameter allowed us to consider a range of statistical scenarios, going from cases in which the interruption duration stays relatively constant but its center point varies (correlation $\approx 1$), to cases in which the center point remains constant but the duration of the interval varies across trials (correlation $\approx -1$). The nature of the variability has an interesting consequence, which is that the width of the dip, as measured empirically from the RT distribution, does not necessarily correspond to the mean interruption time in the model. For instance, for the results in Fig. 2, the mean interruption time is 30 ms, but the dip measured from the RT histogram (panel h, dark shade) is quite wider (~59 ms, width at half depth). This is because of the high SD and strong anticorrelation between $I_{on}$ and $I_{off}$.

In this case, a concurrent motor plan toward the distracter was not explicitly included in the simulation because, given the design of the task, it was always meant to be canceled. And indeed, errors were rare[47]. Presumably, a saccade plan toward the distracter did exist briefly, but there were no data to constrain it.

**Parameter optimization.**    When fitting the model to specific experimental datasets (Figs 2 and 5), optimal values for the model parameters were found by minimizing the mean absolute difference between the empirical and simulated RT distributions, averaged across all the relevant experimental conditions considered. For example, for Fig. 2, the total error was the average of the errors for SOAs of 0, 20, 40, 60, and 80 ms. The parameter combinations tested were picked by exhaustive search over a multidimensional grid.

**Numerical results complementing the theoretical calculations.**    For Fig. 4, the RT associated with each individual rise-to-threshold process was calculated analytically using the interruption duration, $q$, the value of $\alpha$, and the build-up rates (the initial one, inversely related to $Q$, and the build-up rate during the interruption, if nonzero). What was calculated numerically were the averages across trials as these parameters varied.

**The double-step experiment.**    The study by Buonocore and colleagues[39] included 3 experiments, each with 6 different participants. We discuss their experiment 1, but experiments 2 and 3, which used slightly different stimulation conditions, produced highly consistent results. We include and model data from their participant 3 in experiment 1. Results for the other 5 experiment 1 participants were highly similar (their Fig. 2). In double-step blocks, each participant completed 240 control and 240 step trials; in saccadic inhibition (i.e., distracter-step) blocks, each participant completed 240 control and 240 distracter trials. The the SOA for each participant was fixed, with the value optimized to maximize the dip in the RT distribution in saccadic inhibition blocks.

The integration method, i.e., equations (7) and (8), was used to simulate saccade plans in the double-step experiment (Fig. 5). The parameter values for the initial motor plan and for the interruption were as follows: the afferent delay, $T_A$, was drawn from a Gaussian distribution with a mean of 50 and SD of 10 ms; the initial build-up rate was drawn from a Gaussian distribution with a mean of 6.1 and SD of 1.7 AU/ms; the probability of interruption was $p_I = 1$; the build-up rate during the interruption was $r_I = 0$; the onset and offset times of the interruption were Gaussian samples with means of SOA + 53 and SOA + 157 ms, respectively, correlation equal to $-0.8$, and SD equal to 19.2 ms. The SOA was fixed at 120 ms, as was the case for participant 3 in the experimental study[39]. These values were identical between the simulated blocks of distracter-step and target-step trials.

In distracter-step trials, the initial motor plan continued after the end of the pause, and the RT was recorded as the time when it reached threshold. In this case the concurrent plan was omitted because very few errors were observed in the experiment, implying that such plan was almost always successfully canceled after the interruption. In target-step trials, a concurrent plan was initiated during the pause. The build-up rate of this plan was drawn independently for each trial, with a mean of 6.2 and SD of 1.7 AU/ms. In correct distracter-step trials, the first plan was assumed to be canceled after the pause and the RT was recorded as the time at which the second, concurrent plan reached threshold (as shown in Fig. 5d, middle). In error trials, the cancelation failed, the first plan continued after the pause, and the RT was recorded as the time at which this plan reached threshold (as shown in Fig. 5d, right). The rule for generating such failures was that they occured when the build-up rate of the initial plan exceeded that of the concurrent plan by at least 1.5 AU/ms. This value was set so that the simulated error rate matched the experimental rate (26%). Finally, whichever motor plan was still active after the pause, accelerated; that is, its build-up rate increased linearly between the end of the pause and threshold crossing. During distracter-step trials the initial plan accelerated at a rate of 0.03 AU/ms$^2$, and during target-step trials the concurrent plan accelerated at a rate of 0.01 AU/ms$^2$. The effect of these acceleration terms was to curtail the right tails of the RT distributions. They were not essential, as qualitatively similar results were obtained with zero acceleration; however, they are consistent with related saccadic choice models[5,16,76], and slightly improved the fits between the experimental (Fig. 5b,c) and model data (Fig. 5e,f).

