UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Neural Mechanisms of Perceptual Decision Making

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy

in

Psychology

by

Tiffany Cheing Ho

Committee in charge:

Professor John Serences, Chair
Professor Thomas Albright
Professor Steven Hillyard
Professor David Huber
Professor Donald MacLeod

2012

The Dissertation of Tiffany Cheing Ho is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2012

TABLE OF CONTENTS

ACKNOWLEDGMENTS

Six years ago, I never would have imagined entering a doctoral program, much less completing one. Such an accomplishment would not have been possible without the support of the following individuals, to whom I'd like to dedicate this section.

To the greatest advisor in the world, John Serences, thank you for your wisdom, patience, teaching, and support for the past 5 years. Your passion for science has influenced me deeply and it is you that I look to and envision when aspiring to be a successful scientist.

To my collaborators across the ocean, Scott Brown, Leendert van Maanen, Birte Forstmann, and E.J. Wagenmakers, thank you for your mathematical expertise and for reigniting my fascination with quantitative modeling. I truly hope this will not be the last time we work together.

Thank you also to my dissertation committee members, Don MacLeod, Thomas Albright, Dave Huber, and Steven Hillyard for your advice and insight.

No words can adequately express how thankful I am for the Perception & Cognition Lab: Miranda Scolari, Anna Byers, Eddie Ester, Thomas Sprague, Javier Garcia, Sirawaj Itthipuripat, Mary Smith, and Sameer Saproo. You have seen me at my best and my worst and have accepted me despite the latter. You have sparked in me new ideas and new understanding about what it is we are studying and you have helped me do everything from carrying heavy boxes and being guinea pigs for my experiments, to calculating visual angle and providing emergency chocolate during times of abject stress. In short,

you have kept me sane, which is arguably the single most important factor in finishing a doctoral dissertation.

Thank you to my research assistants for putting up with my opaque code and for spending countless hours piloting my projects and collecting data (most of which will never see the light of day, but was necessary in order to design a better experiment, I swear!). Tony Abuyo, Dorothy Yen, Lilly Wu, Joshua Chen, Joanne Ku, Nafees Hamid, Donnie Davis, Shaheen Modir, Chris Reinert, Sherri Conklin, Deanna Shinsky, Haddas Elisha, and Laura Torres: thank you for being wonderful RAs and I wish you all the best.

Last but not least, thank you Mom, Ba, Popo, Christina, and Aleck for being my family and all that that entails. Without your emotional (and financial) support, completion of this dissertation would have not at all been possible.

Chapter 1, in full, is a reprint of the material as it appears in Domain General Mechanisms of Perceptual Decision Making in Human Cortex in *Journal of Neuroscience, 29 (27)*, 8675-8687. Ho, T.C., Brown, S.D., Serences, J.T. (2009). The dissertation author was the primary author of this paper.

Chapter 2, in full, is a reprint of the material as it appears in The Optimality of Sensory Processing During the Speed-Accuracy Tradeoff in *Journal of Neuroscience, 23(32)*, 7992-8003. Ho, T.C., Brown S.D., Van Maanen, L., Forstmann, B.U., Wagenmakers, E.J., Serences, J.T. (2012). The dissertation author was the primary author of this paper.

Chapter 3, in full, is a reprint of the material as it appears in Perceptual Consequences of Feature-based Attentional Enhancement and Suppression in *Journal of Vision*. Ho, T.C., Brown, S. Abuyo, N.A., Ku, E.J., Serences, J.T. (in press).

CURRICULUM VITAE

2006	Bachelor of Arts, University of California, Berkeley
2006-2007	Research Assistant, University of California, Berkeley
2006-2007	Research Assistant, University of California, San Francisco
2007	Student Researcher, Ernest Gallo Clinic and Research Center
2007-2008	Graduate Student, Department of Cognitive Sciences, University of California, Irvine
2009	Master of Arts, University of California, San Diego
2008-2012	Graduate Student, Department of Psychology, University of California, San Diego
2012	Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

Ho, T.C., Brown, S.D., Abuyo, N.T., Ku, E.H., Serences, J.T. (in press). Perceptual consequences of feature-based attentional enhancement and suppression. *Journal of Vision*.

Ho, T.C., Horn, N.A., Huynh, T., Kelava, L., Lansman, J.B. (2012). Evidence TRPV4 contributes to mechanosensitive ion channels in mouse skeletal muscle fibers. *Channels,* 6(4): http://dx.doi.org/10.4161/chan.20719.

Ho, T., Brown, S.D., Van Maanen, L., Forstmann, B.U., Wagenmakers, E.J., Serences, J.T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *Journal of Neuroscience,* 32(23): 7992-8003.

Van Maanen, L., Brown, S.D., Eichele, T., Wagenmakers, E.J., Ho, T., Serences, J., Forstmann, B.U. (2011) Neural correlates of trial-to-trial fluctuations in response caution, *Journal of Neuroscience*, 31(48):17488-17495.

Ho, T.C., Brown, S.D., Serences, J.T. (2009) Domain general mechanisms of perceptual decision making in human cortex. *Journal of Neuroscience*, 29(27):8675-8687.

Serences, J.T., Saproo, S., Scolari, M., Ho, T., Muftuler, T. (2009) Estimating the influence of attention on population response profiles. *Neuroimage*, 44(1):223-31.

ABSTRACT OF THE DISSERTATION

The Neural Mechanisms of Perceptual Decision Making

by

Tiffany Cheing Ho

Doctor of Philosophy in Psychology

University of California, San Diego, 2012

Professor John Serences, Chair

Perceptual decision making (PDM) involves choosing one option among several on the basis of sensory evidence and is a highly adaptive mechanism for organisms to successfully interact with their environments. Such a choice requires integrating and interpreting sensory information for the purpose of guiding subsequent behavior (e.g., seeing a ball move rightward and veering accordingly to catch it). Typical single-unit recording studies examining PDM utilize simple sensorimotor tasks (e.g., a macaque views a noisy array of dots moving in one of two possible directions and deploys a saccade in the chosen – and presumably, perceived – direction) in order to parse various aspects of PDM. With the aid of mathematical models, these experiments have found that the activity of individual neurons involved in motor response generation comprises perceptual decisions, and that PDM can be formalized as an accumulation of sensory evidence towards a particular choice (as represented by an increase in neuronal firing rate) until some threshold is reached. Explaining the mechanisms of PDM at the level of neural populations and linking ensemble patterns of neural activity to perception,

however, still remains unclear. With a combination of visual psychophysics, neuroimaging, and modeling, I present a set of studies that examines the neural correlates subserving PDM in human cortex (Experiment 1), clarifies the relationship between sensory representations in visual cortex and perceptual performance (Experiment 2), and tests the behavioral predictions derived from single-cell recordings (Experiment 3). These findings both challenge and confirm some of the previous neurophysiological work: Experiment 1 provides evidence of a neural mechanism of PDM not based purely on oculomotor regions, Experiment 2 shows that the optimality of activation patterns in visual cortex predicts task performance, and Experiment 3 illustrates that attentional manipulations influence perception in a manner consistent with the enhancement and suppression of distinct neural populations predicted from single-unit recordings. Furthermore, these studies demonstrate the utility of model-based cognitive neuroscience in quantifying psychological processes of interest for each individual and relating between-subject differences with corresponding brain measurements.

INTRODUCTION

The Neural Mechanisms of Perceptual Decision Making

At any given point in time, the state of the external environment is unknown and must be inferred based on noisy sensory input. Behavior is critically dependent on the ability to quickly and accurately decide among these possible states – a process known as perceptual decision making (PDM; von Helmhotz, 1925, Tenenbaum & Griffiths, 2001; Newsome et al., 1989; Salzman and Newsome, 1994; Gold and Shadlen, 2001; Shadlen and Newsome, 2001; Gold and Shadlen, 2007). For instance, deciding whether or not a predator is present in a shadowy corner will dictate an animal's subsequent action and survival. Various factors must be taken into account before committing to a decision and executing the appropriate behavioral response, including the quality of evidence for a particular choice and the subjective costs and benefits for each potential outcome. The decision variable (DV) represents the aggregate of these multiple sources of information (e.g., prior history, current emotional state, amount and quality of available evidence, value of each choice, etc) and the decision rule determines how and when the DV is interpreted to arrive at a particular choice (Gold and Shadlen, 2007). In the following sections, we will review the conceptual framework and quantitative models used to study PDM, summarize the neurophysiological and neuroimaging efforts investigating the neural correlates of the DV (as well as likely decision rules), and discuss the importance of informing neurally inspired PDM models with behavioral data.

Sequential sampling models of perceptual decision making

While the concept of the DV and the decision rule are abstractions, formal mathematical descriptions of these elements in the context of decision making have been proposed to better test and understand PDM. In particular, sequential sampling models have been used to explain response time (RT) and accuracy data on a variety of simple perceptual and cognitive tasks (for a review of these models, see Luce, 1986; Townsend and Ashby, 1983; Vickers, 1971; LaBerge, 1962; Laming, 1968; Vickers, 1979; Usher and McClelland, 2001; Ratcliff and Smith, 2004). The central assumption of these models is that information stored about a stimulus in our sensory systems is inherently noisy; thus, subjects must accumulate successive samples of this noisy stimulus representation until enough evidence towards a particular hypothesis regarding the stimulus is obtained. Under the assumption that distributions of evidence are normal, a model that accumulates samples of evidence in this manner is statistically equivalent to one that computes log-likelihood ratios (Laming, 1968). Mathematically, these models resemble sequential probability ratio tests (SPRT), which also based on likelihood and distributions of evidence that could only be estimated from thousands of observations of a stimulus (Smith and Ratcliff, 2004; Gold and Shadlen, 2001; 2002).

Sequential sampling models can be classified according to whether evidence is accumulated discretely or continuously, sampled in fixed or varying-sized steps, or accrued as single or separate totals (see Figure 1 and Ratcliff and Smith, 2004; Bogacz et al., 2006; Otter et al., 2008 for a comparison of these models). In *random walk models*, relative evidence is accumulated as a single total over time (Ratcliff, 1978; Ratcliff and Rouder, 1998; Ratcliff et al., 1999; Busemeyer and Townsend, 1992; 1993; Smith, 1995; see Figure 1, left). Assuming a binary choice task, relative evidence is

defined as the evidence difference – a single scalar value which changes stochastically (the DV). Evidence for one choice is therefore simultaneously equivalent to evidence against the other choice. Once this measure of relative evidence exceeds – or falls below – criterion boundaries which correspond to each choice alternative, the boundary that is first reached determines the subject's decision and RT (the decision rule). In contrast, *race models* assume that independent accumulators corresponding to each alternative (the DV) accumulate evidence in favor of their respective choice (Usher and McClelland, 2001; Brown and Heathcote, 2005; Brown and Heathcote 2008; see Figure 1, right). The first accumulator to gather the criterion amount of evidence determines the choice and RT (the decision rule).

While many models provide an account of either RT (Townsend and Ashby, 1983; Sternberg, 1969) or accuracy (Green and Swets, 1966), sequential sampling models relate shapes of RT distributions with probabilities of correct and incorrect responses, thereby explaining how RT and accuracy *jointly* vary as a function of the experimental conditions of interest. Moreover, the parameters of these sequential sampling models quantify different aspects of the decision process, including the quality of sensory information (and/or the efficiency by which the system processes sensory information), response caution, and the amount of time spent on processes unrelated to decision formation. Differences in experimental conditions explained by differences in these parameters can thus lend insight into the latent psychological processes beyond those available from mean RT and mean accuracy rate.

For example, consider a simple race model, the Linear Ballistic Accumulator (LBA) model (Brown and Heathcote, 2008; see Figure 2 for a schematic and description). The LBA contains five parameters: the *drift rate* (which corresponds to the rate of sensory evidence accumulation), the *standard deviation* (which corresponds to how much drift rates can vary across trials), the *starting point* (where the decision process begins), the *response threshold* (how much evidence is needed before making a choice), and *non-decision time* (time unrelated to the decision process, such as response execution or information encoding). With this framework, we can immediately see that bias for a particular choice would reduce the distance between the starting point and the response threshold for the accumulator corresponding to that choice (either by raising the starting point, lowering the response threshold, or both) and would be behaviorally manifested by lower accuracy rates with faster and less variable RTs. Similarly, an experimental condition that affected only the non-decision time parameter would affect mean RT, but not RT variability or accuracy. Such subtle situations, while often missed using conventional analyses of accuracy and RT, are easily explained by sequential sampling methods.

Neurophysiological approaches to studying perceptual decision making

More recently, sequential sampling models have been applied to the firing rates of neurons involved with PDM, possibly offering a quantitative bridge between neural and behavioral data (Ditterich et al., 2003; Roitman and Shadlen, 2002; Ratcliff et al., 2003; Hanes et al., 2006; Gold and Shadlen, 2001; Gold and Shadlen, 2002; Salzman et al., 2003; Hanes and Schall, 1996; Schall, 2001; Schall, 2003; Shadlen and Newsome,

2001; Churchland et al., 2008; for reviews, see Gold and Shadlen, 2007; Smith and Ratcliff, 2004; Churchland et al., 2012). One advantage of neurophysiological studies is the ability to explicitly measure spike trains from individual neurons. Prototypical PDM tasks in these single-unit studies involve simple, binary decisions that allow precise quantification between aspects of a sensory stimulus and an observed behavioral response. Here, the focus of this dissertation will be on visual tasks, many of which employ random dot patterns (RDPs, see Figure 3).

RDPs consist of a circular array of moving dots, a proportion of which are moving altogether in a single direction (*motion coherence*). Task difficulty is controlled by varying the level of motion coherence. The RDP is placed in the receptive field (RF) of the recorded neuron while the animal decides between two possible directions of coherent motion. These motion directions are selected based on the preferred and anti-preferred (i.e., null) directions of the recorded neuron (see Figure 3). The animal then indicates its decision with a saccade in the chosen direction (which will either match the neuron's preferred or null directions). This aspect of the RDP discrimination task imposes a link between a perceptual decision and a particular course of action (an eye movement) and treats decision making as a problem of movement selection. As a result, the search for the neural correlates of the DV for this task has centered on brain areas not only related to motion processing (middle temporal area, MT), but also on the preparation and generation of saccades (lateral intraparietal area, LIP; frontal eye fields, FEF; superior colliculus, SC; dorsolateral prefrontal cortex, DLPFC).

Several single-unit studies have documented that the activity of MT neurons correlates with the strength of coherent motion in the RDP displayed within their RFs and the animal's decision about the stimulus (Britten et al., 1992; 1993; Newsome et al., 1989; Newsome et al., 1995) – even on trials with 0% coherent motion (Britten et al., 1996). Lesion (Newsome and Paré, 1988; Rudolph and Pasternak, 1999) and microstimulation (Salzman et al., 1990; 1992; Salzman and Newsome, 1994; Ditterich et al., 2003) studies further established a causal link between MT activity and perceptual performance. In one key experiment, Ditterich et al. (2003) conducted a reaction-time variant of the RDP discrimination task, where monkeys could saccade as soon as arriving upon a choice (rather than viewing the stimulus for a fixed-duration and waiting for a cue before reporting their choice). The authors found that not only did their monkeys choose the preferred direction of the stimulated neuron more often and more quickly, but they also chose the null direction more slowly. This last point is especially significant as it informed a simple sequential sampling model of their neural data. In their proposed model, two processes integrate sensory signals over time in favor of the two direction options, respectively, and the choice and corresponding RT are determined by the first process to accumulate evidence to a threshold level (see Ditterich et al., 2003, Figure 6). Under this model, microstimulation increases the rate of evidence accumulation for the preferred direction and therefore increases the likelihood of and reduces RTs for that choice while decreasing the likelihood of and increasing the RTs for the opposite choice. Thus, the DV is thought to be derived from a comparison of responses from direction selective sensory neurons in MT and that MT neurons with opposite directional preferences must contribute in an opponent fashion to both decision

types. After all, if downward decisions were based solely on information from downward-preferring neurons, Ditterich et al. (2003) would not have observed slower downward decisions during stimulation of upward-preferring MT neurons. Nonetheless, even with the explanatory power of this model, the issue of where the comparison and integration of sensory signals (i.e., the DV) was being computed still remained.

Given that a choice is tantamount to an eye movement in a RDP discrimination task, neurophysiological efforts in determining the neural substrates of the DV have focused on oculomotor areas. In these studies, saccadic targets corresponding to possible direction choices were placed either inside the RF of the recorded neuron (T_{in}) or outside of it (T_{out}). Researchers discovered that activity from these neurons when monkeys performed a fixed-duration (Shadlen and Newsome, 1996; 2001) or reaction-time (Roitman and Shadlen, 2002) task was predictive of the animal's response. Moreover, when Roitman and Shadlen (2002) aligned the recorded responses with the onset of the stimulus, they found that the average firing rates of T_{in} choices rose like a ramp (see Figure 4, left), while the average firing rates steadily declined on T_{out} choices (see Figure 4, left). The dependence of LIP activity on choice was observed even on trials with 0% motion coherence (see blue lines on Figure 4, left). Aligning LIP responses to saccade initiation also revealed that the firing rates of these neurons on T_{in} choice trials were approximately the same ~70 ms prior to the eye movement, irrespective of the RDP's motion coherence (see Figure 4, right). In addition, Huk and Shadlen (2005) presented RDP stimuli to monkeys while briefly perturbing the strength of the motion stimulus during the formation of perceptual decisions. The effects of these bursts of motion change on LIP neurons, choice behavior, and response times could be explained by a

model of near-perfect integration that stops when a criterion amount of sensory evidence is accumulated. These data were taken to suggest that LIP neurons integrate evidence over time from sensory signals (represented in areas such as MT, see shaded insert in Figure 4) to form a decision and that this decision process terminates when the neurons associated with the chosen target reach a critical firing rate.

