Temporal integration of sensory evidence for saccade target selection

Casimir J.H. Ludwig

University of Bristol, Department of Experimental Psychology, 12a Priory Road, Bristol BS8 1TU, UK

ARTICLE INFO

ABSTRACT

Perceptual decision-making may be viewed as a process of integrating sensory evidence over time to a response threshold. Ludwig et al. (2005b) argued against an integration to threshold account for saccadic eye movement decisions based on data from a stochastic contrast discrimination task. They argued that evidence integration was time-limited, with the deadline being independent from the quality of the sensory evidence. In this study, the data from Ludwig et al. (2005b) were fit with a model in which sensory evidence is integrated to a time-varying threshold. The functional form of the threshold variation allowed the model to approximate a constant threshold as well as an abrupt deadline. Sensory evidence was computed on the basis of a temporally blurred representation of the stimulus sequence. The model provides an overall good fit to the latency and accuracy data. Its predictions are consistent with a short window of evidence integration, as proposed in the original study. The model produces qualitatively correct predictions for an experiment in which the availability of sensory evidence is varied systematically. Inspection of the model parameters, however, shows that integration was not terminated by an abrupt deadline, although a rather gradual deadline signal improved the fit for some observers.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Studies of the neural basis of (perceptual) decision-making have highlighted a role for mechanisms that integrate sensory evidence over time in order to come to an informed decision (reviewed in Gold & Shadlen, 2001, 2007). Temporal integration is a particularly useful strategy when there is noise or uncertainty in the sensory signal and system. For instance, in a typical paradigm the observer views a random dot kinetogram (RDK) in which only a subset of dots moves coherently in one direction. Motion discrimination thresholds and the time needed to make a decision typically decrease as the signal-to-noise ratio of the RDK increases (Cook & Maunsell, 2004; Palmer, Huk, & Shadlen, 2005; Roitman & Shadlen, 2002). These findings have been modelled using neural mechanisms involved in programming the motor response that integrate the evidence provided by upstream sensory areas up to some criterion level of evidence (Ditterich, 2006a, 2006b; Huk & Shadlen, 2005; Mazurek, Roitman, Ditterich, & Shadlen, 2003).

In the psychological domain, models based on the accumulation of evidence over time have a long and rich history (Brown & Heathcote, 2005; Busemeyer & Townsend, 1993; LaBerge, 1962; Laming, 1969; Link & Heath, 1975; Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Smith, 1995; Smith & Vickers, 1988; Usher & McClelland, 2001; Vickers, 1970). Although a wide variety of models exist that differ in their precise assumptions and dynamics,

they share the common idea that some noisy internal signal is integrated over time until a decision criterion is exceeded. The broad class of sequential sampling models may be regarded as a natural extension of signal detection theory (Green & Swets, 1966) to the temporal domain. Whereas detection theory assumes decisions are based on a single noisy sample of the information, sequential sampling models assume multiple (noisy) samples are acquired. This extension allows these models to account for response time in addition to choice.

Integration to threshold models have also been used to account for distributions of saccade latencies in a variety of tasks (Carpenter & Williams, 1995; Hanes & Carpenter, 1999; Ludwig, Gilchrist, & McSorley, 2005a; Ludwig, Mildinhall, & Gilchrist, 2007; Reddi, Asrress, & Carpenter, 2003; Reddi & Carpenter, 2000). However, in a recent paper Ludwig, Gilchrist, McSorley and Baddeley (2005b) reported evidence from a saccadic choice task that appeared inconsistent with integration of sensory evidence to threshold. In their task observers were presented with patterns that fluctuated in luminance contrast over time (in 25 ms steps). One of the patterns had, on average, higher contrast than the other and the observers' task was to respond to the high contrast patch (Caspi, Beutter, & Eckstein, 2004, see also). An illustration of the stimulus is provided in Fig. 1A. As with the RDK, the stochastic nature of this stimulus enables more accurate decisions to be made the more temporal samples are taken into account. By relating the noise in the patterns to the saccadic choices on a trial-by-trial basis, Ludwig et al. (2005b) were able to show that only the luminance samples in the first ~100 ms were taken into account when choosing between the patterns (Fig. 1B).

E-mail address: c.ludwig@bristol.ac.uk
URL: http://casimir.psy.bris.ac.uk
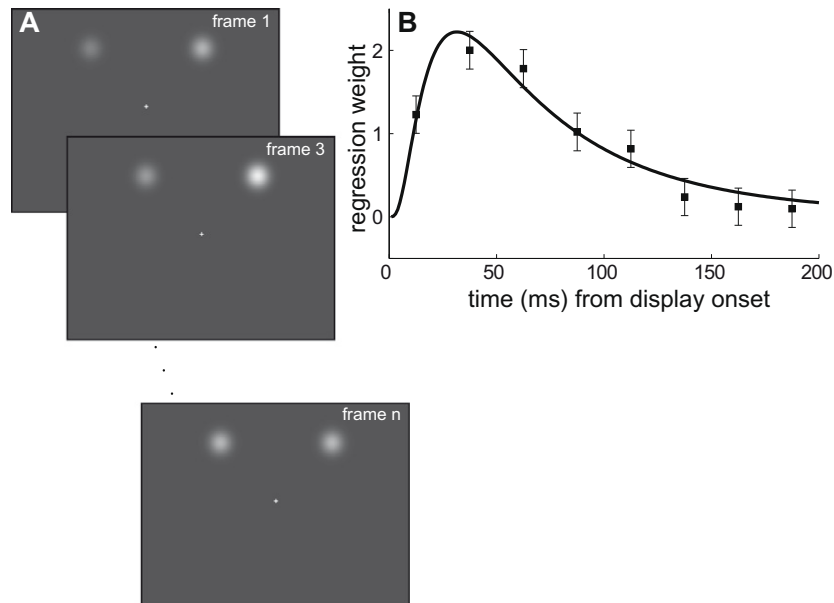
Fig. 1. The temporal noise paradigm (Ludwig et al., 2005b, experiment 1). (A) Two patterns fluctuate in contrast by drawing a new peak contrast from a Gaussian distribution every two video frames. The observer has to look at the pattern that is, on average, of higher contrast (top-right in this example). Only 3 out of a total of 80 video frames are illustrated. (B) Illustrative data from one observer. Shown are logistic regression weights that indicate the importance the observer attributed to the stimulus information over time. The characteristic pattern is that only the first ~100 ms appear to contribute to the decision making process.

In a subsequent experiment the authors modified their stimulus so that the mean contrast of the two patterns only differed in a 100 ms "critical window". This window was placed either right at the beginning of the temporal sequence or after an initial 100 ms had passed. Under these conditions, good performance was maintained when the critical window was presented early, but performance fell to chance when the critical window was presented late. This finding suggested that observers did not extend their integration window in the face of weak sensory evidence (the late critical window condition) to enable more accurate decisions. As such, the results appeared to discount an integration to threshold account of saccadic decisions in this particular context. Ludwig et al. (2005b) proposed that the eye movement decisions were based on integration of sensory information up to an internally set deadline (corresponding to ~100 ms of stimulus information). Instead of adjusting the integration period on a trial-by-trial basis in accordance with the quality of sensory evidence, observers integrate the evidence from the two pattern locations to a deadline, after which the saccade is directed to whichever pattern triggered the largest internal response.

