Behavioral/Systems/Cognitive

The Optimality of Sensory Processing during the Speed–Accuracy Tradeoff

Tiffany Ho,[1] **Scott Brown,**[3] **Leendert van Maanen,**[4] **Birte U. Forstmann,**[4] **Eric-Jan Wagenmakers,**[4] **and John T. Serences**[1,2]

[1]Department of Psychology and [2]Neurosciences Graduate Program, University of California, San Diego, La Jolla, California 92093, [3]School of Psychology, University of Newcastle, Callaghan, New South Wales 2308, Australia, and [4]Department of Psychology, University of Amsterdam, 1018 WS Amsterdam, The Netherlands

When people make decisions quickly, accuracy suffers. Traditionally, speed–accuracy tradeoffs (SATs) have been almost exclusively ascribed to changes in the amount of sensory evidence required to support a response ("response caution") and the neural correlates associated with the later stages of decision making (e.g., motor response generation and execution). Here, we investigated whether performance decrements under speed pressure also reflect suboptimal information processing in early sensory areas such as primary visual cortex (V1). Human subjects performed an orientation discrimination task while emphasizing either response speed or accuracy. A model of choice behavior revealed that the rate of sensory evidence accumulation was selectively modulated when subjects emphasized accuracy, but not speed, suggesting that changes in sensory processing also influence the SAT. We then used fMRI and a forward encoding model to derive orientation-selective tuning functions based on activation patterns in V1. When accuracy was emphasized, the extent to which orientation-selective tuning profiles exhibited a theoretically optimal gain pattern predicted both response accuracy and the rate of sensory evidence accumulation. However, these relationships were not observed when subjects emphasized speed. Collectively, our findings suggest that, in addition to lowered response thresholds, the performance decrements observed during speeded decision making may result from a failure to optimally process sensory signals.

Introduction

Fast decisions are typically more error prone, while precise decisions require more time, a phenomenon known as the speed–accuracy tradeoff (SAT) (Woodworth, 1899; Fitts, 1966; Wickelgren, 1977; Dickman and Meyer, 1988). Traditional models of the SAT hold that fast but premature responses occur when not enough sensory information has been accumulated to support an accurate judgment (i.e., response thresholds are too low) (Bogacz et al., 2006; Ratcliff and McKoon, 2008). On this "response threshold" account, the SAT is mediated by neural mechanisms of late-stage decision processes that immediately precede the initiation of motor responses (Forstmann et al., 2008, 2010; van Veen et al., 2008; Bogacz et al., 2010). A complementary— and largely untested—hypothesis is that speed pressure also influences the efficiency with which sensory evidence is accumulated during decision making ("sensory-readout" hypothesis). This is an important possibility given that the rate of

sensory evidence accumulation necessarily limits the efficacy of downstream decision making and motor control processes.

To investigate the influence of the SAT on sensory processing, we designed a perceptual decision making task that required human observers to discriminate between two orientated grating stimuli (Fig. 1) (see Materials and Methods) under either speed emphasis (SE) or accuracy emphasis (AE) conditions. Importantly, subjects had to discriminate a small rotational offset (5°) between the gratings. Previous psychophysical and neurophysiological studies have shown that the most informative neurons for supporting such fine discriminations are tuned away from the target feature (hereupon termed "off-target" neurons) (Fig. 2) (Regan and Beverley, 1985; Hol and Treue, 2001; Schoups et al., 2001; Purushothaman and Bradley, 2005; Butts and Goldman, 2006; Jazayeri and Movshon, 2006; Navalpakkam and Itti, 2007; Moore, 2008; Scolari and Serences, 2009). This theoretical framework thereby provides a benchmark pattern of optimal sensory gain against which we can compare gain observed under different SAT conditions.

We investigated how the SAT influenced information processing by fitting response time (RT) and accuracy data using two models of choice behavior: the linear ballistic accumulator (LBA) model (Brown and Heathcote, 2008) (Fig. 3) and an extension of the LBA, the single-trial linear ballistic accumulator (STLBA) (van Maanen et al., 2011). These models revealed an impact of task instruction on the amount of information required to initiate a decision ("response caution") and on the rate of sensory evidence accumulation (the "drift rate"); the later effect suggests

Figure 1. Behavioral paradigm. Subjects were presented with an oriented grating at full contrast flickering at 6 Hz (on for 83.33 ms, off for 83.33 ms, etc., for a total of 3 s). On each trial, the target orientation of the grating was selected from one of nine possible orientations (0, 20, 40, 60, 80, 100, 120, 140, 160°), plus or minus an offset randomly selected between 0 and 6°. On one-half of the trials, the same stimulus was presented on every flicker (match trials), but for the remaining trials (mismatch trials), the orientation of the grating was offset by 5° on alternating flickers (either clockwise or counterclockwise, counterbalanced across trials). On alternating runs, subjects had to emphasize either speed or accuracy in their responses. For additional details, see Materials and Methods.

that the SAT may affect sensory processing (Hübner et al., 2010; Vandekerckhove et al., 2011). We then used a forward encoding model (Brouwer and Heeger, 2009, 2011) (for review, see Naselaris et al., 2011; Serences and Saproo, 2012) to examine how feature-selective BOLD response profiles in primary visual cortex (V1) are associated with behavioral performance and with the rate of sensory evidence accumulation under different SAT conditions. Our results suggest that theoretically optimal response patterns in V1 are associated with more efficient sensory evidence accumulation—but only when accuracy is emphasized over speed.

Materials and Methods

Subjects. Sixteen subjects (11 females) were recruited from the University of California, San Diego (UCSD) (La Jolla, CA) community. All had normal or corrected-to-normal vision. Each subject gave written informed consent per Institutional Review Board requirements at UCSD and completed a single 1 h session in a climate- and noise-controlled subject room outside of the scanner and a single 1.5–2 h session in the scanner. Compensation for participation was $10/h for behavioral training and $20/h for scanning. Subjects received an additional reward for task compliance according to a point system described below (mean additional compensation, $6.64). Data from two subjects were discarded due to improper slice stack selection during fMRI scanning that resulted in no data being collected from large portions of primary visual cortex, the main area of interest in this study.

Stimuli and task. Visual stimuli were generated using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) implemented in MATLAB (version 7.1; MathWorks), presented at a frame rate of 60 Hz, and projected onto a screen at the entrance of the scanner bore that subjects viewed through a mirror. Button press responses were made on an fMRI-compatible response box using the index and middle fingers of the right hand.

Subjects were shown a centrally presented oriented grating (with a diameter of ~14°) at full contrast which flickered at 6 Hz (83.33 ms on, 83.33 ms blank interval). On each trial, the orientation of the grating was pseudorandomly selected with equal probability from one of nine possible orientations (0, 20, 40, 60, 80, 100, 120, 140, 160°) with a small amount of pseudorandom jitter added (up to ±6°, selected from a uniform distribution). On one-half of the trials, the same stimulus was presented at every "flicker" ("match" trials), but for the remaining trials ("mismatch" trials), the orientation of the grating was offset by 5° on every other flicker, with the rotational offset of the deviant grating (i.e., clockwise or counterclockwise) fixed on a given trial and counterbal-

anced across trials (Fig. 1). The subject's objective was to make a match/mismatch judgment by pressing one of two buttons held in the right hand. The order of match and mismatch trials was pseudorandomized and counterbalanced within each run. Subjects were allowed to make a response any time after the stimulus onset; the stimulus was present for 3 s, after which it was replaced with a white centrally presented fixation circle for 3.5 s. We omitted all trials in which a subject failed to give a response (<1%) or emitted a response faster than 200 ms (<1.05%). Since the second grating did not appear until 166.67 ms into the trial, we reasoned that responses quicker than 200 ms should be regarded as definite blind guesses. In total, a single run consisted of 50 trials (36 experimental trials and 14 null trials consisting of just a fixation circle) and lasted 336 s including an 11 s period of passive fixation at the end of each run. Across the 36 experimental trials, each of the 9 possible orientations was presented 4 times. Before each run, subjects were instructed by the experimenter to emphasize either response accuracy or speed. Subjects earned points based on their performance: +10 for correct responses on accuracy trials, −10 for incorrect responses on accuracy trials, +10 for correct responses on speed trials within the response deadline, +0 for incorrect responses on speed trials, −10 for any responses exceeding the response deadline. At the end of the experiment, subjects were paid an additional $1 for every 100 points earned during their performance while being scanned (rounded to the nearest dollar). During training in the laboratory, subjects were given trial-by-trial feedback, but feedback was delayed until the end of each run during the scanning session.

Response deadline on speeded trials. All participants were trained before the scan session for a minimum of 180 trials. During training, subjects received point feedback on a trial-by-trial basis according to the reward scheme outlined above. Participants practiced the task without any speed pressure until they felt comfortable and performed at ~90% accuracy. Subjects were then asked to repeat the task by responding as quickly as they could without guessing. The median of their RT distribution on this block was then set as their response deadline for both the subsequent speeded training blocks and the speeded blocks during the fMRI session.

LBA model. Behavioral data were modeled using the LBA, which is a simplified version of the ballistic accumulator model and the leaky competing accumulator model (Usher and McClelland, 2001; Brown and Heathcote, 2005, 2008). Figure 3 illustrates the LBA model schematically. On each trial, two racing accumulators begin with a random activation level (the "starting point") that is independently drawn from a uniform distribution on [0, A]. Activity in each accumulator increases linearly, and a response is triggered as soon as one accumulator crosses the response threshold. The predicted response time is the time taken to reach the threshold, plus a constant offset time ("nondecision time"). The rate at which activation increases in each accumulator is termed the drift rate for that accumulator. These drift rates are drawn from independent normal distributions for each accumulator (with the SD of these distributions being arbitrarily fixed at 1). The means of the normal distributions reflect the quality of the perceptual input. For example, a salient mismatch between the orientated gratings would lead to a large mean drift rate in the accumulator corresponding to a mismatch response (and vice versa). Hence, the LBA model estimates the mean of this drift rate distribution for each accumulator (match or mismatch).