In 2006, Hanks and colleagues conducted a reaction-time variant of the RDP discrimination task while recording and microstimulating from LIP neurons. Microstimulation increased the proportion of T_{in} choices and induced faster RTs, but slowed RTs for T_{out} choices. Critically, microstimulation never directly evoked saccades nor did it change RTs in a simple saccade task, thereby establishing that LIP neurons possess a causal role in the formation of a decision.

Similar results have also been found in other oculomotor and higher order regions, such as SC (Horwitz and Newsome, 1999; 2001; Horwitz et al., 2004; Ratcliff et al., 2003), DLPFC (Kim and Shadlen, 1999), and FEF (Hanes and Schall, 1996; Thompson et al., 1996; Gold and Shadlen, 2000; 2003; Bichot et al., 2000), In one experiment, viewing of a RDP was interrupted at a random time during decision formation by turning off by the stimulus and applying a brief electrical current to the FEF (Gold and Shadlen 2000). The microstimulation evoked eye movements that deviated in the direction of the animal's subsequent choice, implying that FEF signals reflected the accumulated motion information supporting the perceptual decision. The authors applied the same microstimulation technique to another experiment (2003), where monkeys performed a variety of visual tasks based on motion information, but had to respond with a saccade that was either in the same direction as the choice

direction, in the opposite direction, or towards an unpredictable location. They found that the formation of the direction decision was reflected in the activity of FEF neurons only when the monkey could anticipate the needed eye movement. These results further verified that – at least on tasks requiring a specific and well learned behavioral response – a perceptual decision is formed as a direct transformation from sensory information into motor commands.

Several models mirroring random walk models have been proposed to explain findings from these single-unit studies and to quantify the DV and the decision rule (including one described previously by Ditteritch et al., 2003). Gold and Shadlen (2001; 2002) conceived of the DV as the log likelihood ratio (logLR) of the two choice options (i.e., the sum of the logLRs associated with each piece of evidence). The decision rule required updating this DV with new pieces of evidence until reaching a particular bound – a process which can be formalized mathematically as a SPRT (Gold and Shadlen, 2002; 2007). Another related model was proposed by Mazurek et al. (2003), which was able to explain both the neural and behavioral data from Britten et al. (1993); Shadlen and Newsome, (2001); and Roitman and Shadlen (2002). In this three-stage model, evidence (as represented by the response of ensembles of direction-sensitive MT neurons) is assumed to be accumulated by two pools of mutually inhibitory neural populations in LIP until one of the totals reaches the response criterion (see Figure 5). This model was able to predict LIP responses, mean RT for correct responses, and accuracy rates (although not the RT distributions and error RTs). Similarly, Ratcliff et al. (2003) showed that a diffusion model could account for the ramp-like activity in SC neurons, behavioral RTs, and accuracy in a dot separation task. Gold and Shadlen (2003)

also successfully related the magnitude of saccadic deviation observed during microstimulation of FEF with the likelihood of making the correct choice, as a function of motion strength and viewing time. Under this model, the decision is assumed to be the log of the expected difference between neural signals corresponding to each possible alternative (i.e., the DV). For correct choices, both the saccade deviations and the DV increased with longer viewing duration and stronger coherent motion. For incorrect choices, both values increased with longer viewing duration but decreased with stronger coherent motion.

Together, these neurophysiological findings and their corresponding models support the idea that the neural mechanisms of PDM involve tracking sensory signals early visual areas (e.g., MT) and then temporally integrating this information by neurons involved with motor output (e.g., LIP, FEF, etc) until a threshold level of activity is reached and the appropriate behavioral response (e.g., a saccade) is deployed (Hanes and Schall, 1996; Kim and Shadlen, 1999; Gold and Shadlen, 2001; Schall, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Ditterich et al., 2003; Huk and Shadlen, 2005; Gold and Shadlen, 2000; 2003; 2007; Horwitz et al., 2004; Hanks et al., 2006). The robust relationship between neural activity and behavior suggests that decision making is carried out by the same neurons that ultimately initiate the appropriate motor response. However, these studies require macaques to practice making binary decisions about simple visual stimuli for weeks or even months at a time. Such a paradigm stands in stark contrast to the types of more deliberate and less experienced decisions we encounter in real life (e.g., selecting the right stocks to invest, choosing who to marry, etc). Moreover, these RDP discrimination tasks have been

designed to treat the decision process as a problem of movement selection, where a stereotyped course of action (e.g., a saccade) is almost always associated with a particular choice. Thus, the issue of where and how the brain forms perceptual decisions that are not used to select a particular movement still remains unknown.

Neuroimaging approaches to studying perceptual decision making

The paradigm used in the neurophysiology studies surveyed thus far makes it difficult to test whether or not a general mechanism of PDM exists. Another limiting factor of those experiments is the restricted number of recording sites, as well as the fact that the recorded neurons are often chosen circularly (i.e., a neuron exhibits the desired activity profile and is thus chosen for measurement in an experiment). Moreover, it cannot be assumed that the activity of individual neurons can be generalized perfectly to a whole population. Studies simultaneously measuring from populations of neurons typically use neuroimaging techniques such as EEG or BOLD fMRI. Here, we review the BOLD fMRI studies on the neural correlates of the DV.

In order to identify neural correlates of the DV that read out object categorization evidence, Heekeren and colleagues (2004) measured BOLD activation from the ventral temporal cortex in human subjects while they performed a discrimination task between faces and houses masked by noise (two classes of stimuli which differentially activate the ventral temporal cortex). The researchers found that activity in the DLPFC was greatest when sensory evidence was strongest and moreover, activity in the DLPFC tended to covary with the magnitude of the difference in the measured BOLD signal from this region of ventral temporal cortex (which presumably is tracking sensory evidence for faces the way MT would for motion direction). As such, this was one of the

first studies that provided evidence for an amodal mechanism of PDM. Similarly, Pessoa and Padmala (2005) had human subjects detect fear among images of faces at near-threshold. The authors calculated which brain regions correlated with choice on neutral trials (and thus, correlated with a decision that was made independently of the amount of sensory evidence available), similarly to the analyses performed on the activity of MT neurons (Britten et al., 1996). The brain regions which most reliably predicted choice responses were the posterior cingulate cortex (PPC), the medial prefrontal cortex (MPFC), right inferior frontal gyrus (IFG), and left insula. Both of these studies observed brain areas not directly involved with motor processing and thereby interpreted their findings as evidence for a domain general mechanism of PDM. However, subjects in these two studies did not vary *how* they made their responses, so it was unclear if these areas compute the DV in a manner that is independent of the response modality.

To address this issue, Tosoni et al. (2008) had human subjects discriminate between noisy images of faces and places. A peripheral target was also presented along with the stimulus, towards which subjects had to either make a saccade (if they believed the presented image was a face) or point with their finger (if they believed the image was a place). As with the study by Heekeren et al. (2004), Tosoni and colleagues predicted that regions involved with accumulating sensory evidence would be more active on trials with more signal (i.e., less noise). The authors found that the activation of certain effector-specific regions (e.g., regions involved with hand reaching or making eye movements) was modulated by the amount of sensory evidence in the image. Unlike the aforementioned fMRI work, these results suggested that perceptual decisions rely on

an accumulator mechanism that integrates sensory evidence towards a particular motor response – a finding that corroborated much of the prior neurophysiological literature.

These three experiments made predictions about the relationship of choice (and the BOLD responses associated with a region involved in computing the DV) and sensory evidence. Both Heekeren et al. (2004) and Tosoni et al. (2008) predicted areas integrating sensory information ought to have greater activation when more sensory evidence was available in an image, while Pessoa and Padmala did not even correlate sensory evidence with choice (having restricted their primary analyses to neutral trials). However, decision signals should also correlate with behavioral performance – namely, RTs. In 2007, Thielscher and Pessoa investigated this requirement by relating RTs to fMRI responses for a PDM task. Subjects viewed face stimuli that differed in emotional expression and would press a button whenever they perceived a fearful or disgusted face. The authors reasoned that a neural decision signal should not only predict the behavioral choice, but also the time it takes to arrive at that choice. As expected, the RT distributions for their experiment were shaped like an inverted U, being fastest for clearly fearful or disgusted faces, and slower for neutral faces or faces with intermediate levels of fear or disgust. In the context of a decision making model, it would make sense that the accumulation of sensory evidence takes longer when there is little sensory evidence available, and that a brain region representing the DV should not only follow the general pattern of RTs, but also track trial-to-trial fluctuations of RTs. The authors found that the BOLD signals of MFG, IFG and anterior cingulate not only predicted perceptual choice, but also followed an inverted U-shape and tracked trial-to-trial RT. Furthermore, they calculated the probability with which the amplitude of BOLD

activation in single neutral trials (which contain no valence) predicted the subjects' support of a fearful or disgusting percept. A network of regions predicted decisions during neutral trials, including the superior temporal sulcus, MFG, IFG, anterior cingulate, the anterior insula. These results further established that amodal regions were involved in representing the DV.

Thielscher and Pessoa, however, did not have their subjects vary response modality like Tosoni et al. (2008) did. Moreover, none of the previous studies utilized decision making models to support their results. Here, we present a BOLD fMRI study (Chapter 1) investigating whether or not mechanisms of PDM depend on the modality of the motor response. Subjects discerned coherent motion among two overlapping RDPs and alternated between two response modalities (saccades versus button presses). We then fit our accuracy and RT data with the linear ballistic accumulator model (LBA; see Figure 2). The use of a mathematical model allowed us to precisely predict the BOLD activation profile expected from cortical areas that accumulate sensory evidence. Consistent with the reasoning from Thielscher and Pessoa (2007), our analyses revealed that BOLD responses ought to be higher on trials with greater perceptual difficulty. Only one area demonstrated the activation profile predicted for a region involved in accumulating sensory evidence regardless of response modality – the right insula. Thus, while it is clear from neurophysiology that specialized areas do play an important role in fluidly translating sensory information into a motor response, our findings – as well as a handful of other neuroimaging studies conducted since (Kayser et al., 2010; Liu and Pleskac, 2011) – provide strong evidence that a distinct mechanism is computing a more

abstract representation of the DV that is possibly sending continuous input to effector-specific sensorimotor areas (e.g., LIP/IPS, FEF, etc) during the decision making process.

Linking neural patterns to perceptual performance

In addition to failing to rule out a modality-independent mechanism of PDM, another shortcoming of single-unit recordings is that the results rely on the assumption that the activity of a single neuron can be generalized to a population of neurons. A sensory stimulus such as a RDP evokes responses in a large population of neurons. Individual neurons typically respond to only a small fraction of possible sensory inputs (i.e., they are tuned to a limited range of stimulus properties) and these responses are inherently noisy. Thus, the activity of the whole population must be "read out" in order to support perceptual decisions. But how does one determine how ensembles of neurons support whole percepts?

While Gold and Shadlen (2001; 2002) outlined a model for how neurons could compute the DV for a binary choice, their model did not extend to population responses for a range of stimuli. Movshon and Jazayeri (2006) proposed a straightforward model of decoding sensory information (i.e., inferring the stimulus that elicited the observed pattern of neural activity) by computing the likelihood that the response of each neuron was elicited by a particular stimulus and then combining those likelihoods to determine the overall likelihood of that stimulus (and reiterating this calculation for all possible stimuli in order to map the entire likelihood function). This model provided excellent fits for several psychophysical experiments and determined the optimal strategies needed for a variety of tasks. For example, when performing a discrimination between very similar stimuli (e.g., 80° versus 82°), the firing rates of neurons tuned maximally to the feature

values of interest (*on-target neurons*) are highly similar and therefore, uninformative for distinguishing between them (Regan and Beverley, 1985; Hol and Treue, 2001; Navalpakkam and Itti, 2007). Rather, neurons tuned away from the target feature (*off-target neurons*) possess more information for such a task since the differences in firing rates for the two stimuli will be much greater (Navalpakkam and Itti, 2007; Scolari and Serences, 2009; 2010).

In Chapter 2, we capitalize on the differences between on- and off-target neurons to link population responses with perceptual performance. Subjects performed a difficult discrimination between oriented gratings, emphasizing either speed or accuracy while we measured BOLD responses. By using a forward encoding model that maps stimuli to population responses (Brower and Heeger, 2009; 2011), we computed how activation levels from the most informative neurons in visual cortex (the off-target neurons) predicted performance when accuracy was emphasized. A logistic regression further supported this link by revealing a trial-by-trial relationship between behavioral accuracy and BOLD activation levels in off-target populations. In addition, a variant of the LBA (Van Maanen et al., 2011) that provides trial-by-trial estimates of the latent cognitive processes involved in perceptual decision making revealed a correlation between activation levels in off-target neural populations and the rate of sensory evidence accumulation when subjects prioritized accuracy over speed.

Another recent model derived from single-unit recordings is the feature-similarity gain model of attention (Treue and Trujillo, 1999; Treue and Martinez-Trujillo, 2004). These studies demonstrated that when animals attend to a particular

feature value, individual neurons which prefer that value will fire more rapidly, while attending to unpreferred feature values attenuates firing rates (Treue and Trujillo, 1999; Treue and Martinez-Trujillo, 2004; Maunsell and Treue, 2006). The feature-similarity gain model generalized these findings, maintaining that selective attention increases the gain of neuronal populations that are tuned to a relevant feature value, while suppressing the gain of neurons tuned to irrelevant or dissimilar feature values (McAdams and Maunsell, 1999a, Treue and Martinez-Trujillo, 1999, Martinez-Trujillo and Treue, 2004, Maunsell and Treue, 2006). Consistent with these neurophysiology data, psychophysical studies have also documented that feature-based attention selectively increases sensitivity to relevant visual features (Busse et al., 2008; Sàenz et al., 2003; Baldassi and Verghese, 2005; Felisberti and Zanker, 2005; Liu et al., 2007; Liu and Hou, 2011; Ling et al., 2009). However, the behavioral correlates of feature-based attentional suppression are less clear. One study (Ling et al., 2009) suggested that feature-based attention suppresses neurons tuned away from an attended feature, while another study reported only an enhancement of an attended feature value without concurrent suppression of dissimilar features (White and Carrasco, 2011). However, neither study systematically varied the relationship between the target stimulus and the focus of feature-based attention, so the consequence of attentional suppression on the efficiency of processing unattended features was not directly evaluated.

Chapter 3 outlines two experiments testing the perceptual predictions of the feature-similarity gain model using a combination of psychophysics and mathematical modeling. In Experiment 1, subjects searched for an oddball motion target among four RDPs; the LBA model revealed increased rates of sensory evidence accumulation for

validly cued targets compared to neutral targets (facilitation). In addition, drift rates on invalidly cued trials varied as a function of the cue target offset, with highest drift rates observed for invalid targets that most closely matched the cue. However, while these patterns are suggestive, the drift rates on invalid trials were not significantly lower than drift rates on neutral trials. In Experiment 2, we increased competition between the target and distractors in an effort to place additional demands on suppressive attentional mechanisms and an analysis of drift rates revealed extremely robust evidence for both attention-related facilitation and suppression.

Collectively, these behavioral and fMRI results both challenge and confirm the existing neurophysiological literature on perceptual decision making: Chapter 1 provides neural evidence for a general mechanism of PDM, Chapter 2 shows that the integration and readout of sensory information impacts perceptual performance, and Chapter 3 illustrates that attentional manipulations influence perception in a manner consistent with the enhancement and suppression of distinct neural populations predicted from single-unit recordings.

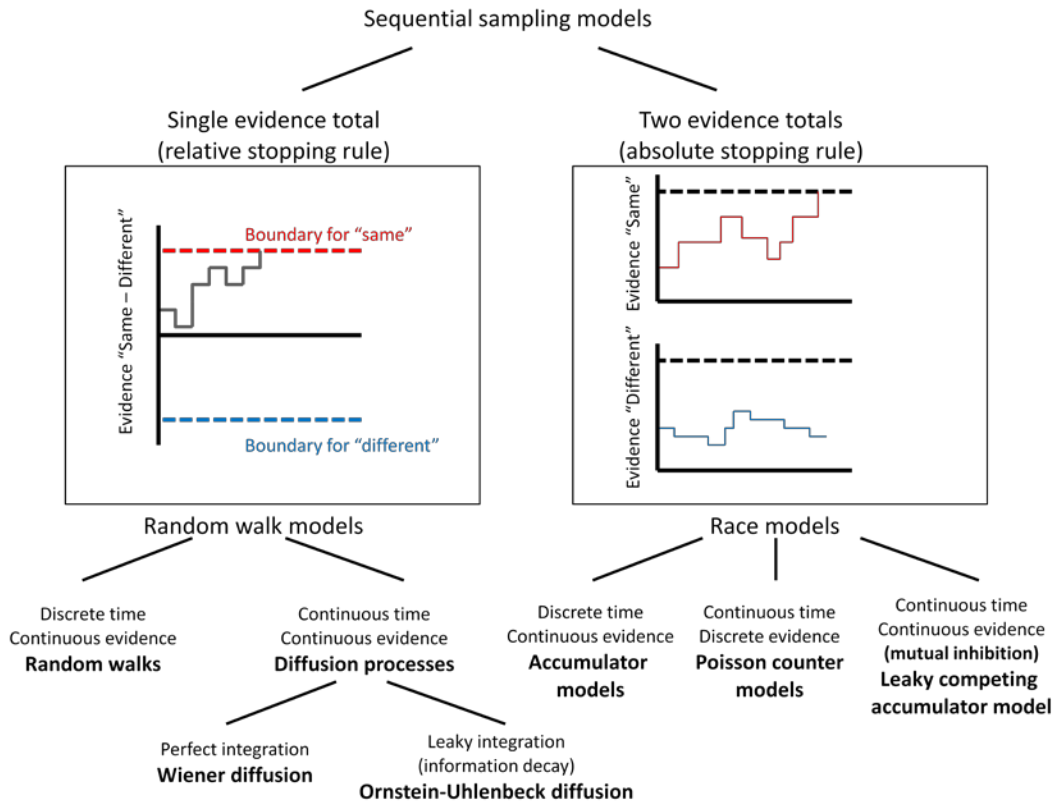

Figure i.1. The different classes of sequential sampling models for two-choice decisions (Smith and Ratcliff, 2004). These models assume that decisions are made by integrating noisy sensory evidence over time until a criterion amount of information is reached and a choice is made.