Integration to threshold is currently the "gold-standard" model of perceptual decision-making. It accounts for a wide variety of behavioural data and is extensively supported by neurophysiological evidence from the RDK paradigm. Before concluding that a radically different decision strategy was used in the contrast discrimination task of Ludwig et al. (2005b) it is fair to demand a more formal, quantitative examination of the integration to threshold model.

This is all the more pressing, because in recent years the model has been modified to include a deadline signal (Churchland, Kiani, & Shadlen, 2008; Ditterich, 2006a, 2006b; Smith, 2000). The basic idea is that the integrated decision variable is brought closer to the response threshold as time elapses, *even in the absence of any sensory evidence in favour of one particular alternative*. One instantiation of this idea is that of a time-varying response threshold, which is moved closer to the starting point of integration. The qualitative deadline model put forward by Ludwig et al. (2005b) may be seen as an extreme version of this idea, in which the

boundary changes rapidly and early (see Fig. 2 and its description below). The important point is that both the standard integration to threshold model and the deadline model can be accommodated within the same framework. As such, fitting a time-varying threshold model to the data from Ludwig et al. (2005b) may shed light on the question of whether saccadic decisions in their paradigm were indeed driven by a deadline signal.

In this regard, the data from Ludwig et al. (2005b, experiment 1) have the potential to be extremely illuminating. The data set is large, with the number of two-choice trials ranging from 4200 to 6200 across four observers. With accuracy around threshold (75% correct) the distributions of erroneous responses were relatively well-defined. Latency distributions of correct and, especially, erroneous choices may be extremely informative and diagnostic with regard to the underlying decision processes (Ratcliff & Rouder, 1998).

In the neural models applied to the RDK data (Ditterich, 2006a, 2006b; Huk & Shadlen, 2005; Mazurek et al., 2003) a direct estimate of the internal sensory response is available through the recorded firing rates of MT neurons with the task-relevant directional selectivities. These firing rates are then used to compute a decision variable, which is integrated over time downstream (in lateral intraparietal area, LIP). As a result, the rate of accumulation (or drift rate) is completely determined by the internal sensory response. In the psychological domain, this internal representation cannot be measured directly and the behavioural task is frequently too complex to allow for computational specification. As a result, the accumulation rate and its variability across different trials are free parameters. In other words, it is typically left unspecified what the accumulated quantity actually represents. However, the noisy luminance contrast stimulus sequence used by Ludwig et al. (2005b) lends itself well to estimation of the momentary sensory evidence: standard models of early temporal vision (for a review see Watson, 1986, chap. 6) may be used to derive the internal sensory response which forms the input into the saccadic system. That is, as in the neural models described above, the accumulation rate and its variability (over time and across trials) are completely determined by the sensory response (Smith, 1995).
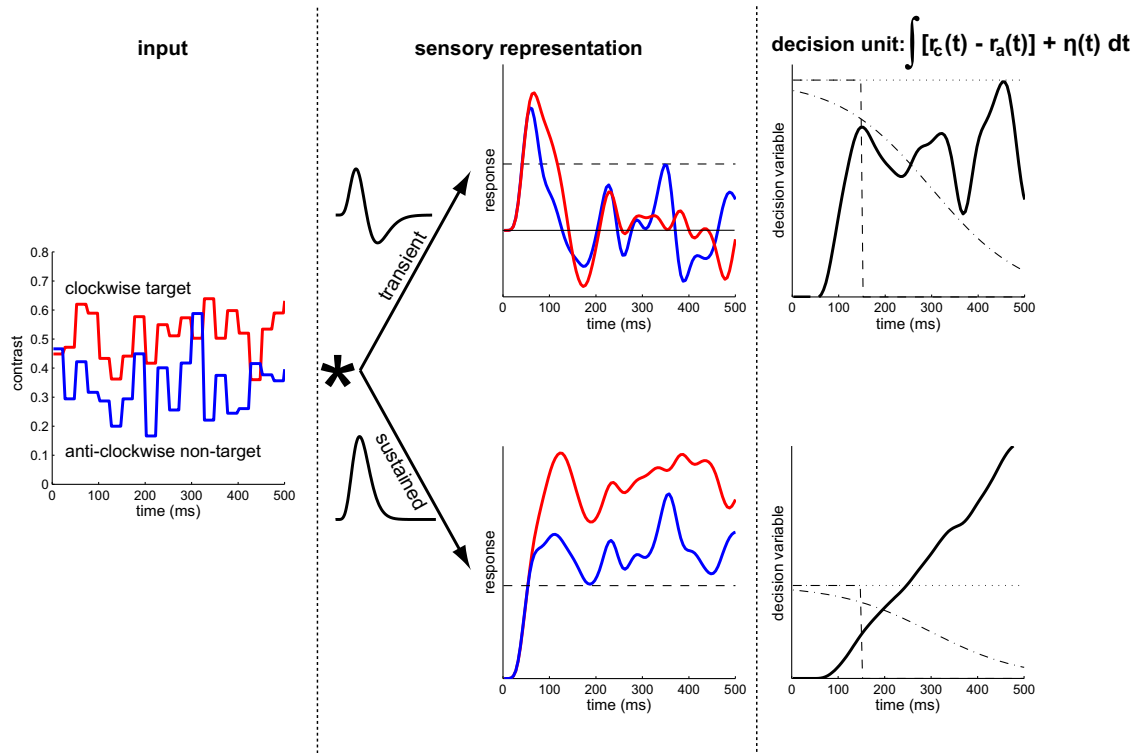
Fig. 2. Model architecture. The visual input (time series of contrast values) is convolved with a single temporal filter. A transient (top pathway) and sustained (bottom pathway) filter are illustrated; in the model fits the degree of "transience" was varied systematically. The sensory representation is a temporally filtered version of the input patterns. An optimal decision variable is the difference between the noise-perturbed internal responses, where a positive difference indicates evidence in favour of the clockwise pattern (the target in this example). Temporal integration of the noisy decision variable only starts when the sensory response to one of the patterns has exceeded some minimum level, indicated by the dashed horizontal lines in the middle panels. A decision is made when the integrated decision variable reaches a boundary, where a positive value indicates a clockwise response (the negative boundary is not shown in the right-hand panels, given the overwhelming evidence in favour of the clockwise pattern). Three different boundaries are shown, which capture a standard, constant threshold model (dotted line), the deadline model proposed by Ludwig et al. (2005b, dashed line), and a more gradual reduction in the threshold that may be regarded as a hybrid between these two extremes (dashed-dotted line). All different thresholds are instantiations of the same underlying logistic function.

2. Model description

Like the neural models described above, the model tested in this study consists of a sensory representation, an integrator and a decision rule. A schematic of the sensory representation and the integrator is provided in Fig. 2. The precise mathematical details of the model and the procedures used to fit it to the data from Ludwig et al. (2005b) are given in Appendices.