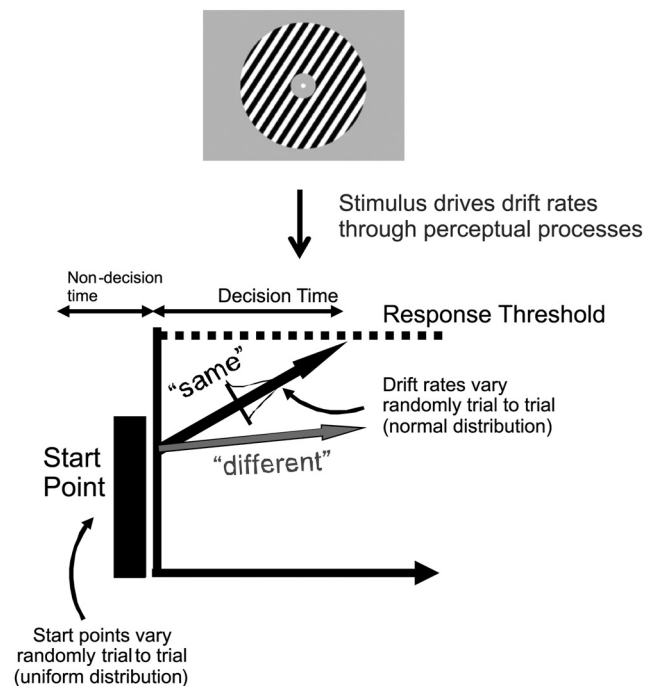The distance from the starting point to the response threshold is a measure of response caution as this distance quantifies the average amount of evidence that needs to be accumulated before a response is initiated. Changes in response caution are usually assumed to originate from adjusting the response threshold; however, adjusting the response threshold is mathematically equivalent to adjusting the starting point;
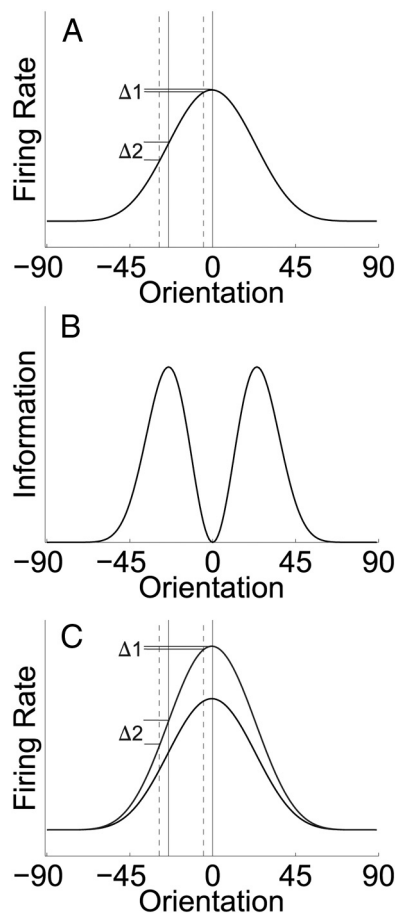
Figure 2. Model of optimality for fine perceptual discrimination tasks. ***A***, On-target neurons (0° offset on the abscissa axis) exhibit small changes in firing rates (Δ1) in response to two similar stimuli (denoted by the vertical solid line and the vertical dashed line). Off-target neurons, however, undergo larger changes in firing rates in response to similarly spaced stimuli (as denoted by Δ2). ***B***, The information available for supporting a fine discrimination—here defined as the slope of the tuning curve of a neuron at the target angle—plotted for neurons selective to all possible orientations (with 0° on the abscissa axis indicating the target). As suggested by ***A***, off-target neurons are potentially more informative when performing a fine discrimination. ***C***, Increasing the gain of the informative off-target neurons serves to further increase their sensitivity to small changes in the stimulus feature, thereby leading to improved discrimination performance.

Figure 3. Schematic of the LBA model. Schematic illustration of the LBA model as applied to the orientation discrimination task. The stimulus grating (top) provides information to two racing accumulators; the first accumulator to reach threshold determines the response, and thus the decision processing time. One accumulator corresponds to each possible response (match shown in black or mismatch shown in gray). The drift rates are assumed to be determined by the stimulus properties. The drift rate for each accumulator varies trial-by-trial based on a normal distribution, and the LBA reports the average rates from this distribution. Response caution determines how much sensory evidence needs to be accumulated before a response is made and is captured by the distance between the response threshold and the starting point. The final response time is the time taken for the first accumulator to reach threshold plus a constant offset (the nondecision time).

therefore, we chose to fix the height of the uniform distribution (*A*) from which the starting points were drawn (although the starting points nevertheless vary trial to trial) (Ratcliff, 1978; Ratcliff and Rouder, 1998; Forstmann et al., 2010; Wolfe and Van Wert, 2010; van Maanen et al., 2011). As a result, we hereon use the response threshold to represent response caution since the maximum of the start point distribution is fixed across the SE and AE conditions.

LBA model analyses. The parameters of the LBA model were estimated using the method of maximum likelihood. Likelihood was optimized using simplex searches (Nelder and Mead, 1965). Initial parameter values for searches were generated two ways: using heuristic calculations based on the data, and using start points determined from the end points of searches for simpler, nested models [full details of these methods and extensive discussion of alternative approaches are provided in the study by Donkin et al. (2011a)]. We fit the match and mismatch trials simultaneously, fixing all parameters between these two trial types to be constant except for the drift rate (which is presumably determined by the stimulus). Different drift rates were estimated for the accumulator corresponding to a mismatch response on trials with a mismatch stimulus (i.e., "correct" drift rate) and on trials with a match stimulus (i.e., "incorrect" drift rate; see Table 2).

Similarly, different drift rates were estimated for the accumulator corresponding to a match response on trials with a match stimulus (correct drift rate) and on trials with a mismatch stimulus (incorrect drift rate; see Table 2). Each different design for constraining model parameters across task conditions was fit separately to each individual subject's data. One subject, however, only made one incorrect response among the AE mismatch trials, thereby providing little constraint on the model estimate for that condition. The parameter estimates for that subject were therefore set to the group average for that condition. The overall grouped BIC value provided very strong support for the design that allowed three parameters to vary between SE and AE conditions [response threshold (*b*), drift rate (*v*), and nondecision time (t_0)]. To quantify that support, we approximated posterior model probabilities based on BIC assuming a fixed effect for subjects (Burnham and Anderson, 2002), which showed this design to be $>10^{10}$ times more likely than the next best design (see Results).

STLBA. Response times and accuracy vary on each trial due to environmental changes and/or internal noise in a subject's cognitive state. It is therefore important not only to map overall mean behavioral patterns with parameters that quantify relevant cognitive processes (as can be done with the standard LBA) but also to link estimates of these parameters and BOLD responses on a trial-by-trial basis.

In the standard LBA model (as in other decision making models) (Ratcliff, 1985; Ratcliff and Rouder, 1998), drift rates are normally distributed across trials, with independent distributions for each respective accumulator. This assumption of normally distributed drift rates implies that drift rates that are close to the mean of the distribution are more likely than values from the tails of the distribution. In addition, the uniform distribution [0, *A*] restricts the range of starting points for each accumulator. These considerations yield the following

maximum-likelihood estimates for a single-trial drift rate (d_i) and a single-trial starting point (a_i) given a trial with response time (t_i):

$$\hat{d}_i = \begin{cases} \dfrac{b-A}{t_i - t_0} & \text{if } t_i \leq \dfrac{b-A}{v} + t_0 \\[2ex] v & \text{if } \dfrac{b-A}{v} + t_0 < t_i < \dfrac{b}{v} + t_0 \\[2ex] \dfrac{b}{t_i - t_0} & \text{if } t_i \geq \dfrac{b}{v} + t_0 \end{cases}$$

$$\hat{a}_i = \begin{cases} A & \text{if } t_i \leq \dfrac{b-A}{v} + t_0 \\[2ex] b - (t_i - t_0) \cdot v & \text{if } \dfrac{b-A}{v} + t_0 < t_i < \dfrac{b}{v} + t_0 \\[2ex] 0 & \text{if } t_i \geq \dfrac{b}{v} + t_0, \end{cases} \quad (1)$$

where b, v, A, and t_0 are the parameters estimated using the standard LBA that correspond to the response threshold (b), the drift rate (v), the height of the distribution of starting points (A), and the nondecision time (t_0), respectively. Note that the assumed independence between estimated parameters that is found in the standard LBA model is no longer preserved with the STLBA. Nevertheless, parameter recovery studies show that the STLBA can explain much of the variance in the true parameter values [van Maanen et al. (2011), see the text surrounding their Fig. 3].

As in the main LBA analysis, we computed single-trial estimates of drift rate based on a model where response threshold (b), drift rate (v), and the nondecision time (t_0) were free to vary between SE and AE trials, whereas the height of the uniform distribution of starting points (A) was fixed (for exact values, see Table 2). Constraining the model in other reasonable ways (e.g., fixing the nondecision time parameter) yielded qualitatively similar results. Note also that the single-trial estimates for the starting point here are mathematically equivalent to single-trial estimates of the response threshold since what is actually being calculated is the relative distance between the two.

fMRI data acquisition and analysis. All scanning was performed on a General Electric MR750 3T scanner equipped with an eight-channel head coil at the W. M. Keck Center for Functional MRI on the main campus at UCSD. Anatomical images were acquired using a FSPGR T1-weighted sequence that yielded images with a $1 \times 1 \times 1$ mm resolution. Whole-brain echo-planar functional images (EPIs) were acquired in 28 (for eight of the subjects) or 26 (for the remaining subjects) oblique slices (TR, 1500 ms; TE, 30 ms; flip angle, 90°; image matrix, 64×64; field of view, 192 mm; slice thickness, 3 mm; 0 mm gap).