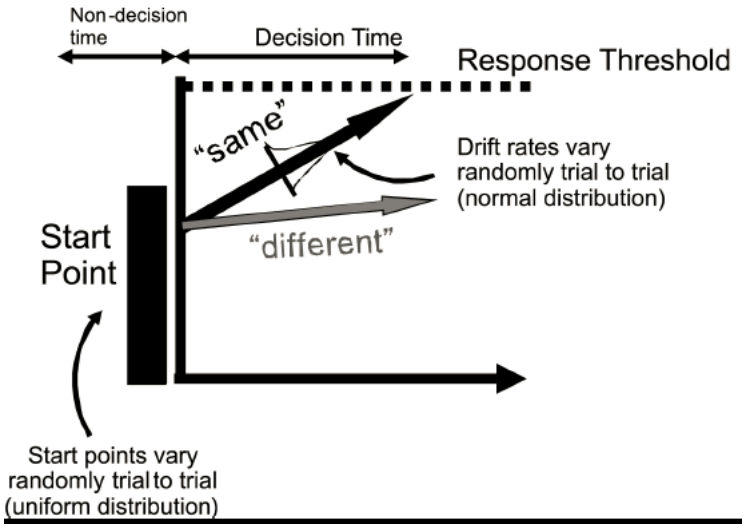

Figure i.2. Schematic of a race model: the linear ballistic accumulator model (LBA; Brown and Heathcote, 2008). One accumulator corresponds to each possible response ("same" shown in black and "different" shown in gray). The *drift rate* is the rate at which sensory evidence is accumulated and is assumed to be determined by the stimulus properties; thus it can be inferred from this example that the stimulus is the "same." Each drift rate for each accumulator varies trial by trial in a normal distribution with a particular *standard deviation*; the drift rate parameter that the LBA estimates is the mean of this distribution. The decision process begins at a *starting point* drawn randomly on a trial to trial basis from a uniform distribution. The *response threshold* is the criterion boundary that determines how much evidence is needed before a decision is made. The final response time is the time it takes for the first accumulator to reach the response threshold plus some constant (*non-decision time*).

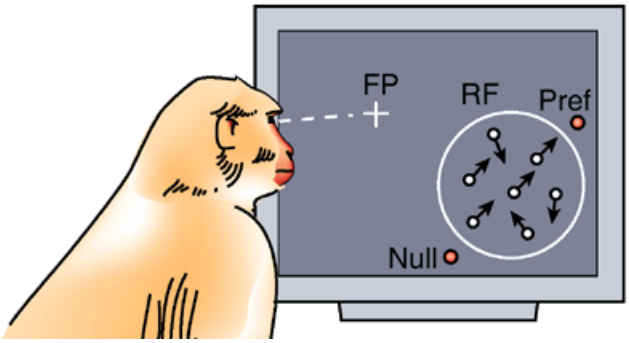

Figure i.3. Prototypical task and stimuli used in neurophysiology studies of PDM. A random dot pattern (RDP) contains coherent motion in one of two possible directions and is placed in the receptive field of the recorded neuron. The two motion directions are selected based on the preferred and null (180° from the preferred) directions of the neuron. Choice targets are positioned relative to the receptive field. That is, if the animal decides the RDP contains coherent motion in the preferred direction, it will make a saccade accordingly (similarly with the null direction). Task difficulty is manipulated by adjusting the percentage of dots moving coherently. Figure adapted from Nadler and DeAngelis, 2005 (Figure 1).

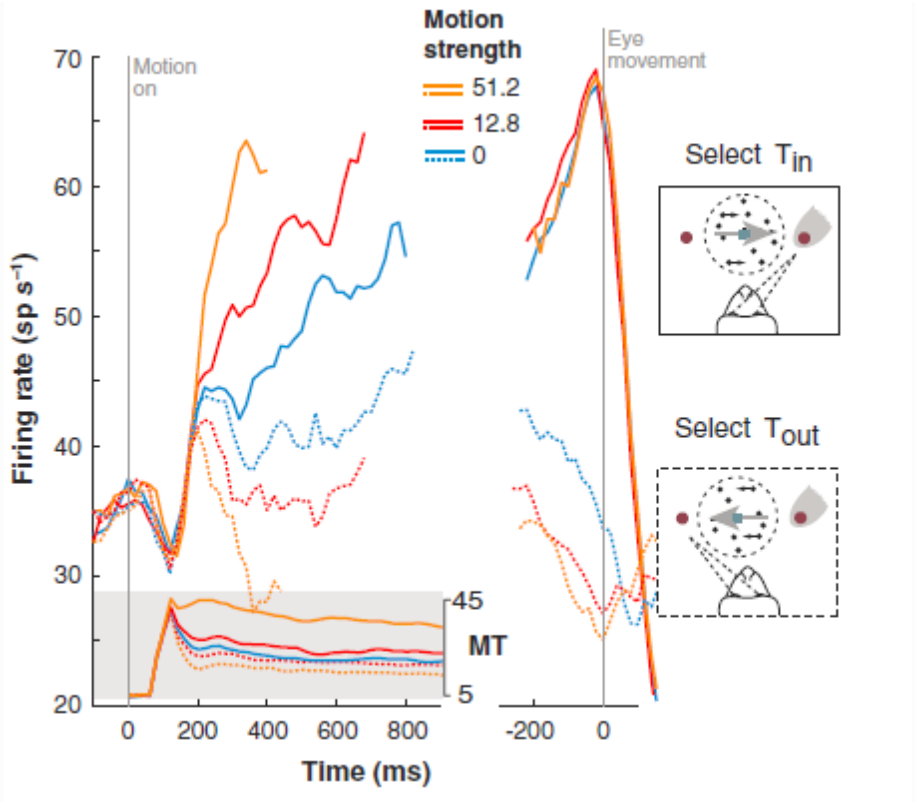

Figure i.4. Neural mechanisms of a perceptual decision of motion direction. Response of LIP neurons during decision formation (Roitman and Shadlen, 2002). Average firing rate from LIP neurons is shown for three levels of difficulty. Responses are grouped by motion strength and direction of choice. *Left*: the responses are aligned to the onset of the RDP. Averages are shown during decision formation (curves are truncated at median RT or 100 ms before the initiation of the saccade). The shaded region shows average responses from direction selective neurons in area MT to motion in the preferred and null directions. After a delay, MT responses appear to asymptote. *Right*: the responses are aligned to the eye movement response. The LIP firing rates approximate the integral of a difference in firing rate between MT neurons with opposite direction preferences. Dashed lines indicate activity from trials where the monkey chose the direction where the saccadic target was inside the RF from the recorded neuron (T_{in}) while solid lines indicate activity from trials where the monkey chose the direction where the saccadic target was outside of the RF from the recorded neuron (T_{out}). Adopted from Gold and Shadlen, 2007 (Figure 5).

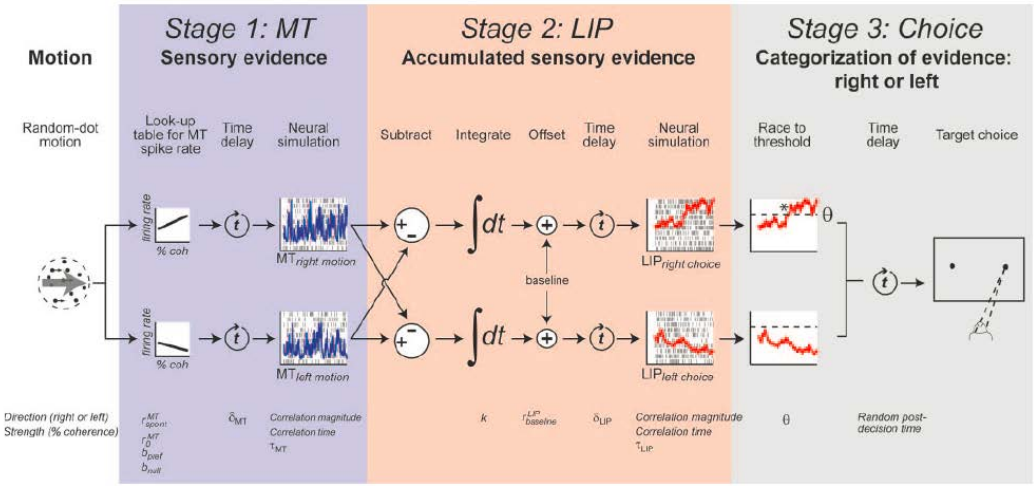

Figure i.5. Model of the decision making process. The model is divided into three stages: representation of sensory evidence, accumulation of sensory evidence into the DV, and comparison of the DV to threshold (the decision rule). In the first stage, the authors simulated the neurons to RDPs of varying motion. For illustrative purposes, all simulations assume coherent motion was either rightward or leftward. Each neuron in the two pools produces a sequence of spikes with an expected rate proportional to the strength of motion, based on values from Britten et al. (1992). The averaged spike rate from the two pools is the output of stage 1. In the second stage, two pools of LIP neurons were simulated: one corresponding to a rightward choice and the other, leftward. Unlike MT neurons, the expected firing rates of these neurons are time-dependent, determined by the integral of the difference in the output of the right and left MT pools. The expected LIP firing rate is calculated by integrating the difference in spike rate signals from MT starting from when the coherence-dependent MT response begins. As with the first stage, the averaged spike rate from the two pools of LIP neurons is the output of stage 2. In stage 3, the LIP responses are compared to a criterion threshold. The two ensemble average spike rates comprising the LIP signals (output of stage 2) race against each other to provide the weight of evidence for their preferred choice direction. The first to reach the decision threshold determines the target choice and RT (see Mazurek et al., 2003). Adapted from Mazurek et al., 2003.

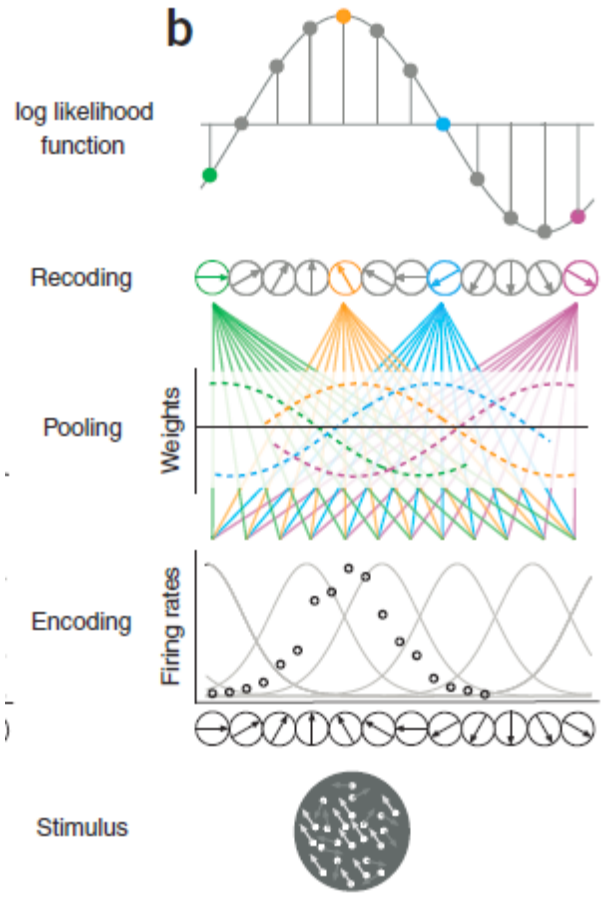

Figure i.6. Computing the likelihood for motion direction. A RDP stimulus (bottom row) produces firing rate responses across differently tuned neural populations in MT (second row). The smooth curves in the second row represent neuronal tuning curves for a particular motion direction, while the open circles indicate the population response on a given trial. The sensory signals from the population of the encoding neurons provide weights into which the signals are recoded (third row). The output layer (top row) consists of an ensemble of neurons where the weighted sensory signals converge. The neurons here represent the log likelihood for all possible directions – the likelihood function. In the top row, the average likelihood profile is shown, with the colored points representing the average likelihoods of four example directions. The peak of the average likelihood function – and thus, the expected maximum likelihood estimate of the stimulus direction – is shown in orange. Adapted from Movshon and Jazayeri, 2006 (Figure 2).

CHAPTER 1: Domain General Mechanisms of Perceptual Decision Making in Human Cortex

Ho, T.C., Brown, S.D., Serences, J.T. (2009). *Journal of Neuroscience*, 29(27), 8675-8687.

Abstract

To successfully interact with objects in the environment, sensory evidence must be continuously acquired, interpreted, and used to guide appropriate motor responses. For example, when driving, a red light should motivate a motor command to depress the brake pedal. Single-unit recording studies have established that simple sensorimotor transformations are mediated by the same neurons that ultimately guide the behavioral response. However, it is also possible that these sensorimotor regions are the recipients of a modality independent decision signal that is computed elsewhere. Here, we used fMRI and human observers to show that the timecourse of activation in a subregion of the right insula is consistent with a role in accumulating sensory evidence independently from the required motor response modality (saccade vs. manual). Furthermore, a combination of computational modeling and simulations of the BOLD response suggests that this region is not simply recruited by general arousal or by the tonic maintenance of attention during the decision process. Our data thus raise the possibility that a modality independent representation of sensory evidence may guide activity in effector-specific cortical areas prior to the initiation of a behavioral response.

Domain General Mechanisms of Perceptual Decision Making in Human Cortex

On a moment-to-moment basis, the brain must infer the most likely state of the world given a variable amount of sensory evidence, a process referred to as *perceptual decision making* (Newsome et al., 1989; Salzman and Newsome, 1994; Gold and Shadlen, 2001; Shadlen and Newsome, 2001). In a prototypical laboratory experiment, observers view a noisy field of moving dots drifting to the left or to the right (a random-dot pattern, or RDP) and indicate the direction with a saccade in the appropriate direction. The firing rate of motion-selective neurons in the middle temporal area (MT) monotonically tracks the quality of the available sensory evidence, which is systematically manipulated by varying the percentage of dots moving in a common direction (termed *motion coherence*; Newsome et al., 1989; Salzman et al., 1992; Britten et al., 1996; Shadlen et al., 1996; Gold and Shadlen, 2001; Shadlen and Newsome, 2001; Ditterich et al., 2003; Mazurek et al., 2003; Gold and Shadlen, 2007). This sensory information is then thought to be temporally integrated by spatially selective oculomotor neurons in areas such as the lateral intraparietal area (LIP), frontal eye fields (FEF), dorsal lateral prefrontal cortex (DLPFC), and superior colliculus (SC) until a threshold level of activity is reached and an appropriate eyemovement response is triggered (Hanes and Schall, 1996; Kim and Shadlen, 1999; Gold and Shadlen, 2001; Schall, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Ditterich et al., 2003; Huk and Shadlen, 2005; Gold and Shadlen, 2007; Churchland et al., 2008; Kiani et al., 2008). Microstimulating oculomotor neurons within some of these regions can also bias the response outcome, implying a causal role in perceptual decision making (Gold and Shadlen, 2000, 2003; Horowitz et al., 2004; Hanks et al., 2006).

The strong coupling between neural activity and behavior suggests that decision making is carried out by the same neurons that ultimately initiate the appropriate motor response (here termed the *modality-dependent* hypothesis). For example, oculomotor regions mediate simple decisions requiring saccadic responses, and somatosensory cortex (S1) mediates vibrotactile decisions (Romo and Salinas, 1999; Romo et al., 2002; Romo and Salinas, 2003; Tegenthoff et al., 2005; Preuschhof et al., 2006). Decisions about complex stimuli, such as images of faces or places, are also mediated by motor-specific cortical areas depending on the response-output modality that is required by the task (Tosoni et al., 2008).

While these studies leave no doubt that specialized motor areas play an important role in translating sensory information into a behavioral response, it is also possible that a separate mechanism computes a more abstract supramodal representation of sensory evidence and sends a continuous input signal to motor-effector specific sensorimotor areas during the course of the decision process (termed the *modality-independent* hypothesis). Here, we show that a region of right insula exhibits an activation profile consistent with the accumulation of sensory evidence during decision making, independent of response modality (saccade vs. manual). This finding raises the possibility that a modality-independent mechanism guides activity in motor-specific regions prior to movement initiation.

Methods

Subjects

Twelve right-handed subjects (9 female) were recruited from the University of California, Irvine (UCI) community, and one right-handed subject (male) was recruited from the University of California, San Diego (UCSD) community. Data from one subject (female) were discarded because the manual and saccadic responses were not recorded correctly during the scanning session. All had normal or corrected-to-normal vision. Each subject gave written informed consent as per Institutional Review Board requirements at either UCI or UCSD and completed two 1 hour training sessions outside the scanner and one 1.5 hour session in the scanner. Compensation for participation was $10/hour for training and $20/hour for scanning.

Stimuli and task

Visual stimuli were generated using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) for Matlab (version 7.1; The Math Works, Natick, MA), presented at a frame rate of 60 Hz, and projected onto a screen at the back of the scanner bore that subjects viewed through a mirror. Buttonpress responses were made on an fMRI-compatible response box using the fingers of the right hand.