The temporal noise stimulus is a sequence of contrast values (Fig. 2—'input'). The sensory representation of the two patterns is a temporally filtered version of the time series of contrast values. There is a wealth of psychophysical evidence for the existence of a limited number (generally estimated to be 2 or 3) of temporal mechanisms (Breitmeyer, 1975; Fredericksen & Hess, 1998; Hess & Snowden, 1992; Legge, 1978; Tolhurst, 1975b, 1975a). The response of a transient mechanism peaks rapidly and then drops towards baseline (Fig. 2—'sensory representation', top row). Subsequent fluctuations in the response are the result of the stochastic nature of the stimulus. The magnitude of these fluctuations are the same for both the target and the non-target. As a result, there is little information in the transient response beyond the initial burst of activity (Muller, Metha, Krauskopf, & Lennie, 2001). Indeed, this is one explanation for the short integration period observed by Ludwig et al. (2005b): if the saccadic system only had access to transient mechanisms there would be little point in integrating the input after the response has returned to baseline. In contrast, a sustained mechanism will respond to the onset of the patterns and then maintain this response at a relatively constant, elevated level (Fig. 2—bottom row).

In the presence of internal noise, a noisy response (sample) from a visual mechanism can be probabilistically assigned to the target or non-target category. The likelihood-ratio (Gold & Shadlen, 2001; Green & Swets, 1966) expresses the relative evidence an internal response sample provides in favour of the target or non-target being in the "receptive field" of the mechanism (Eq. (B1)). With one mechanism response to each pattern, two likelihood-ratios may be combined multiplicatively to form an optimal decision variable.

To work out the probability of a noisy response having been triggered by a target or non-target, the decision unit would need to know about the expected mean internal response (at any given point in time) as well as its variance, to both the target and the non-target. In other words, it would need to know the probability density functions of target and non-target induced responses. If the brain was at all capable of extracting and storing such information, it would need to build up its "database" over a large number of trials. As argued by Gold & Shadlen (2001) it would seem unlikely the brain would go through such efforts, especially because internal responses may be affected by a large number of factors (e.g. background luminance, spatial stimulus characteristics, etc.).

However, as outlined in Appendix B (Eqs. (B.2) and (B.5)) an equivalent decision variable is formed simply by the difference between the two responses. The intuitive explanation is that under conditions of equal variance, the larger the internal response the more likely it was triggered by the target pattern. If the "anticlockwise" response of Fig. 2 is subtracted from the "clockwise" response, a positive difference implies that the internal response to the clockwise pattern was greater than that to the anti-clockwise

pattern. As such, the clockwise pattern is more likely to be the target. A negative response difference implies the opposite: a greater internal response to the anti-clockwise pattern, which is therefore more likely to correspond to the target. The advantage of this equivalent decision variable is that it does not require implausible assumptions about the kind of information the decision unit has access to.

The model assumes the difference between the noisy internal responses is integrated over time at the level of the decision unit. One additional option implemented in the current model is for integration to start once the internal sensory response (to either pattern) has exceeded some minimum level (Purcell et al., 2008). The idea was that the early criterion could act as a gating mechanism so that the decision unit only becomes engaged when the sensorium is sufficiently excited. As it turned out, the sharp onset in the visual responses coupled with relatively large amounts of internal noise ensured that the sensory criterion is exceeded very rapidly.

Temporal integration of the sensory evidence continues until a boundary is reached. The temporal variation in the boundary was modelled as a logistic function, allowing for the wide range of boundary profiles shown in Fig. 2. The figure illustrates how this function can approximate a nearly constant threshold over the range of response latencies of interest (Fig. 2, right panels—dotted lines). The deadline model proposed by Ludwig et al. (2005b) may be approximated by a very sharp change in the threshold (dashed lines). An intermediate, more gradual collapse in the response boundaries may be regarded as a hybrid between these two extremes (dashed-dotted lines).

One important point to note is that the model in its current form predicts the probability density of *decision* latencies. It is clear that not the whole response latency period is taken up by integrating the sensory evidence. The overall response time is simply the sum of the decision and a "non-decisional" delay. The latter is subject to its own (Gaussian) variability and is independent from both the decision latency as well as the stimulus properties.

Table 1 lists the seven free parameters of the model. Best-fitting parameters were found that maximised the log-likelihood of the observed choice and latency data. The fitting procedures used are described in Appendices C and D.

3. Results

Best-fitting model parameters were found for five levels of transience, ζ: 0 (purely sustained), 0.25, 0.50, 0.75, and 1 (purely transient). The intermediate filters may be regarded as a weighted combination of purely sustained and transient filters. Transience was not treated as a free parameter because the goodness-of-fit as a function of transience was in itself of interest. That is, if models with a predominantly transient filter consistently provide a better account of the data, this would suggest a natural explanation for the short evidence integration period (see Section 2). As it turned out, better fitting models lay at the sustained end of the spectrum, with the purely transient model always resulting in the worst fit.

However, there was some evidence of a moderate transient component for observers 2 ($\zeta = 0.50$) and 4 ($\zeta = 0.25$).

The model parameters are listed in Table 2. The corresponding fits to the distributions of correct and error saccade latencies are shown in Fig. 3. Clearly, the model fits follow the observed distributions very closely and, by implication, the model accounts for choice accuracy as well as latency. The model successfully captures the shorter latency of erroneous decisions. Although fast errors are a common empirical observation, accounting for the shape of error latency distributions has been a critical test of models of choice and reaction time (Ratcliff & Rouder, 1998).

The insets in Fig. 3 illustrate the variation of the positive response boundary over time (the negative boundary is the mirror image). It is clear that for 2/4 observers the threshold may be regarded as effectively constant (the abrupt change in threshold for observers 2 and 4 occurred outside the data range). For 2 other observers there is some evidence that the threshold changes during an epoch in which it will impact on the decision outcome. However, even for these observers the threshold does not begin to approximate the abrupt deadline postulated by Ludwig et al. (2005b).

Having optimised the model on the basis of the latency distributions and overall choice accuracy, its predictions regarding the temporal weighting of the different luminance samples were examined. Recall that Ludwig et al. (2005b) showed higher temporal weights for the first 4–5 samples, corresponding to the first ~100 ms after display onset (see Fig. 1). Fig. 4A shows the results from 1000 simulations of the model, using the median parameters across observers. Each simulation involved a similar number of trials as that used in the empirical study (5000). Logistic regression was used to derive the temporal weighting function. Panel A shows the mean weights, along with their 95% confidence intervals. It can be seen that these weights are significantly greater than 0 for the first 100–125 ms. The model predicts that the very first luminance sample is given the largest weight. This is not consistent with human behaviour, as can be seen in Fig. 1. This discrepancy is not surprising, given that the sharp onset in the visual response ensures that evidence integration starts essentially from the beginning. It is plausible that in humans there would be variability in the onset time of integration, which would mainly affect the temporal weight given to the first luminance sample.

The same set of parameters were used to simulate a "critical window" experiment similar to that reported in Ludwig et al. (2005b, Experiment 2). Recall that observers were presented with a mean luminance difference either just in the first 100 ms of the sequence (the early window), in the second 100 ms (the late window), or throughout the entire trial (baseline). The results seemed intuitively inconsistent with an integration to threshold account. Observers did not delay their decisions in the late window condition in order to wait for more useful sensory evidence. As a result, accuracy dropped in the late window and overall latency remained nearly constant.