Data analysis was performed using BrainVoyager QX (version 1.91; Brain Innovation) and custom time series analysis routines written in MATLAB (version 7.11.0.584; MathWorks). Data from the main experiment were collected in 8 or 10 runs per subject (i.e., either 4 or 5 runs per response instruction type, respectively). EPI images were slice-time corrected, motion-corrected (both within and between scans), and high-pass filtered (3 cycles/run) to remove low-frequency temporal components from the time series. The time series from each voxel in each observer was then z-transformed on a run-by-run basis to normalize the mean response intensity across time to zero. This normalization was done to correct for differences in mean signal intensity across voxels (e.g., differences related to the composition of a voxel or by its distance from the coil elements). We then estimated the magnitude of the BOLD response on each trial by shifting the time series from each voxel by four 1.5 s TRs (6 s) to account for the temporal lag in the hemodynamic response function, and then extracting data from the two consecutive 1.5 s TRs that correspond to the duration of each 3 s trial (Kamitani and Tong, 2005; Serences and Boynton, 2007a,b; Serences et al., 2009). The two data points extracted from these two consecutive TRs were then averaged together to compute a single estimate of the response in each V1 voxel on each trial. These trial-by-trial estimates of the BOLD response amplitude were subsequently used as inputs to the forward encoding model (see below, Estimating feature-selective BOLD response profiles using a forward encoding model).

Functional localizer scans. Each subject participated in two runs of an independent functional localizer scan to identify voxels within primary visual cortex that were responsive to the spatial position occupied by the oriented grating stimulus used in the primary experiment. The localizer stimulus was comprised of a full-contrast counter-phase modulated (8 Hz) checkerboard that exactly matched the size of the oriented grating stimulus used in the main task. On each trial, the checkerboard stimulus was presented continuously for 10 s, and the contrast of the checkerboard was reduced by 30% for a single video frame at six pseudorandomly selected time points. Subjects were instructed to make a button press response with their right index finger every time they detected a contrast decrement. Each 10 s trial was then followed by 10 s of passive fixation. Visually responsive regions of primary visual cortex were identified using a general linear model (GLM) with a single regressor that was constructed by convolving a boxcar model of the stimulus sequence with a standard model of the hemodynamic response function (a difference-of-two gamma function model implemented in Brain Voyager; time to peak of positive response, 5 s; time to peak of negative response, 15 s; ratio of positive and negative responses, 6; positive and negative response dispersion, 1). Voxels were retained for analysis in the main experimental task if they passed a false discovery rate corrected single-voxel threshold of $p < 0.05$.

Retinotopic mapping. A meridian mapping procedure consisting of a checkerboard wedge flickering at 8 Hz and subtending 60° of polar angle was used to identify V1 (Engel et al., 1994; Sereno et al., 1995). Subjects were instructed to fixate on the center of the screen and to passively view the peripheral stimulus. The data collected during these scans were then projected onto a computationally inflated representation of each subject's gray/white matter boundary. V1 in each hemisphere was then manually defined according to the representations of the upper and lower vertical meridian following standard practices (Wandell et al., 2007).

Estimating feature-selective BOLD response profiles using a forward encoding model. The goal of encoding models is to adopt an a priori assumption about the important features that can be distinguished using hemodynamic signals within an ROI, and then to use this set of features (or basis functions) to predict observed patterns of BOLD responses (Gourtzelidis et al., 2005; Thirion et al., 2006; Dumoulin and Wandell, 2008; Kay et al., 2008; Brouwer and Heeger, 2009, 2011; Kay and Gallant, 2009; Naselaris et al., 2009) (for review, see Naselaris et al., 2011; Serences and Saproo, 2012). Here, we assumed that the BOLD response in a given V1 voxel represents the pooled activity across a large population of orientation-selective neurons, and that the distribution of neural tuning preference is biased within a given voxel due to large-scale feature maps (Freeman et al., 2011) or to random anisotropies in the distribution of orientation-selective columns within each voxel (Kamitani and Tong, 2005; Swisher et al., 2010). Thus, the BOLD response measured from many of the voxels in V1 exhibit a robust orientation preference (Haynes and Rees, 2005; Kamitani and Tong, 2005; Serences et al., 2009; Brouwer and Heeger, 2011; Freeman et al., 2011).

To estimate orientation-selective response profiles based on activation patterns in V1, we first separated the data from the 8–10 scanning runs obtained for each subject into train and test sets using a "leave two out" cross-validation scheme (i.e., all but one SE and one AE run were used as a training set, and the held-out SE and AE runs were used as a test set). By holding one AE and one SE run out for use as a test set, we ensured that the training set had an equal number of trials of each type. For each run in the training set, we then computed the mean response evoked by each of the nine orientations, separately for each voxel. The mean responses were then sorted based on stimulus orientation and run (i.e., mean response to orientation 1 was first, then orientation 2, . . . , orientation 9). Thus, each training set had 54 observations for subjects who underwent 8 runs in the scanner (6 runs in training set $\times$ 9 orientations), and 72 observations for subjects that underwent 10 runs in the scanner (8 runs in the training set $\times$ 9 orientations). Note that, as described below, data in the test set were not averaged across trials, and a unique channel response function was estimated for every trial.

Adopting the terminology of Brouwer and Heeger (2009, 2011), let m be the number of voxels in a given visual area, n_1 be the number of observations in the training set (either 54 or 72), n_2 be the number of
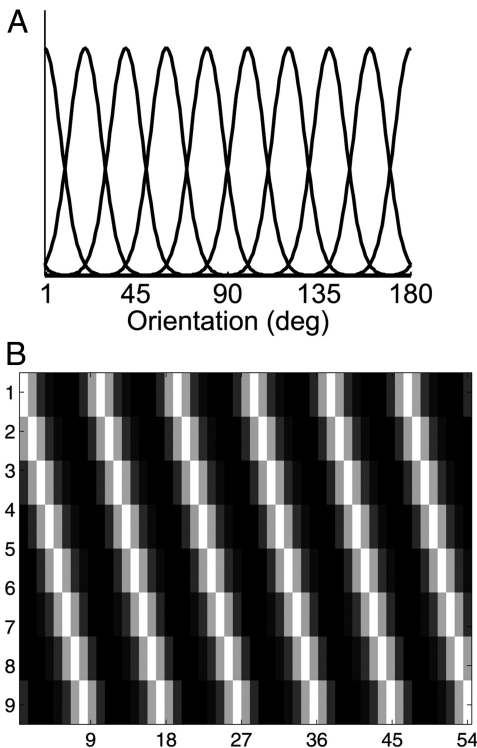
Figure 4. Schematic of the forward encoding model. **A**, Basis set comprised of nine half-sinusoidal functions raised to the sixth power: the functions are evenly distributed across orientation space. **B**, Depiction of the design matrix C_1. There is one column corresponding to each observation in the training data matrix, and each row represents the response profile of one of the nine half-sinusoidal basis functions shown in **A**. This matrix is used to calculate the weight matrix (W) that estimates the magnitude of the response in each voxel in each of the nine hypothetical orientation channels. This weight matrix can then be used to infer the response profile across all nine channels on each trial in the test data set. For more details, see Equations 2–4 (see Materials and Methods, Estimating feature-selective BOLD response profiles using a forward encoding model).

trials in the test set, and k be the number of hypothetical orientation channels. Let B_1 ($m \times n_1$ matrix) be the training data set, and B_2 ($m \times n_2$ matrix) be the test data set. Under the assumption that the observed BOLD signal is a weighted sum of underlying orientation-selective neural responses, we generated a matrix of hypothetical channel outputs (C_1, $k \times n_1$) composed of nine half-sinusoidal functions raised to the sixth power as a basis set (Fig. 4). The training data in B_1 were then mapped onto the matrix of channel outputs (C_1) by the weight matrix (W, $m \times k$) that was estimated using a GLM of the following form:

$$B_1 = WC_1, \qquad (2)$$

where the ordinary least-squares estimate of W is computed as follows:

$$\hat{W} = B_1 C_1^T (C_1 C_1^T)^{-1}. \qquad (3)$$

The channel responses $C_2(k \times n_2)$ were then estimated based on the test data (B_2) using the weights estimated in Equation 3 as follows:

$$\hat{C}_2 = (\hat{W}^T \hat{W})^{-1} \hat{W}^T B_2. \qquad (4)$$

The first steps in this sequence (Eqs. 2, 3) are similar to a traditional univariate GLM in that each voxel is assigned a weight for each feature in the model (in this case, one weight for each hypothetical orientation channel). Equation 4 then implements a multivariate computation because the channel responses estimated on each trial (in C_2) are constrained by the estimated weights assigned to each voxel and by the vector of responses observed across all voxels on that trial in the test set. Thus, one key feature of this approach is that a set of estimated channel responses can be obtained on a trial-by-trial basis so long as the number of

voxels is greater than the number of channels. If there are fewer voxels than channels, then unique channel response estimates cannot be derived as the number of variables being estimated exceeds the number of available measurements. This ability to estimate the orientation-selective tuning profile on each trial is exploited when comparing channel responses on correct and incorrect trials and when correlating channel responses with accuracy and drift rates on a trial-by-trial basis (see Results).

The shape of the basis functions used in C_1 has a large impact on the resulting channel response estimates. In the present experiment, we used half-sinusoidal functions that were raised to the sixth power to approximate the shape of single-unit tuning functions in V1, where the $1/\sqrt{2}$ half-bandwidth of orientation tuned cells is ~20° (although there is a considerable amount of variability in bandwidth) (Schiller et al., 1976; Swindale, 1998; Ringach et al., 2002a,b; Gur et al., 2005). Given that the half-sinusoids were raised to the sixth power, a minimum of seven linearly independent functions was required to adequately cover orientation space (Freeman and Adelson, 1991); however, since we presented nine unique orientations in the experiment, we used a set of nine evenly distributed functions. The use of more than the required seven basis functions is not problematic so long as the number of functions does not exceed the number of measured stimulus values, in which case the matrix C_1 would become rank deficient. While we selected the bandwidth of the basis functions based on physiology studies, all results that we report are robust to reasonable variations in this value (i.e., raising the half-sinusoids to the fifth to eighth power, all of which are reasonable choices based on the documented variability of single-unit bandwidths). Note that, since the magnitude of the channel responses is scaled by the amplitude of the basis functions (which was set to 1 here), the units along the y-axes of all data plots are in arbitrary units. Importantly, however, scaling the basis functions to some other common value would not change the differential response between conditions.