Subjects viewed a display consisting of two overlapping centrally presented RDPs – one comprised of 100 red dots and the other made up of 100 blue dots – against a light gray background (Figure 1). Each small dot subtended 0.1° visual angle, and the circular stimulus aperture subtended 4° of visual angle (radius) with a small circular cutout around fixation (1° radius). On every trial, the coherence level for each RDP was determined by the proportion of dots moving in one of four possible directions – either to the upper left, upper right, lower left, or lower right – while the direction of each remaining dot was selected from a uniform distribution (across 360°). Each RDP moved

in a different direction (pseudo-randomly determined) and contained a motion coherence level of either 40% or 80% so that the total motion signal in the display was equated on every trial (e.g., if the red RDP had 40% coherent motion, the blue RDP would contain 80% and vice-versa). Additionally, there were four small black circles (subtending 1° and centered 11.3° from fixation) arrayed at each corner of the screen that served as saccade targets.

At the start of each trial, a cue was presented for 750ms in the form of a colored fixation cross (either red or blue), indicating which of the two RDPs subjects should monitor. This colored fixation cross remained onscreen throughout the stimulus display. Subjects were asked to judge the direction of the coherent motion of the RDP to which they were cued. If the cued RDP contained 80% coherent motion, then the trial was termed *easy*, and if the cued RDP contained only 40% coherent motion, the trial was termed *hard*. The stimulus remained onscreen for 1500ms, after which only a white fixation cross was displayed for the remainder of the trial. Each trial lasted 5250ms and subjects were allowed to respond any time after stimulus onset and up until the termination of the trial. Each run in the scanner consisted of 32 task trials randomly interleaved with 10 null trials (which were the same duration as a normal trial but only required passively viewing the fixation cross: no RDPs were presented). The color of the cue and the cued motion direction were randomized and counterbalanced within each block and each run ended with 10s of passive fixation.

Response modality was alternated on a run-by-run basis and subjects were informed beforehand whether they were to make their responses via saccades or manual button presses. When making saccadic responses, subjects were instructed to keep their

eyes on the fixation cross and then to make one clean eye movement to one of the four peripheral black circles before redirecting their gaze back to the central cross in preparation for the start of the next trial. When responding with button presses, subjects were instructed to keep their eyes on the fixation cross throughout the entire trial and to press one of four buttons spatially arrayed to correspond to the four possible target directions.

Eye movement data acquisition and analysis

At UCI, eye movements were monitored using an infrared video eye tracker (Applied Science Laboratory, long range optics system, Bedford, MA); at UCSD, an Avotech SV-7021 (Stuart, FL) infrared eye tracker was used. The position of the right eye was sampled at 60Hz and before each run, the eye tracker was recalibrated. Preprocessing and saccade extraction were performed using the ILAB toolbox for Matlab (http://www.brain.northwestern.edu/ilab/, Gitelman, 2002). The raw data were first binned into temporal epochs corresponding to each trial, then blinks (periods when the pupil disappeared), as well as 5 samples on either side of each blink, were marked and removed from the epoched data. The following parameters were used to identify saccades: an initial velocity threshold of 30° per second, a minimum saccade duration of 35ms, and a minimum fixation duration of 100ms at the endpoint of the saccade. Response times on saccade trials were defined as the time between the onset of the stimulus and the first saccadic eye movement that deviated more than 3° from fixation in the direction of one of the four peripheral targets (data were scored by hand on a trial-by-trial basis to ensure accuracy).

fMRI data acquisition and analysis

For 11 of the subjects, MRI scanning was carried out on a Phillips Intera 3-Tesla scanner equipped with an 8-channel head coil at the John Tu and Thomas Yuen Center for Functional Onco Imaging, University of California, Irvine. Anatomical images were acquired using a MPRAGE T1-weighted sequence that yielded images with a 1x1x1mm resolution. Whole brain echo planar functional images (EPI) were acquired in 35 transverse slices (TR = 2000 ms, TE = 30 ms, flip angle = 70°, image matrix = 64 x 64, field of view = 240 mm, slice thickness = 3 mm, 1mm gap, SENSE factor = 1.5). For the remaining subjects, scanning was carried out on a General Electric 3T scanner equipped with an 8-channel head coil at the W.M. Keck Center for Functional MRI, UCSD. Anatomical images were acquired using a MPRAGE T1-weighted sequence that yielded images with a 1x1x1mm resolution. Whole brain echo planar functional images (EPI) were acquired in 33 transverse slices (TR = 2000 ms, TE = 30 ms, flip angle = 90°, image matrix = 64 x 64, field of view = 240 mm, slice thickness = 3 mm, 1mm gap).

Data analysis was performed using BrainVoyager QX (v 1.91; Brain Innovation, Maastricht, The Netherlands) and custom timeseries analysis routines written in Matlab. Data from the main experiment were collected in 8 or 10 runs per subject (i.e., either 4 or 5 runs per response modality, respectively), with each run lasting 230s. EPI images were slice-time corrected, motion-corrected (both within and between scans), high pass filtered (3 cycles/run) to remove low frequency temporal components from the timeseries, and spatially smoothed with a 4mm FWHM kernel. The motion parameters were used to estimate and remove motion induced artifacts in the timeseries of each voxel using a general linear model (GLM). The timeseries from each voxel in each observer was then z-transformed on a run-by-run basis to ensure that the timeseries had a mean of zero. All

anatomical and EPI images were transformed into the atlas space of Talairach and Tournoux (1988) before group analyses were carried out.

Linear ballistic accumulator (LBA) model

Behavioral data were modeled using the LBA, a simplified version of the ballistic accumulator (Brown and Heathcote, 2005), which was in turn a simplified version of Usher and McClelland's leaky competing accumulator (Usher and McClelland, 2001). The simplifications included in the LBA allow it to keep the essential predictive qualities of Usher and McClelland's original model, but with much improved analytic tractability. The simplifying assumptions used in the LBA are similar to those in some other neurally-inspired models of decision-making, most notably the LATER model of Reddi and Carpenter (2000) and the random ray model of Reeves et al. (2005).

In the LBA, each of the four response alternatives (motion directions) is represented by an independent linear accumulator, illustrated in Figure 2. On each trial, each accumulator begins with a random activation level that is independently drawn from a uniform distribution on [0,*A*]. During decision making, activity in each accumulator increases linearly, and a response is triggered as soon as the first accumulator crosses a response threshold (*b*). The predicted response time is simply the time taken to reach the threshold, plus a constant offset time t_0. 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 the four accumulators. To simplify matters, we always assumed that these normal distributions share a common standard deviation (*s*). The means of the normal distributions reflect the perceptual input: when the motion direction of the cued RDP closely matches the response assigned to a particular

accumulator, that accumulator will have a large drift mean rate, and vice versa. We estimated a parameter for the mean drift rate of the accumulator corresponding to the correct response (d_c) and assumed that the other three accumulators had equal mean drift rates (1-d_c)/3, keeping the total of all four drift rates fixed at 1. We also calculated a more detailed analysis with different mean drift rates for the accumulators corresponding to incorrect responses. That analysis showed obvious differences (e.g., the mean drift rate for the response opposite the correct response was about 10% smaller than the mean drift rates for responses that were orthogonal to the correct response) but all of the substantive results were unchanged.

Brown and Heathcote (2008) showed that the LBA accommodates all the benchmark phenomena observed in choice RT paradigms. The LBA is also sufficiently simple in that there are closed form solutions for the densities of predicted RT distributions, making it easy to apply to data such as ours. These solutions were used to calculate likelihood values when fitting the model to data. We assessed the goodness-of-fit between the observed RT distributions and those predicted by the LBA model using the quantile maximum product statistic (Heathcote et al., 2002; Heathcote and Brown, 2004). The parameters of the model were adjusted to maximize the goodness-of-fit using the simplex algorithm (Nelder and Mead, 1965; Brown and Heathcote, 2008).

Predicting blood oxygenation level dependent (BOLD) responses based on the rate of evidence a*ccumulation*

Assuming a different rate of evidence accumulation on easy and hard trials, we generated predictions of the BOLD response profile within regions involved in accumulating sensory evidence during perceptual decision making. The model is

primarily motivated by the work of Shadlen and coworkers, who have shown that the firing rates of neurons in areas such as LIP monotonically increase until a response threshold is achieved and a response is executed. In our simulation, we assumed that the estimated drift rate on easy and hard trials is a proxy for neural activity (Figure 3a); we then convolved this ramping activity profile with a canonical model of the BOLD response (a difference of two gamma functions, time to peak 5s, undershoot ratio 6, time to undershoot peak 15s). Assuming that the firing rate of 'accumulator' neurons in areas like LIP falls off after a response is made (see e.g. Shadlen and Newsome, 2001; Roitman and Shadlen, 2002), the simulation predicts that lower drift rates will produce larger and temporally extended BOLD responses because the response is proportional to the integrated amount of neural activity during the decision process (Figures 3a,b). However, this same effect – larger and temporally extended BOLD responses on hard trials - might also be expected in a region involved in maintaining selective attention to relevant aspects of the stimulus display during decision making, or in a region that more generally participates in sustaining a task set or an aroused state. Thus, it is not possible to distinguish areas involved in accumulating sensory evidence based solely on an increased response associated with perceptual difficulty. Fortunately, the simulation also predicts that the BOLD response should rise more slowly on hard trials compared to easy trials, because hard trials are associated with a more gradual ramping of neural activity (see shaded region of Figure 3b). In contrast, regions that are involved in general attentional processes should be uniformly engaged for the duration of the decision process, resulting in a similar main effect of perceptual difficulty without the accompanying shift in

response latency. Two variants of such attentional accounts – along with the predicted BOLD response profiles – are shown for comparison in Figures 3c-f.

Identifying supramodal mechanisms of information accumulation

The main goal of the analysis was to use a two step inferential process to define regions that (1) exhibit a larger and temporally extended response on hard trials compared to easy trials, and (2) exhibit a temporally delayed BOLD response on hard compared to easy trials (see Figures 3a,b). These properties define regions that are likely performing evidence accumulation, rather than some other role in the decision process.

To identify ROIs that respond more on hard compared to easy trials (step 1 in the analysis), the hemodynamic response function for each event type (easy saccade, hard saccade, easy manual, hard manual) was estimated using a general linear model (GLM) and a finite impulse response model that included separate regressors to estimate the BOLD response at the time of event onset and at each of the next 8 timepoints following that event (times 0s-16s poststimulus, see Dale & Buckner, 1997). Using this approach, the rows in the GLM design matrix correspond to the number of timepoints in a scanning session and the columns correspond to the relative temporal position of each model regressor with respect to the time of event onset. Each of the 9 timepoints was modeled with a ‘1’ in the appropriate row and column of the GLM design matrix, yielding scaled fit coefficients (beta weights) at each timepoint for each event type.

Additional regressors-of-no-interest were included to model the mean response across the 9 timepoints following incorrect trials, collapsed across trial type. A three-way repeated measures ANOVA with response modality (saccade vs. manual), perceptual difficulty (easy vs. hard), and time (0s-16s poststimulus, in nine intervals) as factors was

then performed on the estimated beta weights; ROIs were defined based on the interaction between perceptual difficulty and time, collapsed across response modality. All statistical maps were thresholded at $p<0.05$, after correcting for multiple comparisons using the false discovery rate algorithm implemented in Brain Voyager.

Having identified ROIs where the response is larger and temporally extended on hard compared to easy trials (step 1 of the analysis), we next tested for latency differences in the onset of the BOLD response in each ROI (step 2 in the analysis) by evaluating the interaction between perceptual difficulty and time across only the first two timepoints (0s-2s) of the event-related BOLD responses. A significant interaction across this temporal window indicates a differential slope during the rising phase of the responses, which is consistent with the accumulation of sensory evidence and inconsistent with the maintenance of sustained attention or general arousal.

Because both analytical steps involved evaluating the interaction between perceptual difficulty and time (albeit across different temporal windows), we performed a "leave-one-out" cross validation procedure to ensure that the selection of voxels to include in an ROI during step 1 (larger and temporally extended response on hard compared to easy trials) did not bias the outcome of the statistical test in step 2 (difference in onset latency). Using this procedure, ROIs that exhibited a significant interaction between perceptual difficulty and time (from 0s-16s poststimulus) were identified using data from 11/12 subjects, and then the data from the remaining subject was extracted from each ROI and used for statistical tests (Tables 3,4) and for generating timeseries plots (Figures 6,7,8,10). This procedure was repeated 12 times across all permutations of leaving one subject out, generating 12 sets of ROIs (see Table 2). In

addition to protecting against bias when evaluating differences in response latency, this procedure also ensured that the timecourses are not biased by the inclusion of noise that is favorable to our conclusions (see Kriegeskorte et al., 2009; Vul and Kanwisher, in press; Vul et al., in press). All analyses of the BOLD response used this leave-one-out procedure, with the exception of the results reported from hMT+; however, hMT+ was identified using independent localizers, so bias of this sort was not an issue (see below).

hMT+ Functional Localizer

To identify motion responsive voxels in the human MT region (hMT+), we presented alternating 10s trials of 100% coherent motion moving in one of four directions with 10s trials where the position of each dot was randomly replotted within the circular aperture on every video frame (resembling 'snow' on a television set). The size of the stimulus aperture was the same as the one used in the main experiment. The subject's task was to press a button whenever the speed of the stimulus slowed briefly for 500ms; these target events occurred at three randomly determined intervals in each 10s trial. A GLM that contained a regressor corresponding to each stimulus type was used to identify hMT+ as the contiguous cluster of voxels lateral to the parietal-occipital sulcus that responded more during epochs of coherent motion than to the random-dot stimulus (single voxel threshold was set to $p<0.05$, corrected for multiple comparisons using the false discovery rate algorithm implemented in Brain Voyager). Bilateral regions of hMT+ were identified in 10/12 subjects; only left hMT+ was identified in one of the remaining subjects, and only right hMT+ was identified in the other.

Results

Figure 1 shows a schematic of the four alternative forced choice (4AFC) behavioral task subjects performed while in the scanner. This task was relatively easy when subjects were cued to report the direction of the high-coherence dot field, and relatively hard when they were cued to report the direction of the low-coherence dot field (termed *easy* and *hard* trials, respectively). Importantly, high and low coherence dot fields were simultaneously present on every trial, so the sensory properties of the display were fixed with respect to the total amount of coherent motion. This feature of the design was introduced to avoid simultaneously manipulating sensory factors (i.e., the motion coherence level) and perceptual difficulty. The subject was free to make a response at any point during the trial to indicate the direction of the currently relevant dot field; a saccadic response was required on one-half of the runs, and a manual button-press response was required on the remaining runs. On saccadic-response runs, subjects were required to maintain central fixation until the response was executed; on manual-response runs, central fixation was maintained throughout the trial.

By requiring subjects to use different output response modalities, we were able to search for supramodal signals related to decision making; the observation of this type of signal would support the existence of modality-independent decision variables (e.g., Heekeren et al., 2006). To identify such regions, we used the LBA model (see *Methods*) to make inferences from the behavioral data about how manipulations of perceptual difficulty should influence the BOLD signal originating from areas that play a role in accumulating sensory evidence during decision making. Importantly, these modeling efforts also dissociated cortical regions involved in perceptual decision making from

those more generally involved in attentional processes (i.e., general arousal, task demands, etc.).

Behavioral Results

Separate two-way repeated measures ANOVAs with response modality (saccade vs. manual) and perceptual difficulty (easy vs. hard) were used to assess the accuracy and response time (RT) data collected during the scanning session (see Table 1 for a summary of the group data). Subjects were slightly more accurate when making manual compared to saccadic responses ($F(1,11)=6.2$, $p=0.03$), and there was a robust main effect of perceptual difficulty on accuracy, indicating that deciphering the direction of a low-coherence stimulus on a hard trial was more challenging than deciphering a high-coherence stimulus on an easy trial ($F(1,11)=69.4$, $p<0.001$). Finally, there was no interaction between response modality and discrimination difficulty, indicating that manipulations of perceptual difficulty had a similar influence on both saccade- and manual-response accuracy ($F(1,11)=0.13$, $p=0.73$).

RTs were shorter on saccade trials compared to manual trials, but this effect did not reach significance ($F(1,11)=2.7$, $p=0.13$). RTs were reliably shorter on easy trials compared to hard trials ($F(1,11)=64.6$, $p<0.001$), and there was no interaction between perceptual difficulty and response modality ($F(1,11)=0.14$, $p=0.72$).

Linear ballistic accumulator model of behavioral data

Before analyzing the BOLD fMRI data, we fit our behavioral data using the LBA model (Brown and Heathcote, 2008). The goal was to investigate how manipulations of perceptual difficulty and response modality affected RT distributions. For instance, RTs might be faster on easy compared to hard trials because (1) a change in the rate with

which sensory evidence from the display was accumulated (termed the *drift rate* in the model), or (2) a change in the amount of evidence required to make a decision (termed the *response threshold*), or (3) both. Analysis using a cognitive model allows us to tease apart these separate influences, and to estimate parameters associated with each. By establishing which parameters changed with experimental manipulations, we can then estimate the pattern of BOLD responses expected from a region that is involved in accumulating sensory evidence during the decision process.

We report here fits to data averaged over participants, for simplicity of exposition. However, we repeated the same analyses separately for each individual participant, and obtained broadly similar results (see below). The data were split into 4 within-subject conditions, defined by two factors: response modality (saccade vs. manual) and stimulus coherence (easy vs. hard). For simplicity we collapsed across motion direction (upper left, upper right, lower right, lower left); however, we obtained qualitatively similar results if we included the four motion directions in the analysis to bring the total number of within-subject conditions to 16.