The experiment was simulated 1000 times, with the number of trials per condition set to approximately the same number used in the original study (800). Fig. 4B shows the model predictions for choice accuracy. There is a small drop in performance, relative to

Table 1
Overview of free model parameters.

Parameter	Description
ξ	Gain of the temporal filter
θ	Boundary separation at $t = 0$
m	Location of the logistic threshold function
s	Scale of the logistic threshold function
μ_{nd}	Mean non-decisional delay
σ_{nd}	Standard-deviation of non-decisional delay distribution
λ	Early criterion on the sensory response

Table 2
Best-fitting model parameters.

Observer	ξ	ζ	θ	m	s	μ_{nd}	σ_{nd}	λ
1	4.13	0.00	12.34	515.09	165.15	196.85	22.05	2.09
2	7.29	0.50	9.86	905.04	15.49	164.81	10.82	2.33
3	2.58	0.00	8.16	632.43	96.20	211.74	14.70	1.57
4	3.14	0.25	10.18	1000.00	0.01	182.93	12.54	1.38
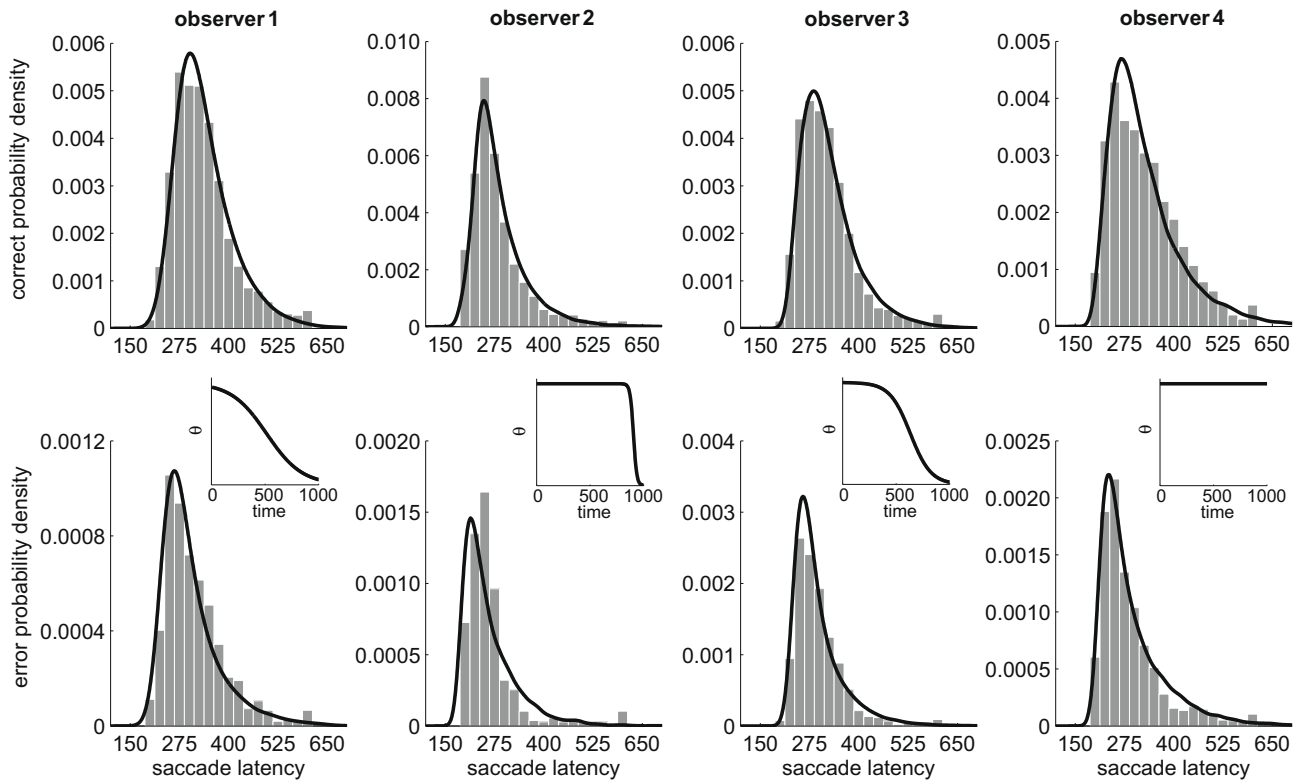
Fig. 3. Observed (histograms) and fitted (solid lines) probability density functions of the latencies of correct (top row) and erroneous (bottom row) saccadic choices. The insets show the variation in the (positive) response boundary over time for each observer.
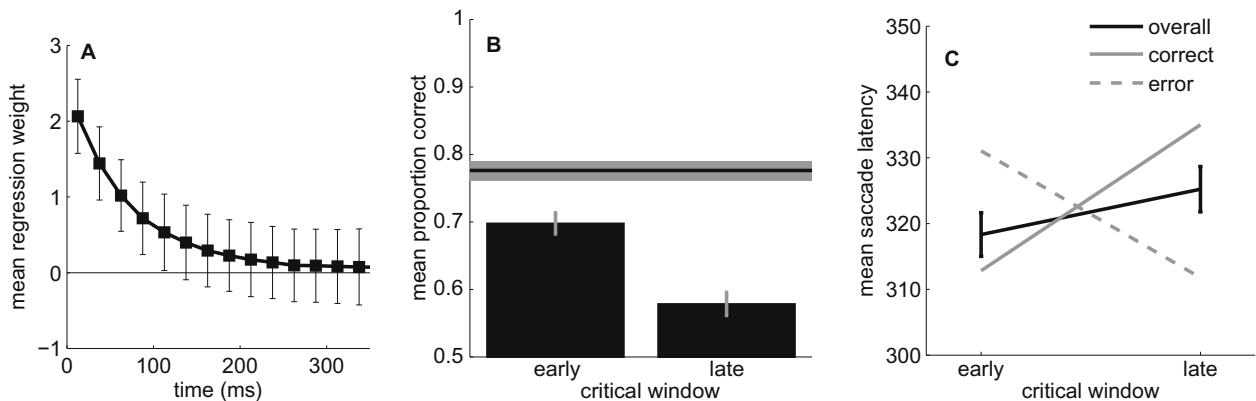
Fig. 4. Model predictions for the temporal weighting function (A), choice accuracy under critical window conditions (B) and saccade latency under critical window conditions (C). The model was simulated 1000 times, each time with a similar number of trials as that used in the original experiments. A—logistic regression weights associated with the noisy luminance samples. Error bars correspond to 95% confidence intervals, because the critical comparison is between the data points and 0. B—predicted accuracy under critical window conditions. The horizontal line indicates baseline performance. Error bars and the shaded region correspond to the standard deviation across the 1000 replications, because the critical comparisons are now between individual data points. C—predicted mean latency for the two critical window conditions. Error bars correspond to the standard deviations across 1000 replications.

baseline, for the early window condition. For the late window condition performance is close to chance. Fig. 4C illustrates the latency predictions. Overall latency increases only slightly in going from the early to the late window condition (by about 7 ms. in this particular simulation). However, given the predicted variability it is clear that this subtle effect will be difficult to detect empirically. The drop in accuracy for the late window and the near constant overall saccade latency were the most salient aspects in the data from Ludwig et al. (2005b).