Using this modeling approach, the center position of each function in the basis set can be systematically shifted across orientation space to estimate the response in a channel centered at any arbitrary orientation (as long as the channels remain linearly independent) (Freeman and Adelson, 1991). While this method of shifting the center of each channel across orientation space could in principle be used to generate channel response profiles with a resolution of 1° (or even smaller), we opted to reconstruct the response functions in 5° steps as no additional insights were gained by estimating the responses at a higher resolution. After generating a channel response function on each trial in 5° steps across orientation space, each function was circularly shifted to a common stimulus-centered reference frame, and these recentered response functions were averaged across left and right V1 and across all trials of a like kind. Thus, by convention, the 0° point along the x-axis in all plots refers to the stimulus that evoked the response profile. Finally, since all channel response functions were found to be symmetrical about their center point, we averaged data from corresponding offsets on either side of the 0° point (i.e., data were averaged from the channels offset by +5 and −5° from the stimulus, +10 and −10°, and so forth) to produce the reported orientation tuning functions. Note that, in the process of collapsing across channels centered on both positive and negative offsets from 0°, we necessarily collapsed across mismatch trials in which there was either a clockwise or a counterclockwise offset between sequentially presented gratings within a trial. However, sorting the data by the rotational offset of the deviant grating had no qualitative impact on any of our results, presumably because the two gratings were flickering back and forth on sequential presentations over the course of the 3 s trial (Fig. 1) and because there was a random jitter of up to ±6° introduced on each trial (see task description above), which was on the same order as the offset between sequential gratings on mismatch trials (±5°).

Bootstrapping/randomization procedure for evaluating statistical significance. Because the basis functions used to estimate channel responses overlapped—thus violating the independence assumption of traditional statistical tests—we estimated statistical significance using a nonparametric bootstrapping/randomization procedure. Note that this bootstrapping/randomization procedure was used for all comparisons related to BOLD tuning functions (see Figs. 6, 7, AE vs SE, correct AE vs incorrect

Table 1. Behavioral accuracy and response times on correct trials during the fMRI experiment for each condition and for the main effect of response emphasis (speed vs accuracy) and the main effect of trial type (match vs mismatch)

Condition	Accuracy	Response time (ms)
Speed (match)	71.37 ± 2.71	947.7 ± 64.2
Speed (mismatch)	61.08 ± 3.08	896.8 ± 49.4
Accuracy (match)	90.00 ± 1.78	1697.8 ± 131.2
Accuracy (mismatch)	82.05 ± 3.28	1443.9 ± 92.9
Speed (match plus mismatch)	66.23 ± 2.24	922.4 ± 40.1
Accuracy (match plus mismatch)	86.02 ± 2.00	1570.8 ± 82.6
Match (speed plus accuracy)	80.69 ± 2.40	1322.8 ± 101.7
Mismatch (speed plus accuracy)	71.57 ± 2.99	1170.4 ± 73.7

For more details, see Results. Shown are means ± 1SEM across subjects.

Table 2. Average LBA parameter estimates for the best BIC model on mismatch and match trials

LBA parameters	Speed (mismatch)	Accuracy (mismatch)	Speed (match)	Accuracy (match)
Starting point (A)	2.31	2.31	2.31	2.31
Nondecision time (t_0)	0.41	0.42	0.41	0.42
Response threshold (b)	2.75	4.43	2.75	4.43
Drift rate for accumulator corresponding to correct response (v_c)	2.67	3.00	2.75	2.58
Drift rate for accumulator corresponding to incorrect response (v_{ic})	1.71	0.69	1.19	−0.07
Differential drift rate ($v_c − v_{ic}$)	0.96	2.31	1.56	2.65

All parameters were fixed to be constant across match and mismatch trials except for the drift rates, which are assumed to be directly related to the quality of the stimuli (for more details, see LBA results). Note that the response threshold parameter is equivalent to what we refer to as response caution in the main text, as the starting point is fixed across conditions (see Materials and Methods). Note that, in the main text, we refer to correct accumulators as the accumulator corresponding to the match response on trials in which the stimulus gratings match, and also the accumulator corresponding to the mismatch response on trials in which the gratings mismatch (and the converse is true for incorrect accumulators). Critically, there was a greater difference in the rate of sensory evidence accumulation (drift rates) between correct and incorrect accumulators on accuracy emphasis trials compared with speed emphasis trials (last row in the table).

AE, correct SE vs incorrect SE, the interaction between AE vs SE based on accuracy, AE logistic regression β weights vs SE logistic regression β weights, and single-trial correlations between AE responses and drift rate vs single-trial correlations between SE responses and drift rate). First, a series of standard paired t tests was performed to determine which points along the two tuning curves differed significantly (using a threshold of $p < 0.05$ for each individual t test). We then generated a new data set by randomly selecting 14 participants with replacement and then reassigning the condition label associated with data from each participant with a probability of 0.5. A series of paired t tests was performed on the resampled and randomized data set using the same procedure applied to the observed data. This resampling plus randomization procedure was then iterated 10,000 times to determine the probability of obtaining the pattern of significant differences obtained using the intact data set under the null hypothesis that the two conditions are equivalent (i.e., interchangeable). The reported p values in the main text thus reflect the proportion of times we observed a pattern of significant t tests in the resampled data that matched the pattern obtained in the observed data. Note that the behavioral data were evaluated using conventional parametric statistical techniques.

Results

Response time and accuracy results

Trials on which RTs were faster than 200 ms were discarded in all subsequent analyses (including model fitting procedures described below) (see Materials and Methods). Two-way repeated-measures ANOVA with factors for response-emphasis (speed vs accuracy emphasis, or SE and AE trials, respectively) and trial type (match vs mismatch) was used to assess accuracy and RT data collected during the scanning session (for a summary of the group data, see Table 1). The task instructions produced a strong SAT effect: participants responded faster on SE trials compared with AE trials ($F_{(1,13)} = 39.168; p < 0.001$; Table 1), and there was a corresponding drop in accuracy on SE trials ($F_{(1,13)} = 71.975$; $p < 0.001$; Table 1).

On average, subjects gave a match response 55% of the time, which is significantly greater than chance (one-sample t test, $t_{(13)} = 2.49$, $p = 0.03$). In addition, RTs were slower and accuracy slightly better on match trials compared with mismatch trials ($F_{(1,13)} = 13.26$, $p < 0.01$; $F_{(1,13)} = 5.4$, $p < 0.05$; Table 1), which is consistent with the bias to respond match over mismatch and commensurate with the well known propensity for making confirmatory responses (Clark and Chase, 1972). There was an interaction between response-emphasis and trial type for RTs ($F_{(1,13)} = 12.6$; $p < 0.01$; Table 1), with selectively long RTs on match AE trials. However, there was no interaction between response-emphasis and trial type for accuracy rates ($F_{(1,13)} = 0.61$; $p = 0.45$; Table 1).

LBA model results

Accuracy rates and RTs might be lower on SE trials compared with AE trials due to differences in response caution and/or in the

rate at which sensory evidence is accumulated. Therefore, we used a mathematical model of decision making (Fig. 3) to investigate how emphasizing either speed or accuracy influenced the rate of sensory evidence accumulation (as captured by the drift rate parameter) and response caution (as captured by the distance between the starting point and the response threshold) (see Materials and Methods). Given that the neuronal mechanisms thought to support fine orientation discriminations are reasonably well characterized (Fig. 2) (see below, Predicting feature-selective response patterns in V1 on mismatch trials), we focused our analyses on mismatch trials (data from match and mismatch trials were nonetheless fit simultaneously) (for more details, see Materials and Methods, LBA model analyses). Eight different versions of the LBA model of Brown and Heathcote (2008) were fit by allowing all combinations of three different parameters (drift rate, response caution, and nondecision time) to either vary freely across SE and AE conditions or to be fixed across those conditions, while keeping the maximum starting point always fixed across AE and SE conditions (Fig. 3) (for more details, see Materials and Methods, LBA model and LBA model analyses). We then used the Bayesian information criterion (BIC) to select the most parsimonious of the eight models, which is a commonly used measure that evaluates the tradeoff between model complexity and goodness of fit (Schwarz, 1978; Raftery, 1995). The model yielding the best BIC was the one that estimated different values for the parameters corresponding to response caution, drift rate, and nondecision time on AE trials compared with SE trials (for a summary of the parameter fits to data averaged over all the subjects, see Table 2). Based on approximated posterior model probabilities (for more details, see Materials and Methods, LBA model analyses), this design was found to be $>10^{10}$ times more likely than the next best design. Individually, this design was also the modal result: the BIC values for 6 of 14 subjects preferred this design. Four subjects had best-BIC designs that included an effect on drift rate but not response threshold, while three had best-BIC designs that included an effect on threshold but not drift rate. Only one subject had a best-BIC design that included no effect at all of the experimental manipulation.