For a single decision condition, the LBA model as described above has five free parameters: t_0, A, b, s, and d_c, but it is not reasonable that all five of these should be estimated separately for all 4 conditions (easy vs. hard, saccade vs. manual). Instead, we fit the LBA model to the data 28 times, using different designs for constraining the parameters. Each design reflects a particular set of psychological assumptions regarding the way our experimental manipulations influenced cognitive processing. For example, the simplest model used a single set of five parameter estimates for all conditions, reflecting the assumption that the data were completely unaffected by the experimental

manipulations. Other designs allowed drift rates (d_c) to be different for easy vs. hard stimuli, or for manual vs. saccadic responses, and so on. We compared the adequacy of all possible designs using the Bayesian Information Criterion (BIC, Schwarz, 1978). The best design, with BIC=18,784.26, used constant values of s=0.227/sec (standard deviation) and A=0.849 (start point parameter) across all conditions. However, the design used higher mean drift rates on easy versus hard trials (d_c,*easy*=0.739 /sec, d_c,*hard*=0. 517/sec, with equal drift rates across modalities), and smaller non-decision times for saccadic responses ($t_{0,s}$=0.053 sec) compared to manual responses ($t_{0,m}$=0.134 sec), likely reflecting the modestly faster movement execution times for saccades. The model also assumed that the response threshold was slightly lower – that is, less cautious – for saccadic than manual responses ($b_{saccade}$=1.212, b_{manual}=1.278). Figure 4 illustrates the observed RT distributions (histograms) along with the predictions from the LBA model (solid lines). The top row shows distributions from high coherence conditions, the bottom row for low coherence conditions. The first two columns show data from trials with saccadic responses, the next two show data from trials with manual responses. The same y-axis scale was used for all histograms, so the heights of the distributions illustrate the relative probabilities of the responses – e.g. there are many more correct than incorrect responses for high coherence trials, so both the observed and predicted distributions are much taller for correct responses. The distributions predicted by the LBA are those corresponding to the best-BIC design described above.

We obtained similar results when we repeated the above analyses separately for each individual subject, although results were more variable due to the smaller sample sizes involved. Most importantly, the best-BIC model from the group data analysis

performed well across the individual subjects. That model had the third best mean BIC score out of the 28 models we tested; the model with the best mean BIC score was identical except with the added constraint that even t_0 should be constant across all conditions (not a surprising outcome given that the BIC tends to favor simpler models for smaller sample sizes). Patterns observed in the mean parameter estimates across individual participants also closely matched those obtained from the group data: s=0.182/sec, A=0.791, $d_{c,easy}$=0.688 /sec, $d_{c,hard}$=0.521/sec), $t_{0,s}$=0.120 sec, $t_{0,m}$=0.203 sec, b_{saccade}=0.993 and b_{manual}=1.078.

Supramodal mechanisms of information accumulation

To identify *candidate* regions that might be involved in perceptual decision making, we first performed a random effects analysis on data from 11/12 subjects to identify cortical areas exhibiting a two-way interaction between perceptual difficulty (easy vs. hard) and time (0s-16s in 2s intervals); this interaction was used to target areas that had a larger and temporally extended response on hard trials compared to easy trials (as in Figure 3b). The timeseries of the response on hard and easy trials was then computed from each ROI in the 12th subject; this leave-one subject out procedure was then repeated so that each subject was left out in turn (the permutation analysis was performed to avoid biasing a subsequent evaluation of response onset latency, see below and *Methods*). We collapsed across response modality because estimated drift rates did not vary between saccade and manual response conditions, and therefore our simulation predicted an identical BOLD response profile on hard compared to easy trials for both response modalities (see Figures 3a,b). This analysis identified regions in the right insula, bilateral intraparietal sulcus (IPS), bilateral frontal eye fields (FEF), a region of medial

frontal cortex (MFC), right inferior frontal gyrus (IFG, just anterior to the insula), right superior frontal gyrus (SFG), and left temporal parietal junction (TPJ, see Figure 5 and Tables 2-3). We also identified a ROI in left superior frontal sulcus (SFS) on 10/12 permutations of leaving one subject out; however, the interaction between perceptual difficulty and time did not reach significance in this region when evaluated in the left-out subjects (see Table 3). In all of the regions identified, the interaction between perceptual difficulty and time was driven by a larger and temporally extended response on hard trials compared to easy trials, with the exception of the left SFS and the left TPJ (a description of these regions is presented in the Discussion and Figure 10).

While a larger and temporally extended response on hard compared to easy trials is consistent with the accumulation of sensory evidence during perceptual decision making, similar effects of perceptual difficulty would also arise from areas involved in maintaining sustained attention or arousal during the decision process (Figures 3c-f). Therefore, we next used data from only the 'left-out' subjects to evaluate the latency of the BOLD response on easy and hard trials in each ROI; a delayed onset on hard trials is a distinguishing characteristic of a neural accumulator (see shaded region in Figure 3b). To test for latency differences, we performed a two-way repeated measures ANOVA with perceptual difficulty and time as factors, but this time we only included data from the first two timepoints of the BOLD response (0s-2s poststimulus). Note that the use of a leave-one-out procedure ensures that this second interaction test is independent from the criterion used to define each ROI. A subregion of the right insula was the only area where the onset of the BOLD response was delayed on hard trials for both response modalities (Table 4, Figure 6), making it a candidate for computing a supramodal decision variable

that might mediate activity in effector-specific regions of sensorimotor cortex. Moreover, three-way repeated measures ANOVA with perceptual difficulty, time (0s-2s), and ROI as factors revealed that the difference in the slope of the BOLD response on hard compared to easy trials was larger in right insula than in any of the other regions (all *F's*(1,11)>5.0, all *p*'s <.05; excluding data from the left SFS and left TPJ). Finally, to further explore the relationship between perceptual difficulty and BOLD response latency in the right insula, we divided RTs into three bins (collapsed across easy and hard trials) and found that the slope of the BOLD response across the first two timepoints decreased systematically with increasing RT (two-way repeated measures ANOVA with RT-bin and time as factors, $F(2,22)=5.11$, $p=0.015$, see Supplemental Figure 1).

In addition, the onset of the BOLD response was delayed on hard trials in bilateral regions of IPS on saccadic response trials (but not on manual response trials), as predicted by previous single-unit recording studies (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; differential effect of perceptual difficulty over the first two timepoints, $F(1,11)=12.3$, $p=0.005$, collapsed across right and left IPS; see Table 4 and Figure 7 for data from each hemisphere). However, no effect of perceptual difficulty on response latency was observed in the FEF on saccade response trials (Table 4, Figure 7).

Modality-dependent accumulator region for manual responses

Although the BOLD response in IPS was temporally delayed on hard saccade trials (See Table 4 and Figure 7), no corresponding modality dependent accumulator region was found on manual response trials. Therefore, based on previous reports (e.g. Meier et al., 2008), we used a two-way ANOVA and a leave-one-out procedure to identify a cluster of voxels in the superior aspect of the left central sulcus that responded

more robustly on manual response trials than on saccade response trials (interaction between response modality and time, $F(8,88)=14.2$, $p<.001$, mean Talairach coordinates: -35, -23, 54, ±1S.E.M. across permutations: 0.6, 0.8, 0.5; mean Volume: 5.2mL, ±1STD 0.637mL, see Figure 8). This region showed a larger and temporally extended response on hard manual trials compared to easy manual trials ($F(8,88)=2.9$, $p<.01$, Figure 8). Moreover, the onset of the BOLD response was delayed on hard trials when manual responses were required, meeting the second requirement for a modality specific neural accumulator (differential effect of perceptual difficulty across the first two timepoints when only considering manual response trials, $F(1,11)=8.0$, $p<0.025$). No such effects were found on saccade response trials (interaction between perceptual difficulty and time across all timepoints: $F(8,88)=1.0$, *n.s.*; interaction between perceptual difficulty and time across only the first two timepoints: $F(1,11)=1.1$, *n.s.*).

Activation profile in motion selective area hMT+

Single unit recording studies have demonstrated that neurons within stimulus-specific regions in early visual cortex – such as area MT for motion – signal the amount of sensory evidence present in the visual field (Newsome et al., 1989; Salzman et al., 1992; Britten et al., 1993; Britten et al., 1996; Ditterich et al., 2003). However, such regions do not integrate sensory evidence over time, suggesting that they primarily function to provide input to sensorimotor regions that are more directly involved in decision making (Roitman and Shadlen, 2002; Romo and Salinas, 2003; Huk and Shadlen, 2005; Hanks et al., 2006; Gold and Shadlen, 2007). If this account applies to hMT+ as well, then we predict a larger and temporally extended BOLD response on hard compared to easy trials because the sensory evidence on hard trials must be

represented for a longer period of time. However, no shift in the latency of activation onset is predicted because the underlying neural activity should be relatively constant for the duration of the stimulus presentation epoch (as opposed to ramping activity, as shown in Figure 3a). We tested this prediction by examining the BOLD activation profile within independently localized regions of hMT+ (see *Methods*). There was a significant interaction between perceptual difficulty and time (from 0s-16s), indicating a larger and temporally extended response on hard trials ($F(8,88)=3.8$, $p<0.005$, collapsed across right and left MT). However, there was no interaction between perceptual difficulty and time over the first two timepoints of the responses, suggesting that onset latency was similar on hard and easy trials ($F(1,11)=0.2$, *n.s.*). These results are consistent with the notion that hMT+ primarily plays a role in relaying information about sensory properties of the display to higher order accumulation centers (see Supplemental Figure 2 for a graphical depiction of the BOLD timecourses from left and right hMT+).

Discussion

Here we examined the neural mechanisms of perceptual decision making using a simple 4AFC task that controlled for sensory factors and a model that allowed us to predict the BOLD activation profile expected from cortical areas that accumulate sensory evidence (Figures 2-3). While the BOLD response in many regions increased with increasing perceptual difficulty, only a subset of these regions exhibited the latency offset predicted for a region involved in accumulating sensory evidence. Of these, only the right insula displayed this characteristic response profile for both tested response modalities. This finding raises the possibility that perceptual decisions are not solely computed by the same neural mechanisms that mediate the ultimate motor response. Instead, the

ramping-up of neural activity in sensorimotor regions such as the LIP may also reflect input from downstream regions that compute an abstract decision variable. Note that this account still allows for a causal influence of sensorimotor areas on decision making (e.g. Romo et al., 2002; Hanks et al., 2006). However, such regions may not be the actual *site* of the decision process, but instead might serve as 'relay stations' that translate abstract decision signals into an appropriate motor response (see Table 4, Figures 7,8). As would be the case with any correlational method, the evidence we provide here in support of this hypothesis is tentative; additional work using converging methodologies will be required to clarify the role of the modality-independent signals that we observed in the right insula.

An alternative account of the temporally delayed onset of the BOLD response in the right insula holds that neural activity might briefly pulse (an impulse response) at a slightly later time on hard compared to easy trials, perhaps signaling the termination of the decision process. For example, de Lafuente and Romo (2005) demonstrated that neurons in the medial prefrontal cortex signal the production of a "yes" response in an all-or-none fashion, such that the amplitude of the response does not correlate with the difficulty of the perceptual decision (in the context of a detection task). However, our data are inconsistent with this type of all-or-none termination signal because if the two temporally shifted impulse responses were equal in amplitude (Figures 9a,b), then we should not see a larger and temporally extended BOLD response on hard compared to easy trials (which we observe, Table 3 and Figures 6-8). On the other hand, if the impulse response on hard trials is temporally delayed *and* larger (Figures 9c,d), then we would expect to see a BOLD response pattern that is similar to the ramping 'accumulator' model

shown in Figures 3a,b. This second hypothesis is not suggested by any data that we are aware of, but one ad-hoc account is that the amplitude of the impulse response is somehow tied to the height of the decision boundary. However, the LBA model we employed estimated that the decision boundary was the same on easy and hard trials, arguing against this hypothesis (i.e. only drift rate differed). In any case, the pattern of activity depicted in Figures 9c,d also implies an important functional role for the right insula as it indicates sensitivity to both the difficulty and the timing of a perceptual decision.

In contrast to the predictions generated by our simulation (Figure 3), at least two previous studies asserted that the magnitude of the BOLD response should be higher on easy trials compared to hard trials because more sensory evidence is present on easy trials. Based on this criterion, Heekeren and coworkers highlighted a region of posterior left SFS/DLPFC as being important for perceptual decision making (Heekeren et al., 2004; Heekeren et al., 2006). Although we identified a region of the left SFS in 10/12 permutations of leaving-one-subject out that tended to respond more on easy than on hard trials, the effect was not significant (see Table 3 and Figure 10). In addition, we also identified a region of the left TPJ – similar to an inferior parietal lobe activation reported by Heekeren et al. (2006) – that responded more on easy trials compared to hard trials (Table 3 and Figure 10). Interestingly, both the left SFS and the left TPJ showed negative response profiles in our experiment (compared to the fixation baseline), with relatively smaller negative deflections on easy compared to hard trials (Figure 10). Thus, in our study at least, the left SFS and TPJ regions do not appear to follow an activation profile

that is consistent with the active accumulation of sensory evidence (i.e. the pattern shown in Figure 3b).

A similar pattern of deactivations was also reported by Tosoni et al. (2008), and we (along with Tosoni et al.) speculate that these regions are functionally related to the 'default' network that is actively suppressed during the performance of a demanding task; this suppression should be longer on hard trials because subjects spend more time trying to discriminate the direction of the target (Greicius et al., 2003; Shulman et al., 2003; Raichle and Snyder, 2007; Buckner et al.,2008; see also Tosoni et al. 2008 for further discussion of this point). Tosoni and coworkers also proposed that activation levels in putative accumulator areas should increase with increasing sensory evidence, contrary to our model simulations. In their study, the primary focus was on identifying regions of parietal and frontal cortex that mediate modality dependent responses (saccade and pointing movements) to arbitrary images (faces and houses); they found that modality sensitive subregions of parietal cortex responded more strongly on easy trials. At first glance, this observation appears at odds with the data we present here which shows larger responses on hard trials when the sensory evidence is weaker. However, because Tosoni et al. (2008) wanted to separate 'sensory' from 'motor' contributions to the BOLD signal, they had subjects delay their decision for 10.5 s following the presentation of the stimulus while awaiting a "go" signal. Since this delay interval is longer than required by the decision process, it is possible that subjects were storing a modality dependent representation of their planned response for much of the trial. Given that the computation of the response occurs more quickly when ample sensory evidence is present, the process of storing the prepared motor response for a longer period of time might have contributed

to increases in activation on easy trials. In contrast, our subjects were required to make speeded perceptual decisions, and thus had little time to engage in cognitive processes not directly related to perceptual decision making. Clearly more work needs to be done to resolve this issue, perhaps by combining Tosoni's methods for precisely mapping manual and saccadic sensitive regions with a task that constrains the cognitive operations subjects engage in during the 'decision making' stage of the task.

Even though we focus on the role of right insula in perceptual decision making, we cannot rule out the possibility that other regions are also involved in accumulating sensory evidence across multiple response modalities. Indeed, the interpretation of activation patterns in other areas is difficult: a larger response on hard compared to easy trials in the absence of a latency shift is equally consistent with a role in general attentional control or a lack of statistical sensitivity to detect a true difference in onset latencies. Therefore, we withhold speculation about other regions in anticipation of future studies that will selectively target candidate areas with converging methodologies to further delineate their role in perceptual decision making. Similar regions of insula have been previously implicated in different aspects of perceptual decision making. Trial-by-trial fluctuations in the left insula predict decisions about near-threshold fearful and non-fearful faces (Pessoa and Padmala, 2005, 2007), even when the sensory evidence is ambiguous and thus equated (Thielscher and Pessoa, 2007). Activation levels in bilateral regions of the anterior insula scale with the amount of differential sensory evidence during vibrotactile decision making (Pleger et al., 2006), increase at the moment of a perceptual decision in an image recognition task (Ploran et al., 2007), and correlate with a non-monotonic RT function during an auditory discrimination task, implying a role in the

decision process as opposed to sensory processing (Binder et al., 2004). Finally, activation levels in insular regions also scale with the amount of 'uncertainty' a subject experiences while discriminating a stimulus, suggesting a role in the process of comparing sensory evidence to a decision criterion (Grinband et al., 2006).

In contrast, other investigators have suggested that insular regions participate in attentional control precisely because more activation is observed on hard compared to easy tasks (Heekeren et al., 2006; Philiastides et al., 2006; Philiastides and Sajda, 2007; Heekeren et al., 2008; Tosoni et al., 2008). However, our simulation (Figures 3a,b) predicts a qualitatively distinct activation profile in 'decision making' areas compared to 'attention' areas, and the profile we observe in the right insula is more consistent with the former. We therefore argue that the present results support the hypothesis that the right insula is involved in coding an abstract decision variable capable of guiding the build-up of activity in effector-specific regions of sensorimotor cortex.

Ultimately, the extent to which regions outside of sensorimotor cortex participate directly in computing perceptual decisions may turn out to depend on the amount of training and the complexity of the task. For example, most single-unit recording studies employ 2AFC paradigms that involve highly stereotyped stimulus-response pairings that are practiced many thousands of times over many months (but see Churchland et al., 2008 for a more complex 4AFC task). In these tasks, making a perceptual decision is tantamount to selecting a motor response, so it is perhaps not surprising that the empirical evidence is consistent with the hypothesis that perceptual decisions are directly computed by sensorimotor neurons. However, in many everyday situations, a combination of motor responses must be issued in response to a single stimulus. For example, when driving, a

red light should motivate both a saccade towards the car immediately in front of you as well as a signal to depress the brake pedal. If perceptual decisions are solely computed and executed by the same mechanisms that mediate the motor response(s), then multiple systems – one for each response modality – must accumulate sensory evidence, translate the evidence into a decision based on current behavioral goals, and then generate two distinct motor responses. An alternative account, and one that is consistent with the present results, holds that a single modality-independent representation of the decision variable is computed and that this representation can then be used to efficiently guide multiple motor responses.

Chapter 1, in full, is a reprint of the material as it appears in Domain General Mechanisms of Perceptual Decision Making in Human Cortex in *Journal of Neuroscience, 29 (27)*, 8675-8687. Ho, T.C., Brown, S.D., Serences, J.T. (2009). The dissertation author was the primary author of this paper.

Table 1.1 Behavioral accuracy and response times on correct trials during the fMRI experiment for each condition and for the main effect of perceptual difficulty (easy vs. hard) and the main effect of response modality (manual vs. saccade). Means ± S.E.M.