Interestingly, the critical window position has striking and differential effects on the latencies of correct and error decisions. Cor-

rect latency increases for the late window condition, whereas the error latency goes down. The reason for the predicted cross-over is that for the early critical window a long integration period would include more noise, thereby enhancing the likelihood of an error decision. The reverse applies to the late critical window: a short integration period would mainly include noise (or at least, relatively less signal) and be associated with a greater likelihood of making an error. This predicted cross-over was not reliably present in the original data. However, note that the original critical window experiment was run on a different set of observers (bar one) than the ones from which the parameters for this simulation were

derived. In addition, the presence or absence of a cross-over depends critically on the position of the decision boundaries (simulations not shown here). This suggests that a critical empirical test of the model should combine the critical window manipulation with a variation in decision criteria (e.g. through speed–accuracy instructions).

4. Discussion

The model used in the current study relies on temporal integration of an optimal decision variable, derived from a temporally blurred version of the stimulus sequence. A decision is made when the integrated decision variable reaches one of two response boundaries, which represent the two alternative saccade target locations. In this sense, the model provides an integrated account of the 'where' and 'when' of saccadic decisions (Findlay & Walker, 1999).

4.1. Interpretation of model parameters

The model accounts well for the distribution of latencies of correct and error saccades, provided that the sensory representation is derived from a relatively low-pass filter. Of course, in the context of this particular task a sustained filter is sensitive to the most informative components of the stimulus, namely the average contrast. Although far from conclusive, the finding that models based on predominantly transient filters do not work nearly as well may suggest that the spatio-temporal properties of the stimulus (low spatial frequencies, relatively high temporal frequencies) did not "oblige" observers to construct their decision variable from a transient internal response. Taking the model parameters from the fits to the defective latency distributions, it naturally predicted a relatively short temporal weighting function. Qualitatively correct predictions concerning choice and latency were produced for "unseen" data from a critical window experiment that was not used to optimise the model parameters.

The finding that the model accounts for shorter-latency error decisions is of particular interest. A diffusion model without trial-to-trial variability in the drift rate and starting point predicts correct and error latencies of the same magnitude (Ratcliff, 1978; Ratcliff & Rouder, 1998). Noise in the drift rate allows such a model to account for long-latency errors, whereas noise in the starting point of integration allows the model to account for short-latency errors. Neither mechanism was included in the present model, apart from the trial-to-trial variations in the drift rate produced by the luminance noise in the stimulus. It turns out that the model accounts for short-latency errors through its temporally blurred sensory response. Unlike in the standard diffusion model where the mean drift for any one trial is constant, the temporal filter ensures that the early sensory responses to both patterns are very similar. Therefore, the emergence of evidence in favour of one or the other alternative is relatively gradual (cf. Fig. 1). As a result, the integrated decision variable hovers around the starting point early on in processing. Given the considerable noise added to the integrator, this is the "best" time for the integrator to end up at the wrong boundary, producing an error. That is, once an internal response difference emerges it is likely to point towards the target, making an error response less and less likely.

Although starting point variability is an entirely reasonable and plausible mechanism in any biological system, it is noteworthy that the present model can account for short-latency errors without including this form of noise. This explanation of short-latency errors is reminiscent of the notion of "pre-stimulus sampling" of Laming (1968), except—of course—that in this instance the 0-drift period occurs after stimulus onset. Ratcliff (2002) simulated data

from a model with a ramped onset in the drift rate and then fit the simulated data with the standard diffusion model that includes starting point variability and assumes an instantaneous onset in the drift. The ramped onset resulted in, among other things, an increase in the estimated starting point variability, suggesting that these different mechanisms may be difficult to distinguish.

The model parameters indicate that a deadline signal may have influenced decision-making for some observers, although this signal is far removed from the abrupt deadline proposed originally. In the current model, the deadline signal was imposed on the response boundary. In neurophysiological data from LIP, a deadline signal was imposed on the activity of the neural integrator itself (Ditterich, 2006a). Functionally and mathematically, a time-dependent reduction in the boundary separation and an increase in the integrated decision variable are equivalent (Smith, 2000).

It may be argued that the existence of some sort of deadline is adaptive under conditions in which trials are terminated by the perceptual response and a reward is collected (Ditterich, 2006a): when the sensory evidence is weak there will come a point in time at which it is better to make a guess and risk losing the reward for that trial, in order to move on to the new reward opportunity provided by the next trial. Although no explicit reward scheme was in operation in Ludwig et al. (2005b), some observers may have invoked a deadline signal to ensure they responded well within the stimulus presentation time (1 s). Perhaps more importantly, integrating sensory information from the peripheral field for a prolonged period of time is quite unnatural behaviour: with limited peripheral acuity it will, after some interval, be more adaptive to actively inspect the peripheral region of interest with the high resolution fovea. The deadline signal that emerged under these conditions may be the mechanism through which, in the natural world, active visual exploration is promoted in favour of prolonged integration of degraded sensory information.

Interestingly, for the model to approximate the empirical data the assumed decision latencies were relatively short (around 125 ms), regardless of the underlying temporal filter. Consequently, the overall non-decisional delay typically made up approximately half of the overall response latency. This is a puzzling finding: What is to stop the eyes from moving as soon as the decision is made, apart from the relatively small efferent delays to get the eye muscles in motion? These efferent delays are typically estimated to lie around 25 ms, based on micro-stimulation of the superior colliculus (Robinson, 1972), frontal eye fields (Robinson & Fuchs, 1969), and posterior parietal cortex (Thier & Andersen, 1998). That said, in similar models as the one tested here non-decisional delays of the same or even greater magnitude are by no means exceptional (Churchland et al., 2008; Ditterich, 2006a, 2006b; Mazurek et al., 2003; Ratcliff, Cherian, & Segraves, 2003; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007).

In the current models, it was important to assume some form of variability in the non-decisional delay. As a result, variability in saccade latency is not very tightly linked to variability in the decision latency on a trial-by-trial basis (Yang & McConkie, 2001). Indeed, this was one of the major conclusions to emerge from the original critical window experiment and the current, relatively indirect, analysis underlines this point once more. One implication is that current models of choice and response time still have some way to go in providing a complete process account of response latency. Instead of lumping the various non-decisional delays together in some parameter that is of relatively minor importance, a complete model would include the computational processes that give rise to the different delays. One critical question in this regard is to what extent the processes that make up the non-decisional delay really are not involved in the choice process (e.g. could there be independent sources of noise that affect choice, even after the "decision threshold" has been crossed?).

4.2. Model mis-specification and non-decisional delay period

It is of course possible that the long non-decisional delay is simply a consequence of mis-specification of the model itself, which may have occurred at a number of levels. Take the sensory response. First, temporal filter parameters were taken from previous fits to psychophysical detection data (Watson, 1986, chap. 6). It is not clear whether the same parameters apply in the supra-threshold domain. Then again, the finding of a long non-decisional delay held regardless of the transience of the filter and also for a number of different filter time constants (fits not reported here).