Figure 5 shows the fit of this best-BIC model to the cumulative response time distributions from the data. This figure estimates the distributions using quantiles plotted against response probability. These plots are also known as "defective cumulative distribution plots" and are a standard method for evaluating the
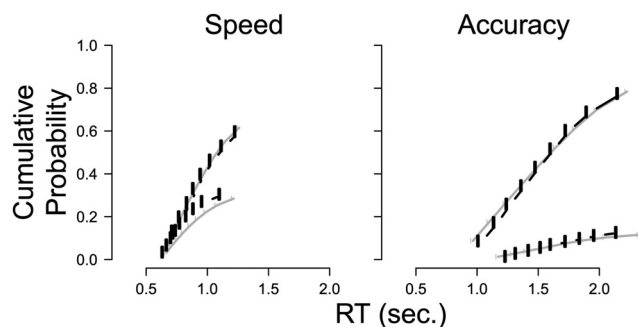
Figure 5. LBA model fits. Cumulative response time distributions estimated from the data using quantiles (black circles and lines) and predicted by the LBA model (gray dots and lines). Data are shown separately for the speed emphasis condition (left panel) and the accuracy emphasis condition (right panel). In each panel, the upper lines and symbols show quantile estimates for correct responses, and the lower set are for incorrect responses. The data quantiles and model predictions were generated separately for each individual participant and then averaged. The height of the graphs shows response probability. Nine quantile estimates are shown in each condition, corresponding to 10, 20, . . . , 90%.

quality of fit for response time models, as they provide a much more rigorous test than histograms [for introductions to this method and related discussion, see the study by Ratcliff and Tuerlinckx (2002) or the study by Donkin et al. (2011a)]. The model fits the data quite well, matching the probability (as indicated by the height on the graph) of each response type in each condition accurately. The latency of each part of the response time distribution (abscissa axis) is also accurately captured by the model. For example, in the SE and AE conditions, the median observed correct RT differs from the median LBA predicted value by <25 ms.

In any choice task, it is possible that participants occasionally make random guesses that are independent of the available stimulus information. This is especially a concern in the SE condition in which error rates were relatively high. However, since the decision model fits the response time distributions from both conditions very well (Fig. 5, left panel), we assume that simple random guessing is not a plausible explanation for observed differences in parameters between the SE and AE conditions. Nevertheless, to avoid having the model results overly biased by contaminant processes such as guessing, we incorporated a mixture process with the assumption that each response had a 98% probability of arising from the LBA choice process, and a 2% probability of arising from a guessing process with random responses and uniform RT over the observed range [for details, see the studies by Ratcliff and Tuerlinckx (2002) and Donkin et al. (2011a)]. With this built-in assumption, the decision model fit the response time distributions from both conditions very well (Fig. 5), consistent with the hypothesis that participants were making informed decisions on the vast majority of trials.

Consistent with most SAT studies, we observed a difference in response caution (Table 2) (Ratcliff, 1985; Ratcliff and Rouder, 1998; Voss et al., 2004; Palmer et al., 2005). Moreover, we observed a larger difference in the rate of evidence accumulation associated with correct and incorrect accumulators on AE trials compared with correct and incorrect accumulators on SE trials ($F_{(1,13)} = 18.27, p < 0.005$, repeated measures two-way ANOVA, with no main effect of stimulus type nor interaction between response type and stimulus type, $F_{(1,13)} = 2.82$ and $F_{(1,13)} = 0.95$, respectively; $p > 0.10$ for both). In the LBA, high accuracy occurs when the accumulator corresponding to the correct response for the stimulus gathers evidence more quickly than the accumulator

corresponding to the incorrect response. The larger difference in drift rates between correct and incorrect accumulators on AE trials therefore suggests that sensory information about the correct response is being selectively accumulated at a higher rate when subjects make decisions in the absence of speed pressure. Such selectivity represents a departure from the typical assumption used by mathematical psychologists that the rate of sensory evidence accumulation is fixed across AE and SE conditions [for extensive discussion, see Ratcliff and Rouder (1998)], as well as the typical assumption that response caution is the only cognitive process involved in the SAT. However, others have also observed evidence for a change in drift rates with task demands (Vandekerckhove et al., 2011), and we speculate that the effect may be even more apparent in the present task because subjects were engaged in a difficult perceptual discrimination in which the quality of sensory representations critically determined behavioral performance [for a more theoretical treatment, see also the study by Hübner et al. (2010)]. Finally, the small observed differences in the time taken for nondecision processing between SE and AE conditions (Table 2) are sometimes observed as a consequence of task instructions, but these differences are not usually of interest when the main purpose of the manipulation is to influence decision processing (Voss et al., 2004; Starns and Ratcliff, 2010).

In general, the parameter estimates from the LBA model have been shown to be in agreement with the corresponding parameters in the Ratcliff diffusion drift model (Donkin et al., 2011b). Nevertheless, to demonstrate that our modeling results are not specific to our choice of model and fitting procedures, we also fit our behavioral data using the Diffusion Model Analysis Toolbox (DMAT) implemented in MATLAB (Vandekerckhove and Tuerlinckx, 2007, 2008). We used DMAT to fit the same eight models tested in our LBA analysis (i.e., all possible combinations of drift rate, response threshold, and nondecision time varying or staying fixed across AE and SE conditions while keeping all other parameters fixed). Group BIC values for each model design were calculated in the same manner as those computed for the LBA models (see Materials and Methods, LBA model analyses). Consistent with the results of the LBA model, the diffusion model design with the best BIC was the one that estimated different values for the parameters corresponding to drift rate and response threshold on AE versus SE trials (AE drift rates were larger than SE drift rates, $t_{(13)} = 4.93, p < 0.01$; and AE response thresholds were higher than SE response thresholds, $t_{(13)} = 5.94, p < 0.01$). We then approximated posterior model probabilities using the group BIC value across all subjects for each model design. The model in which only drift rate and response threshold varied yielded the greatest posterior probability (close to 1, while all other posterior probabilities were close to 0).

Predicting feature-selective response patterns in V1 on mismatch trials

We next sought to establish a relationship between feature-selective BOLD responses in early visual areas and behavior. In situations that require discriminating between two highly similar stimuli (as in the present experiment in which orientations on mismatch trials were offset by only 5°), neurons tuned to off-target orientations provide the most information about the presence of mismatching orientations (Fig. 2) (Regan and Beverley, 1985; Hol and Treue, 2001; Schoups et al., 2001; Purushothaman and Bradley, 2005; Butts and Goldman, 2006; Jazayeri and Movshon, 2006; Navalpakkam and Itti, 2007; Scolari and Serences, 2009, 2010). Hence, we focused our analyses on mismatch
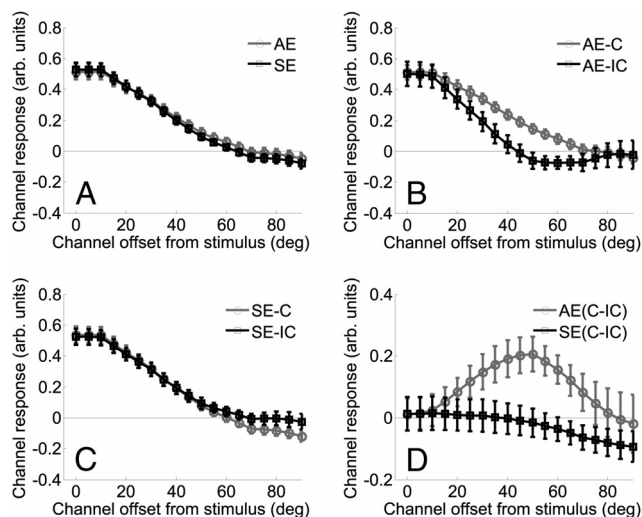
Figure 6. *A*, Orientation-selective tuning functions for mismatch AE and SE trials. Channel TFs for the AE data (gray) and SE data (black) are plotted as a function of the offset of each orientation channel from the stimulus presented on each trial (which is by convention always set to 0°; note that channel responses were collapsed across clockwise and counterclockwise offsets as the functions were symmetrical about the 0° point) (see Materials and Methods). *B*, Tuning functions for mismatch AE trials. A comparison of tuning functions for correct (gray) versus incorrect (black) AE trials. The activation profile for correct responses was significantly different from incorrect responses, with greater activation in channels offset by 35–60° on correct AE trials. All error bars are ±1SEM. *C*, Tuning functions for mismatch SE trials. A comparison of tuning functions for correct (gray) versus incorrect (black) SE trials. There was no significant difference between the two tuning functions. All error bars are ±1SEM. *D*, Interaction between mismatch AE trials and mismatch SE trials on correct and incorrect trials. A comparison of the difference between correct and incorrect trials for AE (gray) and SE (black) tuning functions. There was a significant difference between these two tuning functions, with larger off-target modulations on AE trials. All error bars are ±1SEM across subjects. All *p* values were calculated based on a nonparametric randomization procedure (for more details, see Materials and Methods).

trials in which the activation of off-target neurons is predicted to support such decisions. Given the relatively large difference in drift rates associated with correct and incorrect accumulators on AE mismatch trials (Table 2), we predicted that correct mismatch AE trials should be associated with more activation in off-target neural populations compared with incorrect mismatch AE trials. The difference between the drift rates associated with the correct and incorrect accumulators on SE mismatch trials, however, is much smaller (Table 2). We would therefore expect a small difference between off-target activation on correct SE mismatch trials compared with incorrect mismatch SE trials.

Assessing feature-selective tuning functions in V1 on mismatch trials

We used fMRI and a forward encoding model of BOLD responses (Brouwer and Heeger, 2009, 2011) (for review, see Naselaris et al., 2011; Serences and Saproo, 2012) to estimate how the SAT modulates orientation-selective response profiles in V1 (see Materials and Methods, Estimating feature-selective BOLD response profiles using a forward encoding model). On mismatch trials, we first compared the BOLD-based orientation tuning functions (TFs) associated with AE trials with those associated with SE trials (Fig. 6*A*) and found no significant difference in the shape of the curves (*p* = 0.91; this and all subsequently reported *p* values associated with TFs were estimated using a nonparametric randomization procedure due to the non-independence of adjacent data points; see Materials and Methods). However, when examining only the AE mismatch trials, we found a

significant interaction between channel offset and behavioral accuracy (*p* < 0.01; Fig. 6*B*). In particular, responses in channels tuned ~25–65° away from the target orientation showed larger responses on correct trials compared with incorrect trials. The observation of more activation in these off-target channels on correct trials is consistent with our predictions, as these neural populations should better signal small changes in orientation. In turn, more gain in off-target populations should increase the quality of the information being sent to downstream decision mechanisms and thus increase the probability of a correct response (Fig. 2*C*). In contrast, no differences were observed between channel responses associated with correct and incorrect SE trials (*p* = 0.90; Fig. 6*C*), and the difference between off-target channel responses on correct and incorrect AE trials was significantly larger than the difference on correct and incorrect SE trials (*p* < 0.01; Fig. 6*D*). This interaction is consistent with the relatively large difference in drift rates associated with correct versus incorrect accumulators on AE trials compared with SE trials (Table 2).