## References

1. Hanes, D. P. & Schall, J. D. Neural control of voluntary movement inititation. *Science* **274**, 427–430 (1996).
2. Lo, C. C. & Wang, X. J. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* **9**, 956–963 (2006).
3. Brown, J. W., Hanes, D. P., Schall, J. D. & Stuphorn, V. Relation of frontal eye field activity to saccade initiation during a countermanding task. *Exp. Brain Res.* **190**, 135–151 (2008).
4. Ding, L. & Gold, J. I. Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* **22**, 1052–1067 (2012).
5. Costello, M. G., Zhu, D., Salinas, E. & Stanford, T. R. Perceptual modulation of motor–but not visual–responses in the frontal eye field during an urgent-decision task. *J. Neurosci.* **33**, 16394–16408 (2013).
6. Hauser, C. K., Zhu, D., Stanford, T. R. & Salinas, E. Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. *eLife* pii, e33456, https://doi.org/10.7554/eLife.33456 (2018).
7. Pelz, J. B. & Canosa, R. Oculomotor behavior and perceptual strategies in complex tasks. *Vision Res.* **41**, 3587–3596 (2001).
8. Berg, D. J., Boehnke, S. E., Marino, R. A., Munoz, D. P. & Itti, L. Free viewing of dynamic stimuli by humans and monkeys. *J. Vis.* **9**, 19.1–15 (2009).
9. Phillips, A. N. & Segraves, M. A. Predictive activity in macaque frontal eye field neurons during natural scene searching. *J. Neurophysiol.* **103**, 1238–1252 (2010).
10. Glimcher, P. W. & Sparks, D. L. Effects of low-frequency stimulation of the superior colliculus on spontaneous and visually guided saccades. *J. Neurophysiol.* **69**, 953–964 (1993).
11. Foerster, R. M., Carbone, E., Koesling, H. & Schneider, W. X. Saccadic eye movements in the dark while performing an automatized sequential high-speed sensorimotor task. *J. Vis.* **12** (2012).
12. Zhou, H. & Desimone, R. Feature-based attention in the frontal eye field and area V4 during visual search. *Neuron* **70**, 1205–1217 (2011).
13. Glaser, J. I. *et al.* Role of expected reward in frontal eye field during natural scene search. *J. Neurophysiol.* **116**, 645–657 (2016).
14. Gold, J. I. & Shadlen, M. N. Representation of a perceptual decision in developing oculomotor commands. *Nature* **404**, 390–394 (2000).
15. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).