Second, the visual "system" in the model is entirely linear: there is no rectification or non-linear transformation of the internal responses (Albrecht, Geisler, & Crane, 2003; Gorea & Tyler, 1986). Given that the best models assumed a relatively sustained, elevated sensory response there is no explicit need for rectification. Non-linear compression or expansion of the internal response would predominantly affect the larger sensory response to the target. At the level of the decision unit the effect of such non-linearities would be to decrease (in the case of compression) or increase (in the case of expansion) the mean drift in the decision variable. However, the same effect(s) may be achieved by altering the gain of the filter: e.g. reducing the gain would decrease the mean drift and therefore increase the time it takes to reach a response boundary. In other words, the effects of such simple non-linearities could have been effectively absorbed by a different mechanism in the current architecture.

At the level of the decision unit itself integration was assumed to be perfect. Other models of decision-making assume leaky integrators (Ditterich, 2006b; Usher & McClelland, 2001). Passive decay or leakage in the integrator would act as a force to push its activity back to baseline, thereby potentially prolonging the time it takes to reach the response criterion. Leakage is typically included in models to prevent unlimited growth in the decision variable. In this sense, leakage is difficult to distinguish from the effects of more transient inputs from upstream. Both a transient filter and leakage produce a plateau in the activity of the integrator (although for different reasons—see top pathway in Fig. 2). Given this capacity for transience to mimic leakage, it is likely that there is not much to be gained from the additional inclusion of passive decay in this particular model.

Both non-linear compression and leakage are natural mechanisms to produce longer decision times. However, ultimately the empirical data from Ludwig et al. (2005b) provide a strong empirical constraint, which is not disputed here: whatever the internal evidence integration period is, it has to correspond to only the early section of the stimulus sequence. In other words, the integrated decision variable at a late point in time (say, 200 ms) would need to reflect at least partially the stimulus information presented long before (say, 100 ms). In the current architecture there is only one way to achieve this, namely by altering the temporal characteristics of the filter so that it introduces this large an amount of temporal blur. It is likely that the temporal properties of such a filter would lie beyond the range of physiological and psychophysical plausibility.

A more serious form of mis-specification may have occurred at the level of the model architecture itself. The main reason for using a bounded diffusion model in this study was to test whether the same type of model that successfully accounts for both neural and behavioural data from the motion discrimination paradigm, could account for saccadic decisions in a contrast discrimination task. One alternative is to assume two accumulators, each coding one response alternative, and integrating directly the sensory response associated with that alternative. Such a race model does not rely on the computation of *net* evidence (i.e. a response difference). Lateral inhibition between accumulators may be included to ensure that evidence in favour of one response counts as evidence against the alternative (Brown & Heathcote, 2005; Usher & McClelland, 2001). In this way, a competitive accumulator model may approximate something that looks like the optimal decision variable. At least for quick saccadic decisions it would seem likely that such a rough and ready decision variable is quite adequate. However, it is well established that such accumulator models are difficult to distinguish from the bounded diffusion model on the basis of behavioural data alone (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Ratcliff & Smith, 2004). More importantly, although it is possible that a different—and in some respects neurally more plausible—architecture would result in more prolonged decision times, the same empirical constraint applies. That is, it is not at all clear why a race model would produce longer decision latencies that are still only based on the early portion of the temporal stimulus.

4.3. Conclusions

Ludwig et al. (2005b) presented what appeared to be a challenging data set for models that assume sensory evidence is integrated up to a response threshold. They suggested saccadic eye movement decisions were governed by a temporal deadline that was not adjusted according to the quality of sensory evidence. This radical departure from the prevailing view is not necessary: a time-varying threshold model provides a good account for "seen" (Ludwig et al., 2005b, experiment 1) and "unseen" (the critical window paradigm) data, provided a relatively low-pass temporal filter was used to represent the sensory evidence. Importantly, the temporal variation in the threshold, if present, was rather gradual and did not resemble an abrupt deadline. Finally, it is important to note that the aim of this study was not to prove that a particular model of sensori-motor decision-making is "correct". Rather, model development here served to test the extent to which an empirical data pattern was diagnostic with regard to the theoretical issue at stake. Herein lies the major importance of specifying and fitting computational models of the decision-process.

Acknowledgments

Appendix A. Temporal filter

At the sensory level, the time series of noisy contrast values is convolved with a temporal filter to give the time-varying internal response, r:

$$r(t) = c(t) * h(t), \tag{A.1}$$

where $c(t)$ is the temporally varying stimulus and $h(t)$ denotes the impulse response of the temporal filter. This impulse response is specified by a model of early temporal vision of Watson (1986), chap. 6:

$$h(t) = \xi[h_1(t) - \zeta h_2(t)], \tag{A.2}$$

with

$$h_i(t) = u(t)[\tau_i(n_i - 1)!]^{-1} \left(\frac{t}{\tau_i}\right)^{n_i-1} e^{-\frac{t}{\tau_i}}, \tag{A.3}$$

for $i = 1, 2$ and $u(t)$ is the unit step function. Each component of Eq. (A.2), $h_i(t)$, represents a cascade of n leaky integrators with a time-constant τ. The overall system function is the difference between

the two components (scaled by the gain ξ). The strength of the second filter, ζ, determines the transience of the system. A transient system ($\zeta = 1$) responds to a step input with a burst which then gradually decreases down to baseline (Fig. 2 in the main text, top pathway). A sustained system ($\zeta = 0$) maintains its elevated response to a step input (Fig. 2, bottom pathway).

The complete system function is specified by six parameters, four of which were kept fixed: $\tau_1 = 6$, $\tau_2 = 8$, $n_1 = 9$, $n_2 = 10$. The level of transience was varied systematically: $\zeta = \{0, 0.25, 0.50, 0.75, 1\}$. The values of the constants were chosen as they approximately match those that appear to account well for a variety of basic phenomena in human early temporal vision (e.g. threshold vs. duration functions; pulse-pair sensitivity Watson, 1986, chap. 6). This choice reflects a prior belief that the same mechanisms that serve human performance in such basic detection tasks form the primary input into the downstream saccadic decision unit.

Appendix B. Formulation of an optimal decision variable

The following treatment is directly analogous to the analysis presented in Gold & Shadlen (2001); the reader is referred to this paper for a more detailed overview of the concepts involved. Consider a pair of temporal filters – one filter for each pattern location – with identical temporal properties, giving rise to two time-varying internal responses: $r_{cw}(t)$ and $r_{aw}(t)$, for the clockwise and anti-clockwise positions respectively. These responses are subject to internal noise, a 0-mean, Gaussian process. Let $r'(t)$ be the noise perturbed internal response to a pattern.

Now consider a single, momentary sample from the two noisy internal responses: r'_{cw} and r'_{aw}. The task of the decision mechanism is to decide, given the two noisy samples, which response is more likely to correspond to the target. Suppose the decision mechanism knew what the expected internal response was in the presence of a target or a non-target for this specific temporal sample, as well as the associated variance. Let $\mu_T > \mu_{NT}$ be the mean internal responses to the target and non-target, respectively, with equal variances: $\sigma_T^2 = \sigma_{NT}^2 = \sigma_{int}^2$ (internal noise here may be viewed as a combination of the intrinsic noise within the observer and a filtered version of the external noise inserted in the stimulus).