Condition	Accuracy	Response Time (ms)
Manual (Easy)	93±2	1184±93
Manual (Hard)	66±5	1469±106
Saccade (Easy)	89±3	1035±94
Saccade (Hard)	61±5	1304±124
Easy (Manual + Saccade)	91.0±2	1110±82
Hard (Manual + Saccade)	64±5	1387±102
Manual (Easy + Hard)	79± 3	1327±98
Saccade (Easy + Hard)	75± 4	1170±109

Table 1.2. Anatomical location and volume of areas defined as showing an interaction between perceptual difficulty and time (based on a leave-one-subject-out analysis, see *Methods*). All coordinates from the atlas of Talairach and Tournoux (1988). The location of all regions was extremely consistent across permutations of leaving one subject out; the only exception was the left SFS activation (which was identified in only 10/12 permutations and moved considerably from one permutation to the next).

Region	X, Y, Z	Std. X,Y,Z	Vol. (mL)	Std. Vol	N
RH Insula	41,7,5	4,3,6	0.34	0.27	12
RH IFG	31,17,9	1,1,1	0.67	0.29	12
RH MFG	44,7,27	2,3,5	0.18	0.82	12
RH SFG	30,44,22	1,1,1	1.04	0.48	12
RH IPS	20,-70,34	3,3,5	3.22	1.4	12
RH FEF	26,-4,51	2,3,3	0.75	0.69	12
LH IPS	-22,-74,27	2,2,3	4.13	1.89	12
LH FEF	-27,-2,55	1,1,1	1.59	0.77	12
MFC	2,12,44	1,1,1	2.56	1.24	12
LH-TPJ	-43,-57,32	1,1,1	0.95	0.58	12
LH-SFS	-22,20,35	10,12,19	0.16	0.19	10

Table 1.3. All statistical tests computed by accumulating the data from each subject based on ROIs that were defined based on data from the remaining subjects (and then permuting this leave-one out procedure 12 times, see *Methods*). Of all the regions identified, only the left SFS did not show a significant effect across permutations of leaving one subject out. All numbers in table represent *F*-values with (8,88) degrees of freedom, with the exception of values for left SFS, where there are (8,72) degrees of freedom (because this region was identified in only 10/12 permutations). The three-way interaction between perceptual difficulty, response modality, and time did not reach significance in any region. *p<0.025, **p<0.01, ***p<0.001.

Region	**Difficulty x Time, F(8,88)**	**Diff x Modality x Time, F(8,88)**
RH Insula	3.64***	0.98
RH AntInsula	5.00***	1.54
RH MFG	2.46*	0.97
RH SFG	3.46**	0.39
RH IPS	4.53***	1.59
RH FEF	2.66*	0.98
LH IPS	4.25***	1.80
LH FEF	5.80***	0.77
MFC	6.85***	0.87
LH-TPJ	4.95***	0.94
LH-SFS	0.32	0.60

Table 1.4. All statistical tests computed by accumulating the data from each subject based on ROIs that were defined based on data from the remaining subjects (and then permuting this procedure 12 times). The first column contains F-values for the interaction between perceptual difficulty and time over the first two data points (0s-2s), collapsed across response modality; also reported is the interaction computed separately for each response modality, and the three-way interaction between perceptual difficulty, response modality, and time. All numbers represent F-values with (1,11) degrees of freedom, with the exception of the values for left SFS, where there are (1,9) degrees of freedom. *p<0.05, **p<0.025, ***p<0.01

Region	**Diff x Time**	**Diff x Time (saccades)**	**Diff x Time (manual)**	**Diff x Modality x Time**
RH Insula	10.7***	5.30*	7.57**	0.28
RH AntInsula	0.15	2.04	2.06	3.98 (p=.07)
RH MFG	0.00	0.70	0.18	0.91
RH SFG	0.09	0.40	0.03	0.53
RH IPS	1.86	9.35**	0.24	4.15 (p=.06)
RH FEF	0.00	0.01	0.00	0.01
LH IPS	0.38	8.76**	0.11	3.46 (p=.09)
LH FEF	0.00	0.50	0.81	1.64
MFC	0.03	1.11	0.82	2.18
LH-TPJ	2.19	0.29	2.65	0.38
LH-SFS	1.47	0.30	0.64	.01

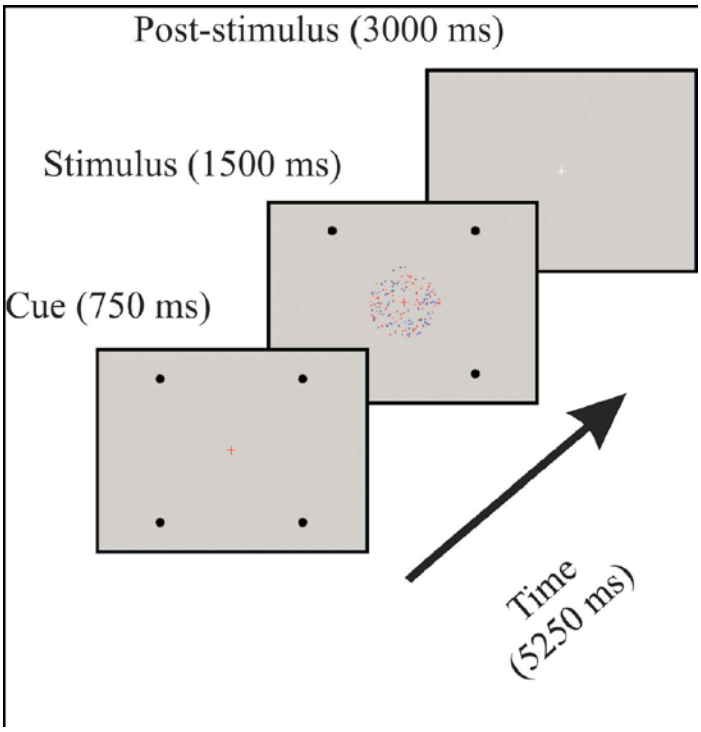

Figure 1.1 Behavioral paradigm. Subjects maintained fixation on the central fixation cross at the start of each trial; the color of the cross cued the subjects to decipher either the dots rendered in red or blue. On every trial, one dot field contained 40% coherent motion (*hard* stimulus) and the other contained 80% coherent motion (*easy* stimulus). On alternating runs, subjects indicated the direction of the relevant dot field with either a saccade to one of the four peripheral position markers or with a button press response. See *Methods* for additional details.

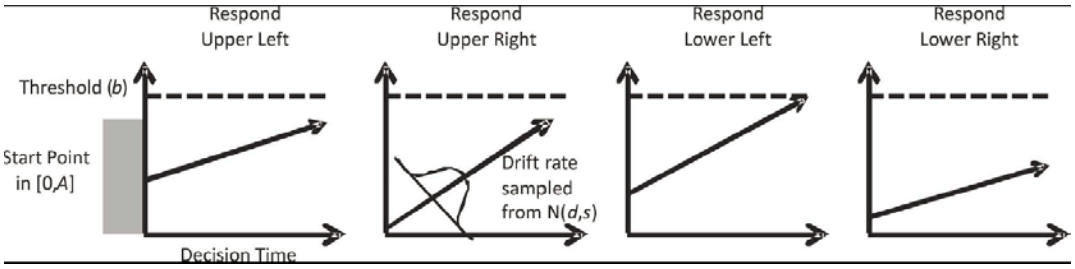

Figure 1.2. Schematic of the linear ballistic accumulator (LBA) model. The choice between the four responses is modeled as a race between four accumulators. Activation in each accumulator begins at a random point between zero and *A* and increases with time. The rate of increase is random from trial to trial, but is (on average) faster for the accumulator whose associated response matches the stimulus. A response is given by whichever accumulator first reaches the threshold *b*, and the predicted response time depends on the time taken to reach that threshold.

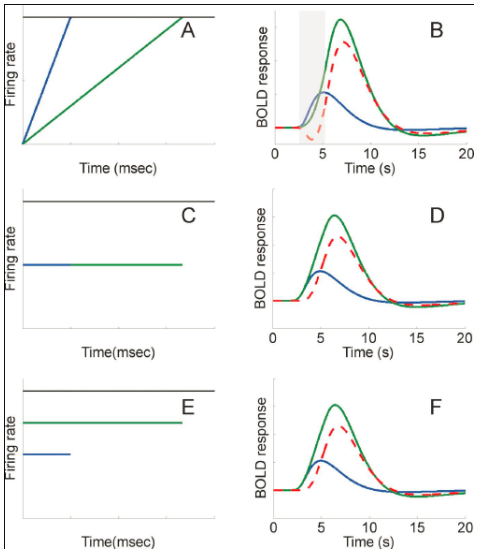

Figure 1.3. Simulated BOLD activation profile in a region involved in accumulating sensory evidence. Neurons in areas such as LIP are known to increase their firing rates to a bound during perceptual decision making tasks; the time taken to reach the bound is determined by the quality of the sensory evidence (e.g. the motion coherence in a RDP, see Shadlen and Newsome, 2001). (**a**) Two hypothetical cases of a fast (blue) and a slow (green) decision process in a sensory integration area like LIP. The blue trace might be expected on easy trials because sensory evidence is abundant, the green trace might be expected on hard trials because sensory evidence is sparse. (**b**) Predicted pattern of BOLD responses associated with each hypothetical case shown in panel (a), computed by convolving the simulated firing rate of the neuron depicted in (a) with a simulated BOLD response function (a 'double-gamma' function). Notice that a larger response is expected when the drift rate is slow because the BOLD response is proportional to the integrated amount of neural activity during the decision process. Moreover, the onset of the BOLD response is delayed when the drift rate is slow (shaded region), which is a distinguishing characteristic of a region involved in accumulating sensory evidence. Dashed red line represents the predicted response on hard trials minus the response on easy trials. (**c**) Hypothetical neural activity in a region that is involved in maintaining attention at a fixed level for the duration of the decision process. (**d**) BOLD response profiles expected on easy and hard trials given the neural profiles shown in (c). As in panel (b), a region involved in maintaining sustained attention should also exhibit a larger and temporally extended response, but without the corresponding offset in response latency. (**e**) Same as (c), but assumes a larger sustained response on hard compared to easy trials due to the increase in task difficulty. (**f**) BOLD response profiles expected on easy and hard trials given the neural profile in (e); again a larger and temporally extended response is predicted on hard trials, but without a shift in the latency of response onset. Note that the absolute scale of BOLD responses is not relevant for the present purposes; instead we focus on the qualitative pattern of the BOLD responses that should be associated with a region involved in accumulating sensory evidence.

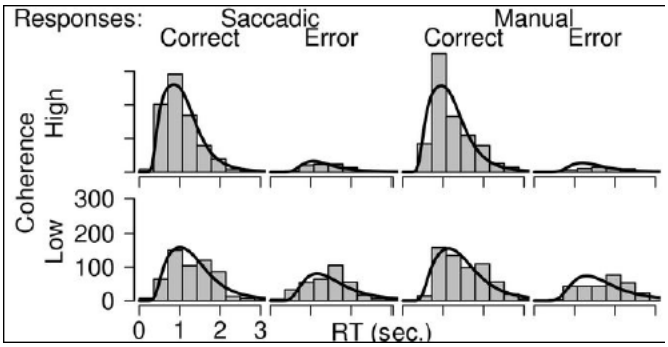

Figure 1.4. Fit of LBA model to the response time (RT) data. The histograms show observed RT distributions for correct and incorrect decisions. The top and bottom rows show distributions of RTs associated with decisions based on high coherence (easy) and low coherence (hard) stimuli, respectively. The left half of the figure shows data from saccadic responses, the right half from manual responses. Solid lines indicate the RT distributions predicted by the LBA model (see text).

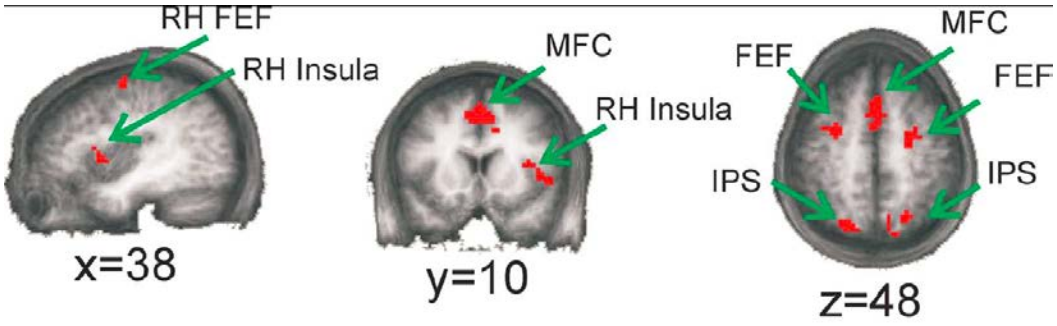

Figure 1.5. Regions exhibiting a larger and temporally extended response on hard compared to easy trials. These maps were generated by averaging the ROIs identified on each permutation of leaving-one-subject out while testing for an interaction between perceptual difficulty and time (0s-16s poststimulus), collapsed across response modality (see *Methods*, Tables 2,3 for statistical values associated with each region). These are candidate areas that may play a role in accumulating sensory evidence during decision making. All activations projected onto an average of the high-resolution anatomical scans from all subjects in our study (applies to Figure 8 as well).

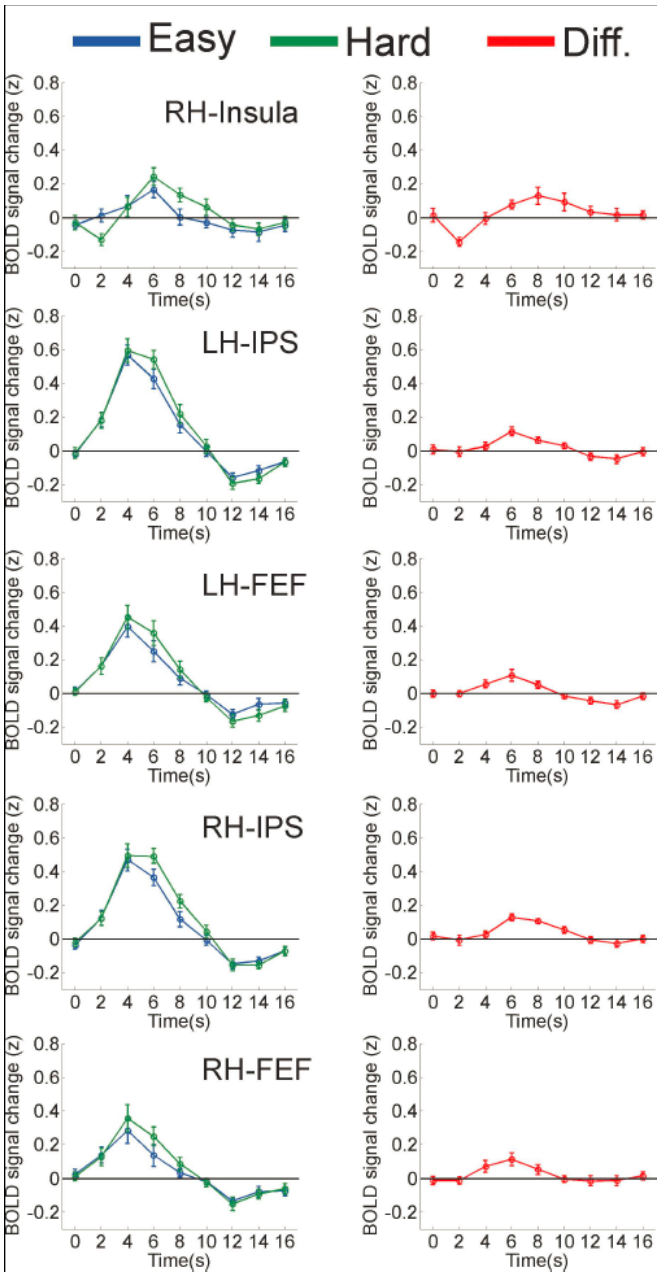

Figure 1.6. Mean timecourses across subjects in each ROI, collapsed across response modality. All timecourses based on data from the left-out-subject after ROIs were identified in 11/12 subjects. Figures in the right column depict the mean difference between responses associated with hard and easy trials (left column), and error bars are ±1S.E.M. across subjects. Note the relative delay in the onset of the BOLD response on hard trials in the right insula (Table 4); this delay is qualitatively similar to the pattern predicted in a region involved in accumulating sensory evidence during decision making (see Figure 3b), and inconsistent with the predicted response in a region involved in sustaining attention or arousal during task performance (Figures 3d,f).

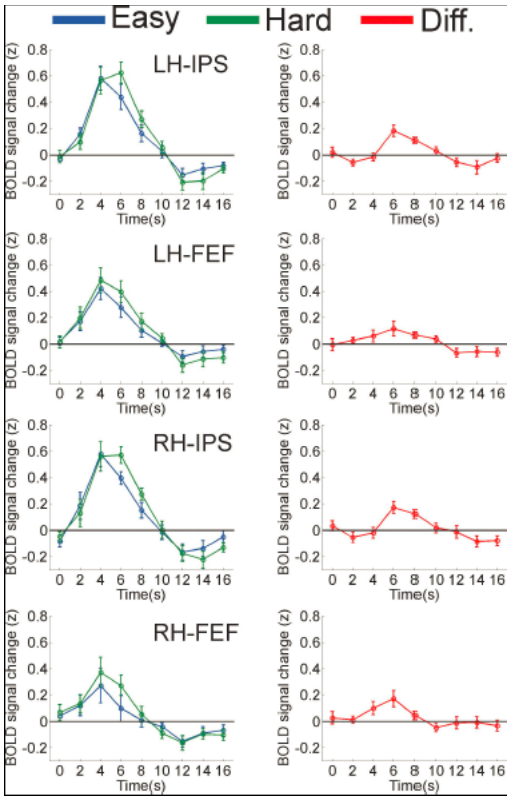

Figure 1.7. Mean timecourses across subjects on saccadic-response trials. Timeseries computed as described in Figure 6; note the temporal delay in the IPS ROIs on hard saccadic-response trials (see Table 4). Error bars ±1S.E.M. across subjects.