16. Stanford, T. R., Shankar, S., Massoglia, D. P., Costello, M. G. & Salinas, E. Perceptual decision making in less than 30 milliseconds. *Nat. Neurosci.* **13**, 379–385 (2010).
17. Hafed, Z. M. & Ignashchenkova, A. On the dissociation between microsaccade rate and direction after peripheral cues: microsaccadic inhibition revisited. *J. Neurosci.* **33**, 16220–16235 (2013).
18. McPeek, R. M., Skavenski, A. A. & Nakayama, K. Concurrent processing of saccades in visual search. *Vision Res.* **40**, 2499–2516 (2000).
19. McPeek, R. M. & Keller, E. L. Short-term priming, concurrent processing, and saccade curvature during a target selection task in the monkey. *Vision Res.* **41**, 785–800 (2001).
20. Katnani, H. A. & Gandhi, N. J. Time course of motor preparation during visual search with flexible stimulus-response association. *J. Neurosci.* **33**, 10057–10065 (2013).
21. Hayhoe, M. & Ballard, D. Eye movements in natural behavior. *Trends Cogn. Sci.* **9**, 188–194 (2005).
22. Ballard, D. H. & Hayhoe, M. M. Modelling the role of task in the control of gaze. *Vis. cogn.* **17**, 1185–1204 (2009).
23. Nothdurft, H. C. Saliency effects across dimensions in visual search. *Vision Res.* **33**, 839–844 (1993).
24. Nothdurft, H. C. Attention shifts to salient targets. *Vision Res.* **42**, 1287–1306 (2002).
25. Theeuwes, J., Kramer, A. F., Hahn, S. & Irwin, D. E. Our eyes do not always go where we want them to go: capture of the eyes by new objects. *Psychol. Sci.* **9**, 379–385 (1998).
26. Theeuwes, J., Kramer, A. F., Hahn, S., Irwin, D. E. & Zelinsky, G. J. Influence of attentional capture on oculomotor control. *J. Exp. Psychol. Hum. Percept. Perform.* **25**, 1595–1608 (1999).
27. Connor, C. E., Egeth, H. E. & Yantis, S. Visual attention: bottom-up versus top-down. *Curr. Biol.* **14**, R850–R852 (2004).
28. Hikosaka, O., Nakamura, K. & Nakahara, H. Basal ganglia orient eyes to reward. *J. Neurophysiol.* **95**, 567–584 (2006).
29. Borji, A., Sihite, D. N. & Itti, L. What stands out in a scene? A study of human explicit saliency judgment. *Vision Res.* **91**, 62–77 (2013).
30. Adeli, H., Vitu, F. & Zelinsky, G. J. A model of the superior colliculus predicts fixation locations during scene viewing and visual search. *J. Neurosci.* **37**, 1453–1467 (2017).
31. Yantis, S. Attentional capture in vision. In *Converging operations in the study of visual selective attention* (eds Kramer, A. F., Coles, M. G. H. & Logan, G. D.) 45–76 (American Psychological Association, 1996).
32. Theeuwes, J. Top-down and bottom-up control of visual selection. *Acta Psychol. (Amst)* **135**, 77–99 (2010).
33. Gottlieb, J. P., Kusunoki, M. & Goldberg, M. E. The representation of visual salience in monkey parietal cortex. *Nature* **391**, 481–484 (1998).
34. Thompson, K. G., Bichot, N. P. & Sato, T. R. Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J. Neurophysiol.* **93**, 337–351 (2005).
35. Ipata, A. E., Gee, A. L., Gottlieb, J., Bisley, J. W. & Goldberg, M. E. LIP responses to a popout stimulus are reduced if it is overtly ignored. *Nat. Neurosci.* **9**, 1071–1076 (2006).
36. Joiner, W. M., Cavanaugh, J., Wurtz, R. H. & Cumming, B. G. Visual responses in FEF, unlike V1, primarily reflect when the visual context renders a receptive field salient. *J. Neurosci.* **37**, 9871–9879 (2017).
37. White, B. J. *et al.* Superior colliculus neurons encode a visual saliency map during free viewing of natural dynamic video. *Nat. Commun.* **8**, 14263 (2017).
38. Reingold, E. M. & Stampe, D. M. Saccadic inhibition in voluntary and reflexive saccades. *J. Cogn. Neurosci.* **14**, 371–388 (2002).