For each internal response sample, the decision mechanism may now compute the probability that sample was elicited by the target or by the non-target. The likelihood-ratio gives the relative evidence in favour of two hypotheses (is this pattern a target or a non-target?), given an internal response sample from that particular location:

$$LR_{T,NT|r'_{cw}} = \frac{f(r'_{cw}|\mu_T, \ \sigma_{int})}{f(r'_{cw}|\mu_{NT}, \ \sigma_{int})} \tag{B.1a}$$

$$LR_{T,NT|r'_{aw}} = \frac{f(r'_{aw}|\mu_T, \sigma_{int})}{f(r'_{aw}|\mu_{NT}, \sigma_{int})}. \tag{B.1b}$$

In Eq. (B1), $f(x|\mu, \ \sigma)$ is the Gaussian probability density with mean μ and standard deviation σ. Writing out f and taking the natural logarithm of both sides gives:

$$\ln LR_{T,NT|r'_{cw}} = -\frac{1}{2\sigma_{int}^2}\left[2r'_{cw}(\mu_{NT} - \mu_T) + \mu_T^2 - \mu_{NT}^2\right] \tag{B.2a}$$

$$\ln LR_{T,NT|r'_{aw}} = -\frac{1}{2\sigma_{int}^2}\left[2r'_{aw}(\mu_{NT} - \mu_T) + \mu_T^2 - \mu_{NT}^2\right]. \tag{B.2b}$$

The combined log-likelihood ratio across the two locations is now:

$$\ln LR_{T,NT|r'_{cw},r'_{aw}} = \ln LR_{T,NT|r'_{cw}} + \ln LR_{T,NT|r'_{aw}}. \tag{B.3}$$

Substituting Eq. (B2) into the right-hand side of Eq. (B.3) yields, after some elementary algebra:

$$\ln LR_{T,NT|r'_{cw},r'_{aw}} = (r'_{cw} - r'_{aw})\frac{(\mu_T - \mu_{NT})}{\sigma_{int}^2}. \tag{B.4}$$

Eq. (B.4) is the critical result: it says that the log-likelihood ratio of the target being in the clockwise or anti-clockwise pattern location, given a pair of internal response samples, is simply proportional to the difference between the two samples:

$$\ln LR_{T,NT|r'_{cw},r'_{aw}} \propto (r'_{cw} - r'_{aw}). \tag{B.5}$$

In other words, the decision mechanism need not know about the momentary expected internal response distributions, as the parameters of these distributions only enter Eq. (B.4) as a scaling term. Instead, the mechanism only needs to keep track of the difference in the internal responses. A positive response difference indicates that the internal response from the clockwise location is larger than that from the anti-clockwise location. A negative response indicates the reverse. Given the equal variance assumption, if one pattern triggers a larger internal response than the other, it is more likely to be the target. As argued by Gold & Shadlen (2001), even when the equal variance assumption is violated Eq. (B.5) still provides a useful approximation to the optimal decision variable.

The analysis so far dealt only with a single point in time. One attractive feature of log-likelihood ratios is that they may be added as more information becomes available. For instance, as time passes the decision mechanism may collect more and more momentary internal response samples r'. The continuous-time definition of the noisy decision variable, D is then:

$$D(t) = \int_a^t [r_{cw}(x) - r_{aw}(x)] + \eta(0, \sigma_{int})dx. \tag{B.6}$$

In Eq. (B.6) η is the 0-mean Gaussian (white noise) process. The noise is independently and identically distributed over time. In the model fit to the data, the lower integration limit, a depends on when the magnitude of the internal sensory response exceeds some critical value λ.

Note that there is only one source of internal noise in the model, although internal noise almost certainly exists at every stage between input and output. All the different sources of noise (e.g. neural and sampling noise at the sensory level; motor noise at the level of the decision unit) are lumped together. As noted in the main text, the standard diffusion model contains between-trial variability in the drift rate as well as the starting point. Section 4.1 contains a discussion of how starting point variability may be mimicked by a gradual onset in the drift. Between-trial variability in the drift, *over and above the variability caused by the external noise in the stimulus*, could have been incorporated through noise in the gain of the filter. In the context of the present data set it is not clear whether the inclusion of additional noise sources would have resulted in a significant improvement in the goodness-of-fit, and—more importantly—alter the theoretical conclusions drawn from the simpler model used here.

At this point the model is over-determined: an increase in noise may be compensated for by an increase in the gain of the filter so that the signal-to-noise ratio at the output of the filter remains constant. It is customary to fix one of these parameters. In the model fits reported here σ_{int}^2 was set to 1. The time step of the discrete approximation was set at 5 ms.

Appendix C. Defective densities for correct and error responses

The model is a time-inhomogeneous diffusion process (Ditterich, 2006b; Smith, 2000). The mean drift is specified by the instantaneous response difference between the two patterns. As a result of the external noise in the stimulus, the exact pattern contrast values varied both across and within trials. Consequently, the mean drift varied across and within trials. Thus, for any given trial the

mean drift varies over time (i.e. is inhomogeneous) in a way that is entirely specified by the internal responses to the two patterns, which − in turn − depend entirely on the nature of the underlying temporal filter (Eqs. A.1, A.2 and A.3).

To capture a wide range of possible temporal variations in the response boundaries (from nearly constant to approximating a step change), a logistic function was used to model the value of the (positive) response boundary over time:

$$\theta(t) = \theta_0 \left[1 - \frac{1}{1 + \exp(-(t - m)/s)} \right], \tag{C.1}$$

where m and s correspond to the mid-point and scale of the logistic function, respectively. The upper starting point of the function is specified by parameter θ_0. A nearly constant boundary may be instantiated by a very large value of the scale s, so that the change in the boundary is extremely gradual. Alternatively, a very sharp deadline through a small value of s may occur at some late point in time, outside the range of the observed latencies by setting m to a large value (e.g. >1000 ms, as was done in the creation of Fig. 2). The deadline model proposed by Ludwig et al. (2005b) would correspond to a small scale coupled with a relatively early value of m (e.g. ~150 ms). The lower boundary, − $\theta(t)$, is simply a negative copy of the upper threshold.

For a given set of parameters, the model predictions were derived through simulation of 50,000 trials. For each of these simulated trials two noisy contrast sequences were generated in exactly the same manner as in Ludwig et al. (2005b, Experiment 1). Internal responses to the two patterns were corrupted by 0-mean, Gaussian noise and the response difference between clockwise and anti-clockwise responses was integrated over time as soon as one of the noise-perturbed responses exceeded the minimum response criterion, λ (as described by Eq. (B.6)).

For each simulated trial, the first boundary crossing determined the target of the model's response and the time of the boundary crossing was the decision time. The response was classified as correct or erroneous, depending on the (simulated) target position. The simulated latency distributions of correct and error decisions cannot be used directly to derive the probability density of the observed response latencies, which is necessary for maximum-likelihood estimation of the model parameters. As a result, the simulated distributions were approximated with a Gaussian kernel (Van Zandt, 2000).