As shown in Figure 6*B*, we observed more activation in off-target channels on correct trials compared with incorrect trials in the AE condition. To further test the relationship between the magnitude of off-target responses and behavior, we performed a between-subject correlation between the change in drift rate and the change in off-target activation on correct and incorrect AE trials (where our measure of off-channel activation was the area between the TFs associated with correct and incorrect responses across channels tuned 25–65° from the target orientation) (Fig. 6*B*). Across subjects, larger differences between correct and incorrect accumulator drift rates were positively correlated with larger differences in off-target activation on correct compared with incorrect AE trials (Fig. 7*A*; $R^2 = 0.36$; $t_{(12)} = 2.59$; $p < 0.025$). This relationship was still observed even when the total area between the TFs associated with correct and incorrect AE trials (i.e., from 0 to 90°) was correlated with the differential drift rates ($R^2 = 0.30$; $t_{(12)} = 2.24$; $p < 0.05$), demonstrating that the positive correlation did not strongly depend on the exact points along the TFs that were entered into the analysis. This between-subjects relationship between BOLD and behavior suggests that individual differences in the degree of off-target activation in V1—and by inference, individual differences in the amount of information encoded about the orientation offset of mismatched stimuli—predicts the speed of evidence accumulation during decision making when subjects are not under speed pressure.

The correlation analysis presented in Figure 7*A* establishes a subject-by-subject relationship between off-target responses in V1 and the rate of sensory evidence accumulation. To further explore this relationship on a within-subject basis, we next used logistic regression to map fluctuations in the magnitude of the response in each orientation channel to accuracy on a trial-by-trial basis. A positive fit coefficient (β coefficient) indicates that higher activation in a given channel predicts a higher probability of a correct behavioral response; negative β coefficients indicate an inverse relationship between BOLD activation levels and behavioral performance. On AE mismatch trials, larger responses in channels tuned to the target (0° offset) were associated with a higher probability of incorrect responses, whereas larger responses in channels tuned ~40–60° from the target were associated with a higher probability of a correct response (Fig. 7*B*). In contrast, the β coefficients on SE trials fluctuated around zero. This pattern gave rise to a significant crossover interaction between the AE and SE β coefficient curves (*p* = 0.021; Fig. 7*B*). As with the increased off-target activation on correct AE trials (Fig. 6*B*) and the corresponding relationship with the rate of sensory
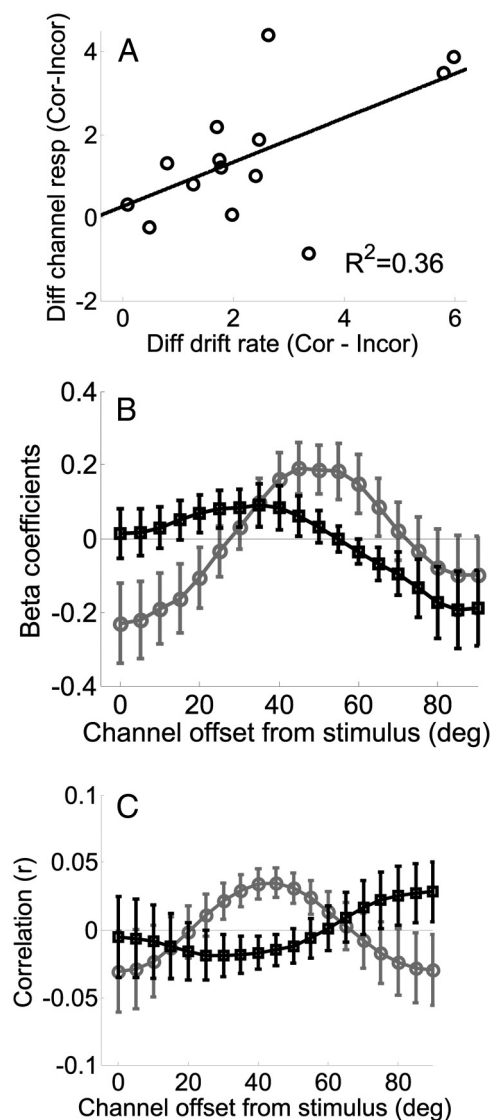
Figure 7. *A*, Between-subject correlation between differences in off-target channel responses and differences in drift rates. The difference in the drift rates associated with correct and incorrect accumulators is positively correlated with the difference in off-target activation on correct and incorrect trials. *B*, Logistic regression results. β coefficients computed using a logistic regression to relate BOLD channel responses and behavioral performance (correct/incorrect) on a trial-by-trial basis. Positive coefficients indicate that larger BOLD signals correspond to correct responses, while negative coefficients indicate that larger BOLD signals correspond to incorrect responses. On AE trials (gray), larger signals in channels that were offset by 40–60° from the stimulus predicted a higher probability of correct responses. The fit coefficients on SE trials (black) clustered around zero, except in the 0° channel, giving rise to a significant crossover interaction. All error bars are ±1 SEM across subjects. *C*, Correlations between STLBA drift rates with channel responses on correct mismatch trials. Correlation coefficients (*r* values) between STLBA drift rates and BOLD channel responses for the AE data (gray) and SE data (black) as a function of channel orientation. On each trial, we estimated both the channel response and the drift rate and then computed a correlation coefficient across trials per subject. The *r* values plotted here are averaged across subjects. The STLBA parameters were estimated based on a model in which response threshold and drift rate were allowed to vary between task conditions. All error bars are ±1 SEM across subjects. All *p* values were calculated based on a nonparametric randomization procedure (for more details, see Materials and Methods).

evidence accumulation on a between-subject basis (Fig. 7*A*), this trial-by-trial coupling between the magnitude of off-target channel responses and behavioral performance suggests that perceptual decisions are tightly coupled to activation levels across informative off-target sensory neurons, but only when subjects emphasize accuracy over speed.

Correlating STLBA estimates with off-target activation levels

Given the data presented thus far that off-target activation levels in V1 predict behavior on AE trials, we would also predict a positive correlation between trial-by-trial estimates of the rate of sensory evidence accumulation and the magnitude of the BOLD response in off-target channels. To evaluate this relationship, we correlated trial-by-trial estimates of off-target channel responses and trial-by-trial estimates of drift rates derived from the STLBA model on a within-subject basis.

As in the standard LBA model described above, we let the rate of sensory evidence accumulation (drift rate), response caution, and nondecision time vary freely across AE and SE conditions (and as in the standard LBA model, the height of the starting point distribution was fixed). Next, we estimated both channel responses and single-trial drift rates on each correct mismatch trial and then computed a correlation between these metrics across all trials for each subject. We observed larger correlation coefficients on AE compared with SE trials in channels tuned 30–55° away from the target. This pattern produced a significant crossover interaction between task instruction and correlation coefficient (*p* = 0.01; Fig. 7*C*) and suggests that larger off-target responses selectively predict higher rates of evidence accumulation on AE trials. This finding converges with the prior analyses of both channel response amplitude (Fig. 6*B,D*), between-subject correlation (Fig. 7*A*), and within-subject logistic regression (Fig. 7*B*), and is consistent with the idea that responses in informative sensory neurons are strongly coupled with behavioral performance, but only in the absence of speed pressure. However, this analysis more directly links trial-by-trial fluctuations in off-target channel responses with the rate of sensory evidence accumulation during decision making.

Note that the correlations shown in Figure 7*C* were expected to be small because both measures (model parameters and BOLD responses) are extremely variable when estimated on a trial-by-trial basis. Nevertheless, even though they were small in magnitude, the observed correlations were robust to reasonable changes in assumptions about how model parameters were constrained across conditions. For example, the same general pattern was observed using a variant of the STLBA model in which the nondecision time was fixed across trials. These results were also specific to trial-by-trial estimates of the model parameter corresponding to the rate of sensory evidence accumulation: correlating V1 channel responses to raw response times or to trial-by-trial estimates of the parameter corresponding to response caution did not yield robust correlations. The selectivity of the correlations presented in Figure 7*C* thus illustrates the explanatory power of the rate of sensory evidence accumulation on the SAT data and further supports the relationship between optimal response patterns in V1 and decision making when subjects emphasize accuracy over speed.

Discussion

When subjects emphasized accuracy, higher off-target activation levels predicted larger differential rates of sensory evidence accumulation (Fig. 7*A*), logistic regression revealed a trial-by-trial relationship between behavioral accuracy and BOLD activation levels in off-target orientation channels (Fig. 7*B*), and a model that provides trial-by-trial estimates of the latent cognitive processes involved in perceptual decision making (van Maanen et al., 2011) revealed a correlation between activation levels in off-target channels and the rate of sensory evidence accumulation (Fig. 7*C*).

The observation that off-target activation levels consistently predict behavioral performance on AE trials suggests that decision mechanisms can selectively pool inputs from the most informative sensory neurons (Purushothaman and Bradley, 2005; Law and Gold,

2009). However, this reliance on informative off-target channels during decision making only appears to happen on AE trials, as fluctuations in off-target responses do not predict behavior under speed pressure. This observation leads to an interesting prediction: given the low overall accuracy under speed pressure, we might have expected that off-target activation levels on SE trials more closely match off-target activation levels on incorrect AE trials (Fig. 6, compare *A*, *B*). Contrary to this prediction, we instead observed that tuning functions in the SE condition more closely resemble tuning functions on correct AE trials. This suggests that poor performance on SE trials is not related to low overall signal in off-target channels per se, but instead is caused by a failure to rely on informative populations of sensory neurons in an optimal manner during decision making. Although further investigation is clearly warranted, this apparent failure to rely on informative off-target neural responses on speeded trials may reflect a heuristic that enables a quick but imprecise readout of sensory information when response speed is at a premium.