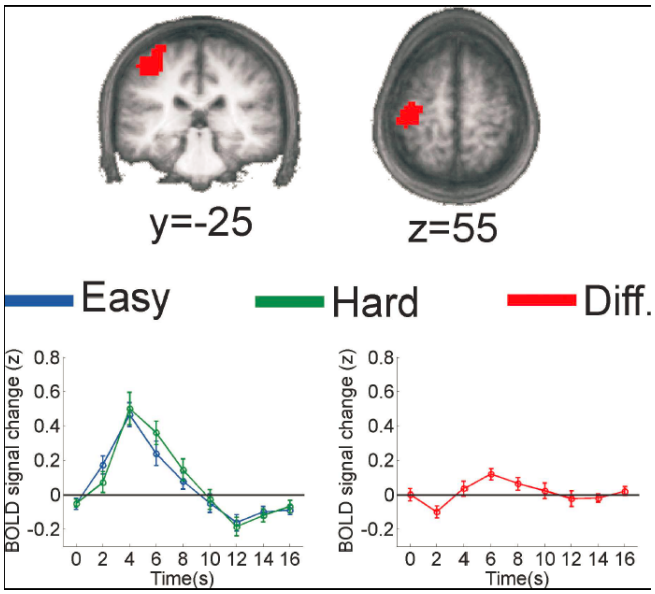

Figure 1.8. Sub-region of the superior left central sulcus that shows an interaction between response modality (saccade vs. manual) and time (0s-16s poststimulus). The depicted region of superior left central sulcus is the mean ROI averaged across all permutations of leaving-one subject out that responded more on manual compared to saccade trials (all subjects responded with their right hand, and a similar region of the central sulcus has been previously shown to respond when the contralateral fingers are moved, see e.g. Meier et al., 2008). Timecoures are based on data from the left-out subject and were computed using only manual-response trials (error bars ±1S.E.M. across subjects).

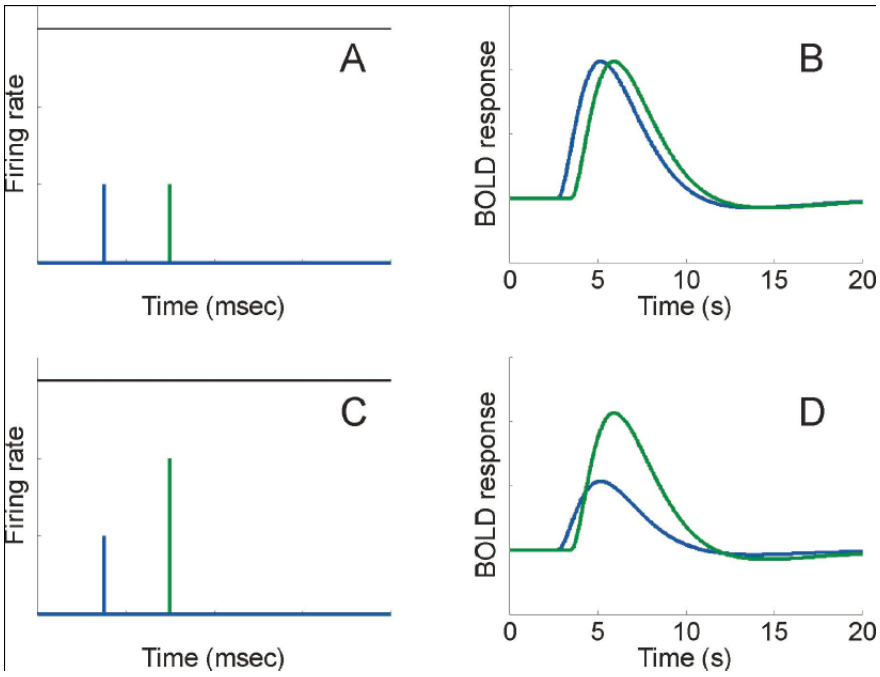

Figure 1.9. Simulated BOLD activation profiles in a region that issues a delayed impulse response on hard compared to easy trials. (**a**) A temporally delayed impulse response (brief response) of neural activity is issued later on hard trials than on easy trials, perhaps signaling the end of the decision process. (**b**) Predicted pattern of BOLD responses associated with each hypothetical case shown in panel (a). Although the response on hard trials is temporally delayed, it is not larger or temporally extended compared to the response on easy trials. (**c**) Same as (a), but now the impulse response associated with a hard trial is both temporally delayed and larger. (**d**) Predicted pattern of BOLD responses associated with each hypothetical case shown in panel (c). The predicted BOLD responses are similar to those predicted by ramping neural activity in an accumulator area during decision making (Figures 3,a,b).

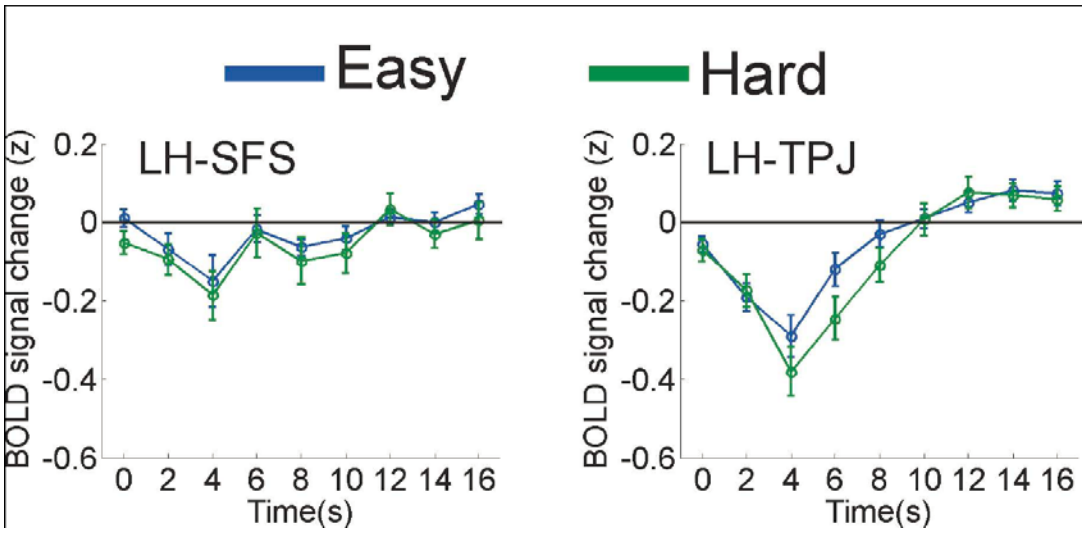

Figure 1.10. Regions of left SFS and left TPJ. Although the interaction between perceptual difficulty and time (0s-16s) was not significant in the left SFS (see Table 3), both of these areas show a 'deactivation' compared to the fixation baseline, and the deactivation is relatively attenuated on easy compared to hard trials.

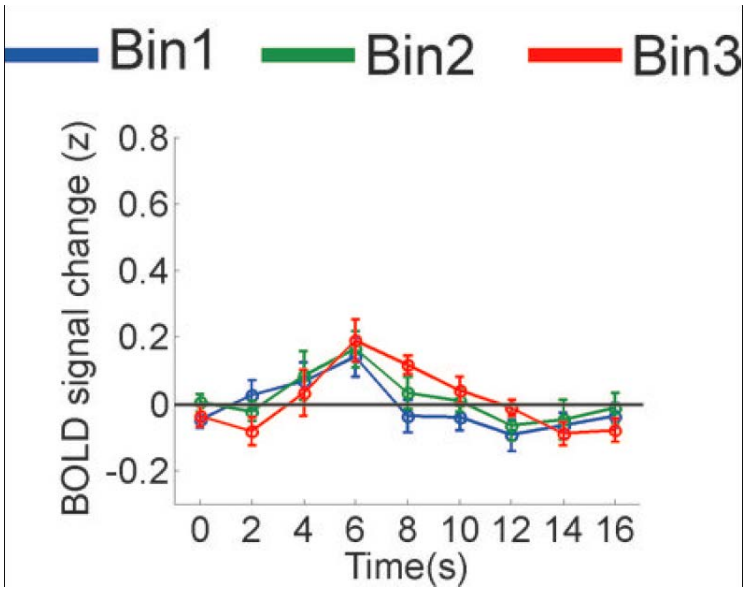

Supplemental Figure 1.1. Mean timecourses across subjects in right insula, broken down into three bins based on reaction time (collapsed across easy and hard trials). The onset of the BOLD response is progressively delayed with increasing RT. Error bars ±1S.E.M. across subjects.

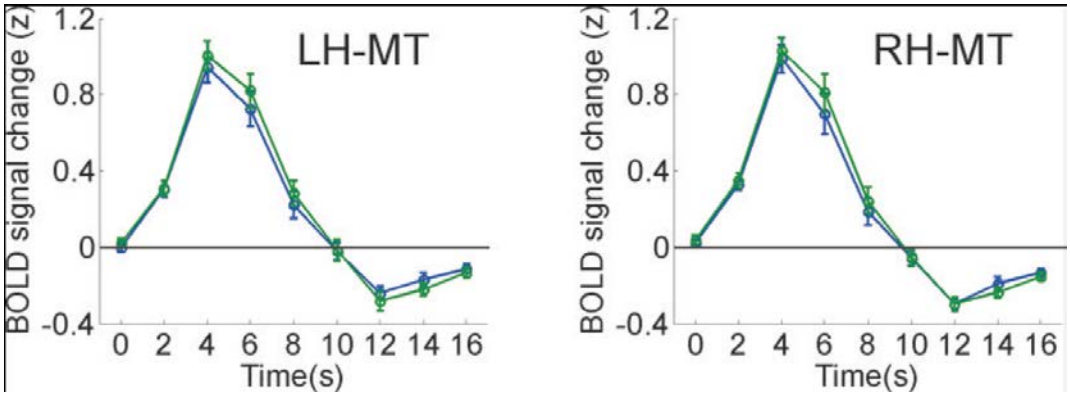

Supplemental Figure 1.2. Mean timecourses across subjects in left and right hMT+. Error bars ±1S.E.M. across subjects.

CHAPTER 2: The Optimality of Sensory Processing During the Speed-Accuracy Tradeoff

Ho T, Brown SD, Van Maanen L, Forstmann BU, Wagenmakers EJ, Serences JT (2012). *Journal of Neuroscience,* 32(23): 7992-8003.

ABSTRACT

When people make decisions quickly, accuracy suffers. Traditionally, speed-accuracy tradeoffs (SAT) 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 functional MRI (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.

The Optimality of Sensory Processing During the Speed-Accuracy Tradeoff

Fast decisions are typically more error prone, while precise decisions require more time, a phenomenon known as the speed-accuracy tradeoff (or 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 (Van Veen et al., 2008; Forstmann et al., 2008; Forstmann et al., 2010; 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 (see Figure 1 and *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; see Figure 2; Hol and Treue, 2001; Butts and Goldman, 2006;

Jazayeri and Movshon, 2006; Navalpakkam and Itti, 2007; Scolari and Serences, 2009; Moore, 2008; Purushothaman and Bradley, 2005; Regan and Beverley, 1985; Schoups et al., 2001). 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 model (LBA; Brown and Heathcote, 2008; see Figure 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 that the SAT may affect sensory processing (see also Vandekerckhove et al., 2011 and Hubner et al., 2010). We then used a forward encoding model (Brouwer and Heeger, 2009; 2011; reviewed in Naselaris et al., 2011; Serences and Saproo, 2011) 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

16 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 hour session in a climate and noise controlled subject room outside of the scanner and a single 1.5-2 hour session in the scanner. Compensation for participation was $10/hour for behavioral training and $20/hour 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; The Math Works, Natick, MA), 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 approximately 14°) at full contrast which flickered at 6 Hz (83.33 milliseconds on, 83.33 milliseconds 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 pseudo-random jitter added (between ± 0°-6°, selected from a uniform distribution). On half 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 counterbalanced across trials (see Figure 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 seconds, after which it was replaced with a white centrally presented fixation circle for 3.5 seconds. We omitted all trials in which a subject failed to give a response (less than 1%) or emitted a response faster than 200 milliseconds (less than 1.05%). Since the second grating did not appear until 166.67 milliseconds into the trial, we reasoned that responses quicker than 200 milliseconds 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 seconds including an 11 second 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. Prior to 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 lab, 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 prior to 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 approximately 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.

Linear ballistic accumulation (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 (see Brown and Heathcote, 2005; 2008; Usher and McClelland, 2001). 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 (*non-decision 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 standard deviation 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, 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; see also Ratcliff 1978; Ratcliff & Rouder, 1998; Forstmann et al., 2010; Van Maanen et al., 2011; Wolfe and Van Wert, 2010). 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 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 non-decision time (t_0)). To quantify that support, we approximated posterior model probabilities based on BIC assuming a fixed effect for subjects (see Burnham and

Anderson, 2002), which showed this design to be more than 10^{10} times more likely than the next best design (see *Results* section).

Single-trial linear ballistic accumulator (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 to not only 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, see Ratcliff, 1985; Ratcliff and Rouder, 1998), drift rates are normally distributed across trials, with different distributions for each respective accumulator. This assumption of normally distributed drift rates implies that drift rates which 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 (MLE) for a single-trial drift rate (d_i) and a single-trial starting point (a_i) given a trial with response time (t_i):

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

$$\hat{a}_i = \begin{cases} A & \text{if } t_i \leq \frac{b-A}{v} + t_0 \\ b - (t_i - t_0) \cdot v & \text{if } \frac{b-A}{v} + t_0 < t_i < \frac{b}{v} + t_0 \\ 0 & \text{if } t_i \geq \frac{b}{v} + t_0, \end{cases}$$

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 non-decision 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 (see the text surrounding Figure 3 in Van Maanen et al., 2011).

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 non-decision 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 (see Table 2 for exact values). Constraining the model in other reasonable ways (e.g., fixing the non-decision 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 carried out on a General Electric MR750 3T scanner equipped with an 8-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 1x1x1mm resolution. Whole brain echo planar functional images (EPI) were acquired in (for 8 of the subjects) or 26 (for the remaining

subjects) oblique slices (TR = 1500 ms, TE = 30 ms, flip angle = 90°, image matrix = 64 x 64, field of view = 192 mm, slice thickness = 3 mm, 0mm gap).

Data analysis was performed using BrainVoyager QX (v 1.91; Brain Innovation, Maastricht, The Netherlands) and custom timeseries analysis routines written in Matlab (version 7.11.0.584; The Math Works, Natick, Massachusetts). 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 timeseries. The timeseries 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 a voxel's composition or by its distance from the coil elements). We then estimated the magnitude of the BOLD response on each trial by shifting the timeseries from each voxel by four 1.5 second TRs (6 seconds) to account for the temporal lag in the hemodynamic response function, and then extracting data from the two consecutive 1.5 second TRs that correspond to the duration of each 3 second trial (see 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 *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 employed in the primary experiment. The localizer stimulus was comprised of a full-contrast counter-phase modulated (8Hz) 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 10s, and the contrast of the checkerboard was reduced by 30% for a single video frame at 6 pseudo-randomly 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 10s trial was then followed by 10s 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: 5s, time to peak of negative response: 15s, 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 was then projected onto a computationally inflated representation of each subject'ss 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 (Brouwer and Heeger, 2009, 2011; Dumoulin and Wandell, 2008; Gourtzelidis, et al., 2005; Kay and Gallant, 2009; Kay et a., 2008; Mitchell, et al., 2008; Naselaris, et al., 2009; Thirion, et al., 2006; reviewed in Naselaris, et al., 2011; Saproo and Serences, 2011). 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 9 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 x 9 orientations), and 72 observations for subjects that underwent 10 runs in the scanner (8 runs in the training set x 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 trials in the test set, and k be the number of hypothetical orientation channels. Let B_1 (m x n_1 matrix) be the training data set, and B_2 (m x 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 x n_1) comprised of nine half-sinusoidal functions raised to the 6th power as a basis set (see Figure 4). The training data in B_1 were then mapped onto the matrix of channel outputs (C_1) by the weight matrix (W, m x k) that was estimated using a GLM of the form:

$$\mathbf{B_1} = \mathbf{W C_1} \qquad (\mathbf{1})$$

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

$$\hat{\mathbf{W}} = \mathbf{B_1 C_1^T (C_1 C_1^T)^{-1}} \qquad (\mathbf{2})$$

The channel responses C_2 (k x n_2) were then estimated for the test data (B_2) using the weights estimated in (2):

$$\hat{\mathbf{C}}_2 = (\hat{\mathbf{W}}^T \hat{\mathbf{W}})^{-1} \hat{\mathbf{W}}^T \mathbf{B_2} \qquad (\mathbf{3})$$

The first steps in this sequence (equations 1-2) 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 3 then implements a multivariate computation because the channel responses estimated on each trial (in C2) 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 C1 has a large impact on the resulting channel response estimates. In the present experiment, we used half-sinusoidal functions that were raised to the 6th power to approximate the shape of single-unit tuning functions in V1, where the $1/\sqrt{2}$ half-bandwidth of orientation tuned cells is approximately 20° (although there is a considerable amount of variability in bandwidth, see Ringach et al., 2002a; Ringach et al., 2002b; Gur et al, 2005; Schiller 1976). Given that the half-sinusoids were raised to the 6th 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 5^{th}-8^{th} 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 re-centered 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 3s trial (see Figure 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 non-parametric bootstrapping/randomization procedure. Note that this bootstrapping/randomization procedure was used for all comparisons related to BOLD tuning functions (see Figures 6 and 7, AE v. SE, correct AE v. incorrect AE, correct SE v. incorrect SE, the interaction between AE v. SE based on accuracy, AE logistic regression beta weights v. SE logistic regression beta weights, and single-trial correlations between AE responses and drift rate v. 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 re-assigning 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 re-sampled and randomized data set using the same procedure applied to the observed data. This re-sampling 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

re-sampled 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 (RT) and accuracy results

Trials on which RTs were faster than 200 milliseconds were discarded in all subsequent analyses (including model fitting procedures described below, see *Materials and Methods*). Two-way repeated-measures analysis of variance (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 (see Table 1 for a summary of the group data). The task instructions produced a strong SAT effect: participants responded faster on SE trials compared to 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).