Let $\mathbf{T} = \{T_1, T_2, \ldots, T_M\}$ be the ordered (from small to large) vector of simulated latencies of a particular response type (correct, error), containing M out of a total of N simulated responses. A Gaussian kernel estimator provides an estimate of the continuous probability density function underlying the simulated sample $\mathbf{T}$:

$$f(t) = \frac{1}{Nh} \sum_{i=1}^{M} -\frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2} \left(\frac{T_i - t}{h} \right)^2 \right], \tag{C.2}$$

where h sets the width of the Gaussian averaging window according to

$$h = \frac{.51}{N^{.2}} \min\left(\sqrt{\mathrm{var}(\mathbf{T})}, \frac{iqr(\mathbf{T})}{1.349} \right),$$

where iqr denotes the inter-quartile range of the simulated sample. To obtain the probability distributions of correct and error *response* times, the decision time densities are convolved with a Gaussian distribution of residual, non-decisional latencies:

$$g(srt) = \int f(t) w(srt - t | \mu_{nd}, \sigma_{nd}) dt, \tag{C.3}$$

where srt is the saccade latency and w is the Gaussian distributed non-decisional delay.

Appendix D. Model fitting

For a given set of parameters, the likelihood of each observed response was computed from the appropriate defective density:

$$L_i = y_i g_{correct}(srt_i) + (1 - y_i) g_{error}(srt_i), \tag{D.1}$$

with $y = 1$ and $y = 0$ for correct and error responses, respectively, and for $i = 1, \ldots, N_{obs}$ trials. The overall log-likelihood is then:

$$\ln L = \sum_{i=1}^{N_{obs}} \ln L_i. \tag{D.2}$$

An initial grid-search was used to find a set of starting parameters that produced reasonable predictions. Subsequently, a simplex optimisation routine (Nelder & Mead, 1965) was initialised with these parameters to minimise the deviance (-2 times the log-likelihood). Each iteration of the algorithm involved the simulation of 50,000 choices. The random number generator seed was set to the same value for each simulation. Without this step the likelihood would vary even when the model parameters are kept constant upon repeated simulation. Such variation impedes convergence of the optimisation routine.

References

Albrecht, D., Geisler, W., & Crane, A. (2003). Nonlinear properties of visual cortex neurons: Temporal dynamics, stimulus selectivity, neural performance. In L. Chalupa & J. Werner (Eds.), *The visual neurosciences* (pp. 747–764). Cambridge, MA: MIT Press.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review, 113*, 700–765.

Breitmeyer, B. (1975). Simple reaction time as a measure of the temporal response properties of transient and sustained channels. *Vision Research, 15*, 1411–1412.

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

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

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

Caspi, A., Beutter, B. R., & Eckstein, M. P. (2004). The time course of visual information accrual guiding eye movement decisions. *Proceedings of the National Academy of Sciences of the United States of America, 101*, 13086–13090.

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience, 11*, 693–702.

Cook, E. P., & Maunsell, J. H. R. (2004). Attentional modulation of motion integration of individual neurons in the middle temporal visual area. *Journal of Neuroscience, 24*, 7964–7977.

Ditterich, J. (2006a). Evidence for time-variant decision making. *European Journal of Neuroscience, 24*, 3628–3641.

Ditterich, J. (2006b). Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks, 19*, 981–1012.

Findlay, J., & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences, 22*, 661–674.

Fredericksen, R. E., & Hess, R. F. (1998). Estimating multiple temporal mechanisms in human vision. *Vision Research, 38*, 1023–1040.

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences, 5*, 10–16.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30*, 535–574.

Gorea, A., & Tyler, C. (1986). New look at Blochs law for contrast. *Journal of the Optical Society of America A, 3*, 52–61.

Green, D., & Swets, J. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

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

Hess, R., & Snowden, R. (1992). Temporal properties of human visual filters: Number, shapes and spatial covariation. *Vision Research, 32*, 47–59.

Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience, 25*, 10420–10436.

LaBerge, D. (1962). A recruitment theory of simple behavior. *Psychometrika, 27*, 375–396.

Laming, D. R. J. (1968). *Information theory of choice reaction time*. New York: Wiley.

Laming, D. (1969). Subjective probability in choice-reaction time experiments. *Journal of Mathematical Psychology, 6*, 81–120.

Legge, G. (1978). Sustained and transient mechanisms in human vision: Temporal and spatial properties. *Vision Research, 18*, 69–81.

Link, S., & Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika, 40*, 77–105.

Ludwig, C. J. H., Gilchrist, I. D., & McSorley, E. (2005a). The remote distractor effect in saccade programming: Channel interactions and lateral inhibition. *Vision Research, 45*, 1177–1190.

Ludwig, C. J. H., Gilchrist, I. D., McSorley, E., & Baddeley, R. J. (2005b). The temporal impulse response underlying saccadic decisions. *Journal of Neuroscience, 25*, 9907–9912.

Ludwig, C. J. H., Mildinhall, J., & Gilchrist, I. D. (2007). A population coding account for systematic variation in saccadic dead time. *Journal of Neurophysiology, 97*, 795–805.

Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex, 13*, 1257–1269.

Muller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (2001). Information conveyed by onset transients in responses of striate cortical neurons. *Journal of Neuroscience, 21*, 6978–6990.

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

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision, 5*, 376–404.

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Logan, G. D., Schall, J. D., & Palmeri, T. J. (2008). Modeling interactions between visually-responsive and movement-related neurons in fef during saccade visual search. *Journal of Vision, 8*, 1080.

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

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin and Review, 9*, 278–291.

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

Ratcliff, R., Hasegawa, Y., Hasegawa, R., Smith, P., & Segraves, M. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness discrimination task. *Journal of Neurophysiology, 97*, 1756–1774.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science, 9*, 347–356.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111*, 333–367.

Reddi, B. A. J., Asrress, K. N., & Carpenter, R. H. S. (2003). Accuracy, information, and response time in a saccadic decision task. *Journal of Neurophysiology, 90*, 3538–3546.

Reddi, B. A. J., & Carpenter, R. H. S. (2000). The influence of urgency on decision time. *Nature Neuroscience, 3*, 827–830.

Robinson, D. (1972). Eye movements evoked by collicular stimulation in the alert monkey. *Vision Research, 12*, 1795–1808.

Robinson, D., & Fuchs, A. (1969). Eye movements evoked by stimulation of frontal eye fields. *Journal of Neurophysiology, 32*, 637–648.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience, 22*, 9475–9489.

Smith, P. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review, 102*, 567–593.

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology, 44*, 408–463.

Smith, P. L., & Vickers, D. (1988). The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology, 32*, 135–168.

Thier, P., & Andersen, R. A. (1998). Electrical microstimulation distinguishes distinct saccade-related areas in the posterior parietal cortex. *Journal of Neurophysiology, 80*, 1713–1735.

Tolhurst, D. (1975a). Reaction times in the detection of gratings by human observers: A probabilistic mechanism. *Vision Research, 15*, 1143–1149.

Tolhurst, D. (1975b). Sustained and transient channels in human vision. *Vision Research, 15*, 1151–1155.

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

Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review, 7*, 424–465.

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics, 13*, 37–58.

Watson, A. (1986). *Temporal sensitivity. Handbook of perception and human performance* (Vol. 1). New York: Wiley-Interscience. pp. 1–43.

Yang, S. N., & McConkie, G. W. (2001). Eye movements during reading: A theory of saccade initiation times. *Vision Research, 41*, 3567–3585.