39. Buonocore, A., Purokayastha, S. & McIntosh, R. D. Saccade reorienting is facilitated by pausing the oculomotor program. *J. Cogn. Neurosci.* **18**, 1–13 (2017).
40. Anderson, B. A., Laurent, P. A. & Yantis, S. Value-driven attentional capture. *Proc. Natl. Acad. Sci. USA* **108**, 10367–10371 (2011).
41. Fecteau, J. H. & Munoz, D. P. Warning signals influence motor processing. *J. Neurophysiol.* **97**, 1600–1609 (2007).
42. Carpenter, R. H. S. & Williams, M. L. L. Neural computation of log likelihood in control of saccadic eye movements. *Nature* **377**, 59–61 (1995).
43. Reddi, B. A. & Carpenter, R. H. The influence of urgency on decision time. *Nat. Neurosci.* **3**, 827–830 (2000).
44. Brown, S. D. & Heathcote, A. The simplest complete model of choice response time: linear ballistic accumulation. *Cog. Psych.* **57**, 153–178 (2008).
45. Cisek, P., Puskas, G. A. & El-Murr, S. Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11560–11571 (2009).
46. Palmer, J., Huk, A. C. & Shadlen, M. N. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* **5**, 376–404 (2005).
47. Bompas, A. & Sumner, P. Saccadic inhibition reveals the timing of automatic and voluntary signals in the human brain. *J. Neurosci.* **31**, 12501–12512 (2011).
48. Reingold, E. M. & Stampe, D. M. Saccadic inhibition in reading. *J. Exp. Psychol. Hum. Percept. Perform.* **30**, 194–211 (2004).
49. Edelman, J. A. & Xu, K. Z. Inhibition of voluntary saccadic eye movement commands by abrupt visual onsets. *J. Neurophysiol.* **101**, 1222–1234 (2009).
50. Buonocore, A. & McIntosh, R. D. Modulation of saccadic inhibition by distractor size and location. *Vision Res.* **69**, 32–41 (2012).
51. Stampe, D. M. & Reingold, E. M. Influence of stimulus characteristics on the latency of saccadic inhibition. *Prog. Brain Res.* **40**, 73–87 (2002).
52. Glaholt, M. G. & Reingold, E. M. Perceptual enhancement as a result of a top-down attentional influence in a scene viewing task: Evidence from saccadic inhibition. *Q. J. Exp. Psychol. (Hove)* **2**, 1–9 (2016).
53. Harrison, J. J., Freeman, T. C. & Sumner, P. Saccade-like behavior in the fast-phases of optokinetic nystagmus: an illustration of the emergence of volitional actions from automatic reflexes. *J. Exp. Psychol. Gen.* **143**, 1923–1938 (2014).
54. Harrison, J. J., Sumner, P., Dunn, M. J., Erichsen, J. T. & Freeman, T. C. Quick phases of infantile nystagmus show the saccadic inhibition effect. *Invest. Ophthalmol. Vis. Sci.* **56**, 1594–1600 (2015).
55. Dorris, M. C., Olivier, E. & Munoz, D. P. Competitive integration of visual and preparatory signals in the superior colliculus during saccadic programming. *J. Neurosci.* **27**, 5053–5062 (2007).
56. White, B. J. *et al.* Competitive integration of visual and goal-related signals on neuronal accumulation rate: a correlate of oculomotor capture in the superior colliculus. *J. Cogn. Neurosci.* **25**, 1754–1768 (2013).
57. Khan, A. Z., Heinen, S. J. & McPeek, R. M. Attentional cueing at the saccade goal, not at the target location, facilitates saccades. *J. Neurosci.* **30**, 5481–5488 (2010).
58. Khan, A. Z., Munoz, D. P., Takahashi, N., Blohm, G. & McPeek, R. M. Effects of a pretarget distractor on saccade reaction times across space and time in monkeys and humans. *J. Vis.* **16**, 5 (2016).
59. Lauwereyns, J., Watanabe, K., Coe, B. & Hikosaka, O. A neural correlate of response bias in monkey caudate nucleus. *Nature* **418**, 413–417 (2002).
60. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).
61. Brunton, B. W., Botvinick, M. M. & Brody, C. D. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
62. Uchida, N., Kepecs, A. & Mainen, Z. F. Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat. Rev. Neurosci.* **7**, 485–491 (2006).