One interpretation of the relationship between behavior and off-target modulations on AE trials holds that top-down attentional signals originating in frontal and parietal cortex differentially bias activation levels in off-target channels on a trial-by-trial basis. This type of attentional-feedback account is consistent by many theories of attentional control (for review, see Desimone and Duncan, 1995; Kastner and Ungerleider, 2000; Corbetta and Shulman, 2002; Serences and Yantis, 2006; Yantis, 2008; Noudoost et al., 2010) as well as recent evidence that the frontal operculum plays a causal role in governing attentional modulations in visual cortex and concomitant changes in performance across observers (Higo et al., 2011), and that subregions of frontal cortex mediate perceptual decisions (Heekeren et al., 2004, 2008; de Lafuente and Romo, 2005, 2006; Gold and Shadlen, 2007; Ho et al., 2009; Hernández et al., 2010; Kayser et al., 2010; Lemus et al., 2010; Purcell et al., 2010). However, since we did not directly manipulate attention in this study, it is difficult to dissociate sources of variability in V1 that are due to fluctuations in top–down biasing signals as opposed to sources of variability that are local to visual cortex. Future studies could more critically examine this issue by pairing a SAT task with either a valid or a neutral attention cue to determine whether speed pressure selectively impairs a subjects' ability to use prior information to appropriately bias population response profiles in visual cortex.

In addition to suboptimal usage of sensory information during decision making, it is likely that performance under speed pressure in our task is further limited by other neural mechanisms that operate outside of primary visual cortex. Several studies have found increased activation in the striatum when speeded responses are emphasized (Forstmann et al., 2008, 2010; van Veen et al., 2008), consistent with a response threshold account that only motor and frontal areas are involved in mediating the SAT (Ratcliff, 1985; Ratcliff and Rouder, 1998; Forstmann et al., 2008, 2010; Ivanoff et al., 2008; van Veen et al., 2008; Wenzlaff et al., 2011). In contrast, our findings provide support for the sensory-readout account, which posits that perceptual performance under speed pressure is also limited by how efficiently sensory information is integrated during decision making.

The forward encoding model that was used to estimate responses in different orientation channels is a proxy for the actual neural activity in the underlying populations of sensory neurons. This leads to an inevitable loss of resolution as a single V1 voxel contains many neural populations, and the bandwidth of V1 neurons can be highly variable. Therefore, it is difficult to pinpoint the exact orientation offset at which off-target BOLD modula-

tions would peak given a perfectly optimal modulation of underlying neural responses. However, the observation of increased responses starting in channels tuned 25–30° from the target is reasonable given the known tuning function properties of cells in V1 (Schiller et al., 1976; Ringach et al., 2002a,b; Gur et al., 2005). More generally, the robust relationship between off-channel activation levels and behavior supports the functional importance of the observed modulations, and is consistent with established models of optimal gain during fine discriminations (Fig. 2).

Generating channel tuning functions also depends critically on the ability of fMRI to reliably measure orientation-selective responses in primary visual cortex. In V1, it is likely that these feature-selective response biases depend to a large degree on relatively coarse maps of orientation space that unfold across the cortical surface (Leventhal, 1983; Schall et al., 1986; Sasaki et al., 2006; Mannion et al., 2009; Zhang et al., 2010; Freeman et al., 2011). For instance, there is a radial orientation bias in V1 (Sasaki et al., 2006; Zhang et al., 2010; Freeman et al., 2011). Thus, neurons with spatial receptive fields in (say) the upper right visual field tend to respond more to oblique orientations around 45°, and so on. Given the robust retinotopic organization of V1, this radial bias would generate an orderly representation of orientation across patches of cortex that represent each visual quadrant (Sasaki et al., 2006; Zhang et al., 2010; Freeman et al., 2011). In addition to this coarse orientation map across V1, voxel-level orientation selectivity may also reflect contributions from random anisotropies in the distribution of orientation-selective columns within a voxel (Haynes and Rees, 2005; Kamitani and Tong, 2005; Swisher et al., 2010) (for a useful graphical illustration, see Boynton, 2005). Thus, there is growing evidence that the combination of BOLD fMRI and encoding models can be used to index feature-selective responses arising from neural signals at both coarse and fine spatial scales.

Despite this link, we do not claim that orientation-selective response functions are solely related to neural spiking activity, as the BOLD signal is modulated by many sources including synaptic input from both local and distant inputs, tuned local field potentials, and even responses in astrocytes (Heeger et al., 2000; Logothetis et al., 2001; Buxton, 2002; Heeger and Ress, 2002; Logothetis and Wandell, 2004; Schummers et al., 2008; Sirotin and Das, 2009; Kleinschmidt and Muller, 2010; Das and Sirotin, 2011; Handwerker and Bandettini, 2011a,b; Jia et al., 2011). However, given that neurons in early sensory areas like V1 are massively interconnected (Douglas and Martin, 2007), changes in the BOLD signal related to synaptic activity should be highly correlated with changes in local spiking activity. Despite these caveats, the robust predictive relationship between off-target channel modulations and behavior strongly supports the functional significance of these indirect BOLD assays of neuronal activation.

The instruction-dependent change in the reliance of decision mechanisms on off-target channels in V1 is consistent with other recent studies of perceptual decision making. For instance, Kahnt et al. (2011) found that training-related improvements in performance on a difficult perceptual discrimination task could be explained by a model in which sensory information is read out more effectively, thereby improving the representations of the decision variables leading up to the ultimate choice (Purushothaman and Bradley, 2005; Law and Gold, 2008, 2009; Pestilli et al., 2011). Similarly, Rahnev et al. (2011) observed that manipulating prior expectation increased functional connectivity between posterior and frontal areas, consistent with an increase in the rate of sensory evidence transfer from earlier visual areas to putative deci-

sion mechanisms. Thus, the present results complement other recent studies that emphasize the importance of efficient sensory readout in perceptual decision making, and suggest that the optimality of readout breaks down under speed pressure.

References

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev 113:700–765.

Bogacz R, Wagenmakers EJ, Forstmann BU, Nieuwenhuis S (2010) The neural basis of the speed-accuracy tradeoff. Trends Neurosci 33:10–16.

Boynton GM (2005) Imaging orientation selectivity: decoding conscious perception in V1. Nat Neurosci 8:541–542.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433–436.

Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. J Neurosci 29:13992–14003.

Brouwer GJ, Heeger DJ (2011) Cross-orientation suppression in human visual cortex. J Neurophysiol 106:2108–2119.

Brown S, Heathcote A (2005) A ballistic model of choice response time. Psychol Rev 112:117–128.

Brown SD, Heathcote A (2008) The simplest complete model of choice response time: linear ballistic accumulation. Cogn Psychol 57:153–178.

Burnham KP, Anderson DR (2002) Model selection and multimodel inference: a practical information theoretic approach, Ed 2. New York: Springer.

Butts DA, Goldman MS (2006) Tuning curves, neuronal variability, and sensory coding. PLoS Biol 4:e92.

Buxton RB (2002) Introduction to functional magnetic resonance imaging: principles and techniques. Cambridge, UK: Cambridge UP.

Clark HH, Chase WG (1972) On the process of comparing sentences against pictures. Cogn Psychol 3:472–517.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci 3:201–215.

Das A, Sirotin YB (2011) What could underlie the trial-related signal? A response to the commentaries by Drs. Kleinschmidt and Muller, and Drs. Handwerker and Bandettini. Neuroimage 55:1413–1418.

de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. Nat Neurosci 8:1698–1703.

de Lafuente V, Romo R (2006) Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc Natl Acad Sci U S A 103:14266–14271.

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. Annu Rev Neurosci 18:193–222.

Dickman SJ, Meyer DE (1988) Impulsivity and speed-accuracy tradeoffs in information processing. J Pers Soc Psychol 54:274–290.

Donkin C, Brown SD, Heathcote A (2011a) Drawing conclusions from choice response time models: a tutorial. J Math Psychol 55:140–151.

Donkin C, Brown S, Heathcote A, Wagenmakers EJ (2011b) Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? Psychon Bull Rev 18:61–69.

Douglas RJ, Martin KA (2007) Recurrent neuronal circuits in the neocortex. Curr Biol 17:R496–R500.

Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. Neuroimage 39:647–660.

Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. Nature 369:525.

Fitts PM (1966) Cognitive aspects of information processing: III. Set for speed versus accuracy. J Exp Psychol 71:849–857.

Forstmann BU, Dutilh G, Brown S, Neumann J, von Cramon DY, Ridderinkhof KR, Wagenmakers EJ (2008) Striatum and pre-SMA facilitate decision-making under time pressure. Proc Natl Acad Sci U S A 105:17538–17542.

Forstmann BU, Anwander A, Schäfer A, Neumann J, Brown S, Wagenmakers EJ, Bogacz R, Turner R (2010) Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. Proc Natl Acad Sci U S A 107:15916–15920.

Freeman J, Brouwer GJ, Heeger DJ, Merriam EP (2011) Orientation decoding depends on maps, not columns. J Neurosci 31:4792–4804.

Freeman WT, Adelson EH (1991) The design and use of steerable filters. IEEE Trans Pattern Anal Mach Intell 13:891–906.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535–574.

Gourtzelidis P, Tzagarakis C, Lewis SM, Crowe DA, Auerbach E, Jerde TA, Uğurbil K, Georgopoulos AP (2005) Mental maze solving: directional fMRI tuning and population coding in the superior parietal lobule. Exp Brain Res 165:273–282.

Gur M, Kagan I, Snodderly DM (2005) Orientation and direction selectivity of neurons in V1 of alert monkeys: functional relationships and laminar distributions. Cereb Cortex 15:1207–1221.

Handwerker DA, Bandettini PA (2011a) Hemodynamic signals not predicted? Not so: a comment on Sirotin and Das (2009). Neuroimage 55:1409–1412.