63. Zariwala, H. A., Kepecs, A., Uchida, N., Hirokawa, J. & Mainen, Z. F. The limits of deliberation in a perceptual decision task. *Neuron* **78**, 339–351 (2013).
64. Gawne, T. J., Kjaer, T. W. & Richmond, B. J. Latency: another potential code for feature binding in striate cortex. *J. Neurophysiol.* **76**, 1356–1360 (1996).
65. Edwards, R., Xiao, D., Keysers, C., Földiák, P. & Perrett, D. Color sensitivity of cells responsive to complex stimuli in the temporal cortex. *J. Neurophysiol.* **90**, 1245–1256 (2003).
66. Bodelón, C., Fallah, M. & Reynolds, J. H. Temporal resolution for the perception of features and conjunctions. *J. Neurosci.* **27**, 725–730 (2007).
67. Salinas, E. & Stanford, T. R. The countermanding task revisited: fast stimulus detection is a key determinant of psychophysical performance. *J. Neurosci.* **33**, 5668–5685 (2013).
68. Wolfe, J. M. What can 1 million trials tell us about visual search? *Psychological Science* **9**, 33–39 (1998).
69. Wolfe, J. M., Alvarez, G. A., Rosenholtz, R., Kuzmova, Y. I. & Sherman, A. M. Visual search for arbitrary objects in real scenes. *Atten. Percept. Psychophys.* **73**, 1650–1671 (2011).
70. Sommer, M. A. The spatial relationship between scanning saccades and express saccades. *Vision Res.* **37**, 2745–2756 (1997).
71. McPeek, R. M. & Keller, E. L. Superior colliculus activity related to concurrent processing of saccade goals in a visual search task. *J. Neurophysiol.* **87**, 1805–1815 (2002).
72. Klapetek, A., Jonikaitis, D. & Deubel, H. Attention allocation before antisaccades. *J. Vis.* **16**, 11 (2016).
73. Becker, W. & Jürgens, R. An analysis of the saccadic system by means of double step stimuli. *Vision Res.* **19**, 967–983 (1979).
74. Murthy, A. *et al.* Frontal eye field contributions to rapid corrective saccades. *J. Neurophysiol.* **97**, 1457–1469 (2007).
75. Fecteau, J. H. & Munoz, D. P. Correlates of capture of attention and inhibition of return across stages of visual processing. *J. Cogn. Neurosci.* **17**, 1714–1727 (2005).
76. Salinas, E., Shankar, S., Costello, M. G., Zhu, D. & Stanford, T. R. Waiting is the hardest part: comparison of two computational strategies for performing a compelled-response task. *Front. Comput. Neurosci.* **4**(153), 1–17 (2010).
77. Shankar, S. *et al.* Tracking the temporal evolution of a perceptual judgment using a compelled-response task. *J. Neurosci.* **31**, 8406–8421 (2011).
78. Steinberg, B. R. *et al.* Parsing the contributions of perception, motor planning, and cognition to anti-saccade performance in human subjects. *Society for Neuroscience Abstracts 2017*, Program No. 493.08, http://www.abstractsonline.com/pp8/#!/4376/presentation/12653 (2017).
79. Wessel, J. R. & Aron, A. R. Unexpected events induce motor slowing via a brain mechanism for action-stopping with global suppressive effects. *J. Neurosci.* **33**, 18481–18491 (2013).
80. Wessel, J. R. & Aron, A. R. On the globality of motor suppression: unexpected events and their influence on behavior and cognition. *Neuron* **93**, 259–280 (2017).
81. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife* **4**, e04677 (2015).
82. van den Berg, R. *et al.* A common mechanism underlies changes of mind about decisions and confidence. *Elife* **5**, e12192 (2016).
83. Thura, D. & Cisek, P. Modulation of premotor and primary motor cortical activity during volitional adjustments of speed-accuracy trade-offs. *J. Neurosci.* **36**, 938–956 (2016).
84. Stanford, T. R. Decision making, motor planning. In *Encyclopedia of Computational Neuroscience* (eds Jaeger, D. & Jung, R.) https://doi.org/10.1007/978-1-4614-7320-6_313-1 (Springer, 2014).
85. Sato, T. R. & Schall, J. D. Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron* **38**, 637–648 (2003).

## Acknowledgements

## Author Contributions

E.S. and T.R.S. designed the study; E.S. performed analyses and simulations; E.S. and T.R.S. wrote the manuscript.

## Additional Information

**Supplementary information** accompanies this paper at https://doi.org/10.1038/s41598-018-32224-5.

**Competing Interests:** The authors declare no competing interests.

**Publisher's note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.