Handwerker DA, Bandettini PA (2011b) Simple explanations before complex theories: alternative interpretations of Sirotin and Das' observations. Neuroimage 55:1419–1422.

Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat Neurosci 8:686–691.

Heeger DJ, Ress D (2002) What does fMRI tell us about neuronal activity? Nat Rev Neurosci 3:142–151.

Heeger DJ, Huk AC, Geisler WS, Albrecht DG (2000) Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? Nat Neurosci 3:631–633.

Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG (2004) A general mechanism for perceptual decision making in the human brain. Nature 431:859–862.

Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. Nat Rev Neurosci 9:467–479.

Hernández A, Nácher V, Luna R, Zainos A, Lemus L, Alvarez M, Vázquez Y, Camarillo L, Romo R (2010) Decoding a perceptual decision process across cortex. Neuron 66:300–314.

Higo T, Mars RB, Boorman ED, Buch ER, Rushworth MF (2011) Distributed and causal influence of frontal operculum in task control. Proc Natl Acad Sci U S A 108:4230–4235.

Ho TC, Brown S, Serences JT (2009) Domain general mechanisms of perceptual decision making in human cortex. J Neurosci 29:8675–8687.

Hol K, Treue S (2001) Different populations of neurons contribute to the detection and discrimination of visual motion. Vision Res 41:685–689.

Hübner R, Steinhauser M, Lehle C (2010) A dual-stage two-phase model of selective attention. Psychol Rev 117:759–784.

Ivanoff J, Branning P, Marois R (2008) fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. PLoS One 3:e2635.

Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. Nat Neurosci 9:690–696.

Jia X, Smith MA, Kohn A (2011) Stimulus selectivity and spatial coherence of gamma components of the local field potential. J Neurosci 31:9390–9403.

Kahnt T, Grueschow M, Speck O, Haynes JD (2011) Perceptual learning and decision-making in human medial frontal cortex. Neuron 70:549–559.

Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. Nat Neurosci 8:679–685.

Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. Annu Rev Neurosci 23:315–341.

Kay KN, Gallant JL (2009) I can see what you see. Nat Neurosci 12:245.

Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. Nature 452:352–355.

Kayser AS, Buchsbaum BR, Erickson DT, D'Esposito M (2010) The functional anatomy of a perceptual decision in the human brain. J Neurophysiol 103:1179–1194.

Kleinschmidt A, Müller NG (2010) The blind, the lame, and the poor signals of brain function—a comment on Sirotin and Das (2009). Neuroimage 50:622–625.

Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci 11:505–513.

Law CT, Gold JI (2009) Reinforcement learning can account for associative and perceptual learning on a visual-decision task. Nat Neurosci 12:655–663.

Lemus L, Hernández A, Luna R, Zainos A, Romo R (2010) Do sensory cortices process more than one sensory modality during perceptual judgments? Neuron 67:335–348.

Leventhal AG (1983) Relationship between preferred orientation and re-

ceptive field position of neurons in cat striate cortex. J Comp Neurol 220:476–483.

Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. Annu Rev Physiol 66:735–769.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412:150–157.

Mannion DJ, McDonald JS, Clifford CW (2009) Discrimination of the local orientation structure of spiral glass patterns early in human visual cortex. Neuroimage 46:511–515.

Moore BC (2008) Basic auditory processes involved in the analysis of speech sounds. Philos Trans R Soc Lond B Biol Sci 363:947–963.

Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL (2009) Bayesian reconstruction of natural images from human brain activity. Neuron 63:902–915.

Naselaris T, Kay KN, Nishimoto S, Gallant JL (2011) Encoding and decoding in fMRI. Neuroimage 56:400–410.

Navalpakkam V, Itti L (2007) Search goal tunes visual features optimally. Neuron 53:605–617.

Nelder JA, Mead R (1965) A simplex algorithm for function minimization. Comput J 7:308–313.

Noudoost B, Chang MH, Steinmetz NA, Moore T (2010) Top-down control of visual attention. Curr Opin Neurobiol 20:183–190.

Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. J Vis 5:376–404.

Pelli DG (1997) The Video Toolbox software for visual psychophysics: transforming numbers into movies. Spat Vis 10:437–442.

Pestilli F, Carrasco M, Heeger DJ, Gardner JL (2011) Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. Neuron 72:832–846.

Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ (2010) Neurally constrained modeling of perceptual decision making. Psychol Rev 117:1113–1143.

Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. Nat Neurosci 8:99–106.

Raftery AE (1995) Bayesian model selection in social research (with discussion). Sociol Methodol 25:111–193.

Rahnev D, Lau H, de Lange FP (2011) Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. J Neurosci 31:10741–10748.

Ratcliff R (1978) A theory of memory retrieval. Psychol Rev 85:59–108.

Ratcliff R (1985) Theoretical interpretations of the speed and accuracy of positive and negative responses. Psychol Rev 92:212–225.

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput 20:873–922.

Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. Psychol Sci 9:347–356.

Ratcliff R, Tuerlinckx F (2002) Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. Psychon Bull Rev 9:438–481.

Regan D, Beverley KI (1985) Postadaptation orientation discrimination. J Opt Soc Am 2:147–155.

Ringach DL, Bredfeldt CE, Shapley RM, Hawken MJ (2002a) Suppression of neural responses to nonoptimal stimuli correlates with tuning selectivity in macaque V1. J Neurophysiol 87:1018–1027.

Ringach DL, Shapley RM, Hawken MJ (2002b) Orientation selectivity in macaque V1: diversity and laminar dependence. J Neurosci 22:5639–5651.

Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RB (2006) The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. Neuron 51:661–670.

Schall JD, Vitek DJ, Leventhal AG (1986) Retinal constraints on orientation specificity in cat visual cortex. J Neurosci 6:823–836.

Schiller PH, Finlay BL, Volman SF (1976) Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specificity and ocular dominance. J Neurophysiol 39:1320–1333.

Schoups A, Vogels R, Qian N, Orban G (2001) Practising orientation identification improves orientation coding in V1 neurons. Nature 412:549–553.

Schummers J, Yu H, Sur M (2008) Tuned responses of astrocytes and their influence on hemodynamic signals in the visual cortex. Science 320:1638–1643.

Schwarz G (1978) Estimating the dimension of a model. Ann Stat 6:461–464.

Scolari M, Serences JT (2009) Adaptive allocation of attentional gain. J Neurosci 29:11933–11942.

Scolari M, Serences JT (2010) Basing perceptual decisions on the most informative sensory neurons. J Neurophysiol 104:2266–2273.

Serences JT, Boynton GM (2007a) Feature-based attentional modulations in the absence of direct visual stimulation. Neuron 55:301–312.

Serences JT, Boynton GM (2007b) The representation of behavioral choice for motion in human visual cortex. J Neurosci 27:12893–12899.

Serences JT, Saproo S (2012) Computational advances towards linking BOLD and behavior. Neuropsychologia 50:435–446.

Serences JT, Yantis S (2006) Selective visual attention and perceptual coherence. Trends Cogn Sci 10:38–45.

Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT (2009) Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. Neuroimage 44:223–231.

Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in human revealed by functional magnetic resonance imaging. Science 268:889–893.

Sirotin YB, Das A (2009) Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. Nature 457:475–479.

Starns JJ, Ratcliff R (2010) The effects of aging on the speed-accuracy compromise: boundary optimality in the diffusion model. Psychol Aging 25:377–390.

Swindale NV (1998) Orientation tuning curves: empirical description and estimation of parameters. Biol Cybern 78:45–56.

Swisher JD, Gatenby JC, Gore JC, Wolfe BA, Moon CH, Kim SG, Tong F (2010) Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. J Neurosci 30:325–330.

Thirion B, Duchesnay E, Hubbard E, Dubois J, Poline JB, Lebihan D, Dehaene S (2006) Inverse retinotopy: inferring the visual content of images from brain activation patterns. Neuroimage 33:1104–1116.

Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. Psychol Rev 108:550–592.

Vandekerckhove J, Tuerlinckx F (2007) Fitting the Ratcliff diffusion model to experimental data. Psychon Bull Rev 14:1011–1026.

Vandekerckhove J, Tuerlinckx F (2008) Diffusion model analysis with MATLAB: a DMAT primer. Behav Res Methods 40:61–72.

Vandekerckhove J, Tuerlinckx F, Lee MD (2011) Hierarchical diffusion models for two-choice response times. Psychol Methods 16:44–62.

van Maanen L, Brown SD, Eichele T, Wagenmakers EJ, Ho T, Serences J, Forstmann BU (2011) Neural correlates of trial-to-trial fluctuations in response caution. J Neurosci 31:17488–17495.

van Veen V, Krug MK Carter CS (2008) The neural and computational basis of controlled speed-accuracy tradeoff during task performance. J Cogn Neurosci 20:1952–1965.

Voss A, Rothermund K, Voss J (2004) Interpreting the parameters of the diffusion model: an empirical validation. Mem Cognit 32:1206–1220.

Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. Neuron 56:366–383.

Wenzlaff H, Bauer M, Maess B, Heekeren HR (2011) Neural characterization of the speed-accuracy tradeoff in a perceptual decision-making task. J Neurosci 31:1254–1266.

Wickelgren W (1977) Speed-accuracy tradeoff and information-processing dynamics. Acta Psychol 41:67–85.

Wolfe JM, Van Wert MJ (2010) Varying target prevalence reveals two, dissocial decision criteria in visual search. Curr Biol 20:121–124.

Woodworth RS (1899) The accuracy of voluntary movement. Psychol Rev 3:1–114.

Yantis S (2008) The neural basis of selective attention: cortical sources and targets of attentional modulation. Curr Dir Psychol Sci 17:86–90.

Zhang E, Huang M, Xiang X, Yan Y, Shen MY, Chen M, Xu X, Li W (2010) Anisotropic orientation selectivity of monkey V1 neurons in Cartesian and polar coordinates. Soc Neurosci Abstr 36:483.12.