Linear Ballistic Accumulator (LBA) Model Results

Accuracy rates and RTs might be lower on SE trials compared to 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 (see Figure 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 (see Figure 2 and *Predictions* section below), we focused our analyses on mismatch trials (data from match and mismatch trials were nonetheless fit simultaneously, see *LBA model analyses* under *Materials and Methods* for more details). 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 non-decision 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 (see Figure 3 *and Linear Ballistic Accumulator mod*el and *LBA model analyses* under *Materials and Methods* for more details). 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 trade-off 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 non-decision time on AE trials compared to SE trials (see Table 2 for a summary of the parameter fits to data averaged over all the subjects). Based on approximated posterior model probabilities (see *LBA model analyses* under *Materials and Methods* for more details), this design was found to be more than 10^{10} times more likely than the next best design. Individually, this design was also the modal result: the BIC values for 6 out 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 which 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 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 Ratcliff and Tuerlinckx, 2002 or 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 less than 25 milliseconds.

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 where error rates were relatively high. However, since the decision model fits the response time distributions from both conditions very well (see Figure 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 (see Ratcliff and Tuerlinckx, 2002 and Donkin et al., 2011a, for details). With this built in assumption, the decision model fit the response time distributions from both conditions very well (Figure 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; see 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 employed 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 (see also Hubner et al., 2010 for a more theoretical treatment). Finally, the small observed differences in the time taken for non-decision processing between SE and AE conditions (see 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 (see Starns and Ratcliff, 2010; Voss et al., 2004).

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, in order 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 non-decision 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 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 where 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 where 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 (see Figure 2; Hol

and Treue, 2001; Butts and Goldman, 2006; Jazayeri and Movshon, 2006; Purushothaman and Bradley, 2005; Navalpakkam and Itti, 2007; Regan and Beverley, 1985; Schoups et al., 2001; Scolari and Serences, 2009; 2010). Hence, we focused our analyses on mismatch 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 (see Table 2), we predicted that correct mismatch AE trials should be associated with more activation in off-target neural populations compared to incorrect mismatch AE trials. The difference between the drift rates associated with the correct and incorrect accumulators on SE mismatch trials, on the other hand, is much smaller (see Table 2). We would therefore expect a small difference between off-target activation on correct SE mismatch trials compared to 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 (see Brouwer and Heeger, 2009, 2011; reviewed by: Naselaris et al., 2011; Serences and Saproo, 2011) to estimate how the SAT modulates orientation selective response profiles in V1 (see section entitled *Estimating feature-selective BOLD response profiles using a forward encoding model* under *Materials and Methods*). On mismatch trials, we first compared the BOLD-based orientation tuning functions (TFs) associated with AE trials with those associated with SE trials (Figure 6a) 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 non-parametric 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 ($p<0.01$, Figure 6b). In particular, responses in channels tuned approximately 25°-65° away from the target-orientation showed larger responses on correct trials compared to 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 population should increase the quality of the information being sent to downstream decision mechanisms and thus increase the probability of a correct response (see Figure 2c). In contrast, no differences were observed between channel responses associated with correct and incorrect SE trials ($p=0.90$, Figure 6c), 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 ($p<0.01$, Figure 6d). 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 (see Table 2).

As shown in Figure 6b, we observed more activation in off-target channels on correct trials compared to 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, see Figure 6b). Across subjects, larger

differences between correct and incorrect accumulator drift rates were positively correlated with larger differences in off-target activation on correct compared to incorrect AE trials (Figure 7a, R2=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 7a 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 (beta coefficient) indicates that higher activation in a given channel predicts a higher probability of a correct behavioral response; negative beta 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

approximately ~40°-60° from the target were associated with a higher probability of a correct response (Figure 7b). In contrast, the beta coefficients on SE trials fluctuated around zero. This pattern gave rise to a significant cross-over interaction between the AE and SE beta coefficient curves (p=0.021, Figure 7b). As with the increased off-target activation on correct AE trials (Figure 6b) and the corresponding relationship with the rate of sensory evidence accumulation on a between-subject basis (Figure 7a), 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 single-trial LBA (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 non-decision 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 to 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, Figure 7c), 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 (Figure 6b,d), between-subject correlation (Figure 7a), and within-subject logistic regression (Figure 7b), 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 7c 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 where the non-decision 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 7c thus illustrate 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 (Figure 7a), logistic regression revealed a trial-by-trial relationship between behavioral accuracy and BOLD activation levels in off-target orientation channels (Figure 7b), 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 (Figure 7c).

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 (compare Figures 6a and 6b). 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 (reviewed in: Corbetta and Shulman, 2002; Desimone and Duncan, 1995; Kastner and Ungerleider, 2000; Noudoost et al., 2010; Serences and Yantis, 2006; Yantis, 2008) 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., 2010), and that subregions of frontal cortex mediate perceptual decisions (Purcell et al., 2010; Gold and Shadlen, 2007; Heekeren et al., 2004; de Lafuente and Romo, 2005; 2006; Lemus et al., 2010; Hernandez et al., 2010; Ho et al., 2009; Kayser 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 if speed pressure selectively impairs a subject's 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 (Van Veen et al., 2008; Forstmann et al., 2008; Forstmann et al., 2010), consistent with a response threshold account that only motor and frontal areas are involved in mediating the SAT (Van Veen et al., 2008; Ivanoff et al., 2008; Forstmann et al., 2008, 2010; Wenzlaff et al, 2008; Ratcliff, 1985; Ratcliff and Rouder, 1998). 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 modulations 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 (see Ringach et al., 2002a; Ringach et al., 2002b; Gur et al, 2005; Schiller 1976). 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 (Figure 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 (Freeman et al., 2011; Mannion et al., 2010; Leventhal, 1983; Sasaki et al., 2006; Schall et al., 1986; Zhang et al., 2011). For instance, there is a radial orientation bias in V1 (Freeman et al., 2011; Sasaki et al., 2006; Zhang 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 (Freeman et al., 2011; Sasaki et al., 2006; Zhang 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 (Kamitani and Tong, 2005; Haynes and Rees, 2005; Swisher, et al., 2010; see Boynton, 2005 for a useful graphical illustration). 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; Heeger and Ress, 2002; Logothetis et al., 2001; Buxton, 2002; Logothetis and Wandell, 2004; Sirotin and Das 2009; Das and Sirotin, 2011; Handwerker and Bandettini, 2011a, 2011b; Jia et al., 2011; Kleinschmidt and Muller, 2010; Schummers et al., 2008). However, given that neurons in early sensory areas like V1 are massively interconnected (e.g., 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 and colleagues (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 (see also: Law and Gold, 2008; 2009; Purushothaman and Bradley, 2005; 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 decision 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.

Chapter 2, in full, is a reprint of the material as it appears in The Optimality of Sensory Processing During the Speed-Accuracy Tradeoff in *Journal of Neuroscience, 23(32)*, 7992-8003. Ho, T.C., Brown S.D., Van Maanen, L., Forstmann, B.U., Wagenmakers, E.J., Serences, J.T. (2012). The dissertation author was the primary author of this paper.

Table 2.1. 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). See *Results* for more details. Means ± 1S.E.M. across subjects.

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+Mismatch)	66.23±2.24	922.4±40.1
Accuracy (Match+Mismatch)	86.02±2.00	1570.8±82.6
Match (Speed+Accuracy)	80.69±2.40	1322.8±101.7
Mismatch (Speed+Accuracy)	71.57±2.99	1170.4±73.7

Table 2.2. Average LBA parameter estimates for the best BIC model on mismatch and match trials. 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 (see *LBA results* for more details). 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 where the stimulus gratings match, and also the accumulator corresponding to the "mismatch" response on trials where 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 to speed-emphasis trials (last row in the table).

LBA parameters	**Speed (Mismatch)**	**Accuracy (Mismatch)**	**Speed (Match)**	**Accuracy (Match)**
Starting point (*A*)	2.31	2.31	2.31	2.31
Non-decision 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_e)	1.71	0.69	1.19	-0.07
Differential drift rate (v_c - v_e)	0.96	2.31	1.56	2.65

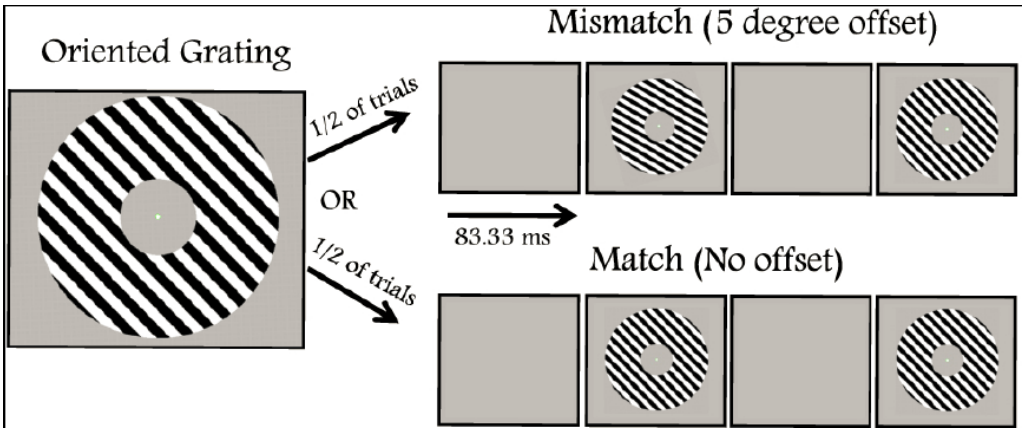

Figure 2.1. Behavioral paradigm. Subjects were presented with an oriented grating at full contrast flickering at 6 Hz (on for 83.33 milliseconds, off for 83.33 millisecond, etc. for a total of 3 seconds). On each trial, the target orientation of the grating was selected from one of nine possible orientations (0°, 20°, 40°, 60°, 80°, 100°, 120°, 140°, 160°), plus or minus an offset randomly selected between 0° and 6°. On half 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. See *Materials and Methods* for additional details.

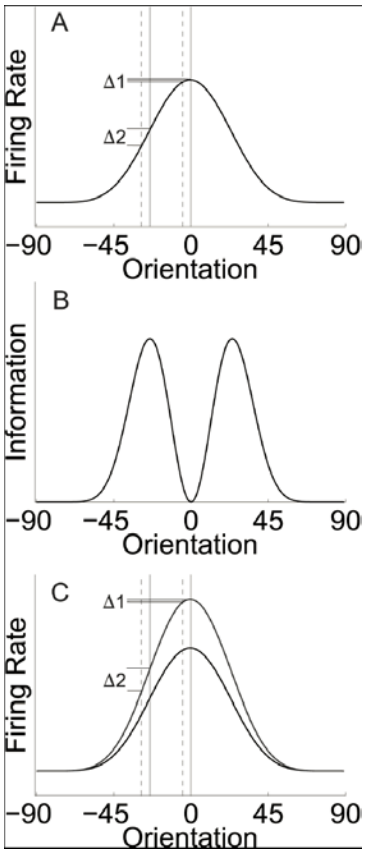

Figure 2.2. Model of optimality for fine perceptual discrimination tasks. (a) 'On-target' neurons (0° offset on the abscissa axis) exhibit small changes in firing rates (Δ1) in response to two similar stimuli (denoted by the vertical solid line and the vertical dashed line). Off-target neurons, on the other hand, undergo larger changes in firing rates in response to similarly spaced stimuli (as denoted by Δ 2). (**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 panel (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.

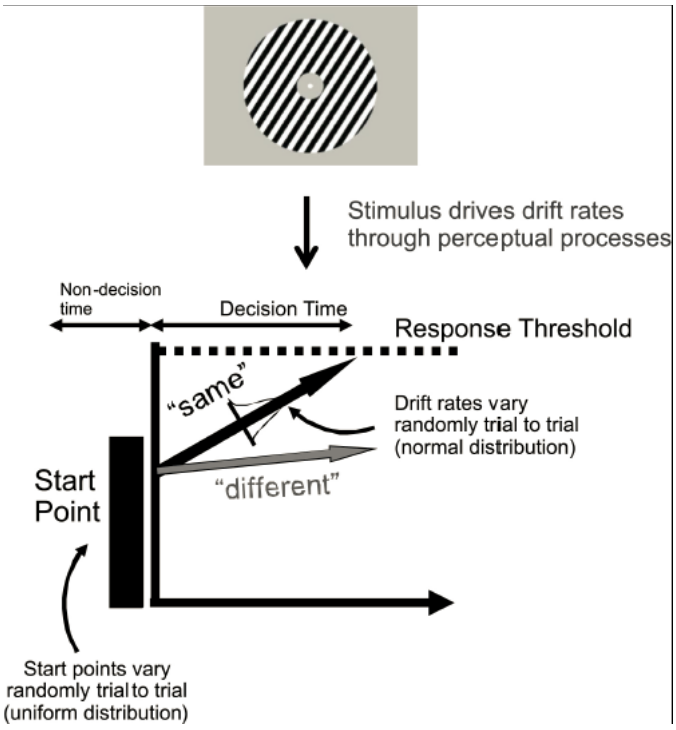

Figure 2.3. Schematic of the linear ballistic accumulator (LBA) model. Schematic illustration of the LBA model as applied to the orientation discrimination task. The stimulus grating (top) provides information to two racing accumulators; the first accumulator to reach threshold determines the response, 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 non-decision time).

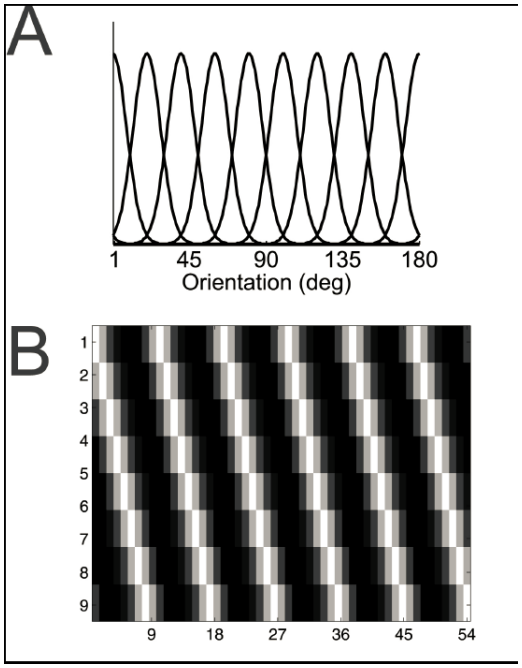

Figure 2.4. Schematic of the forward encoding model. (**a**) Basis set comprised of nine half-sinusoidal functions raised to the 6^{th} power: the functions are evenly distributed across orientation space. (**b**) Depiction of the design matrix C_1. There is one column corresponding to each observation in the training data matrix, and each row represents the response profile of one of the 9 half-sinusoidal basis functions shown in panel (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 9 hypothetical orientation channels. This weight matrix can then be used to infer the response profile across all 9 channels on each trial in the test data set. See equations 1-3 under *Estimating feature-selective BOLD response profiles using a forward encoding model* for more details.

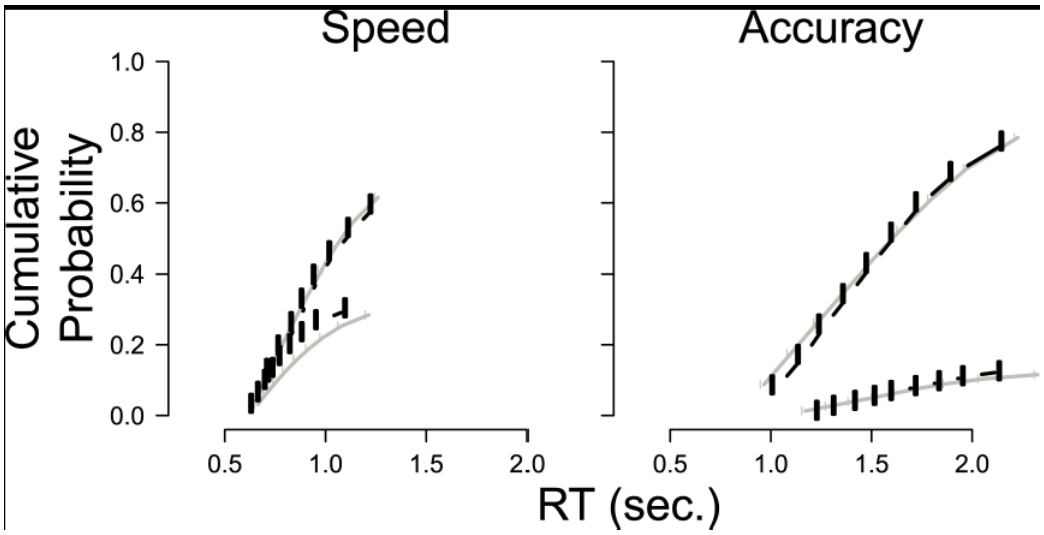

Figure 2.5. LBA model fits. Cumulative response time distributions estimated from the data using quantiles (black circles and lines) and predicted by the LBA model (gray dots and lines). Data are shown separately for the speed emphasis condition (left panel) and the accuracy emphasis condition (right 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%.

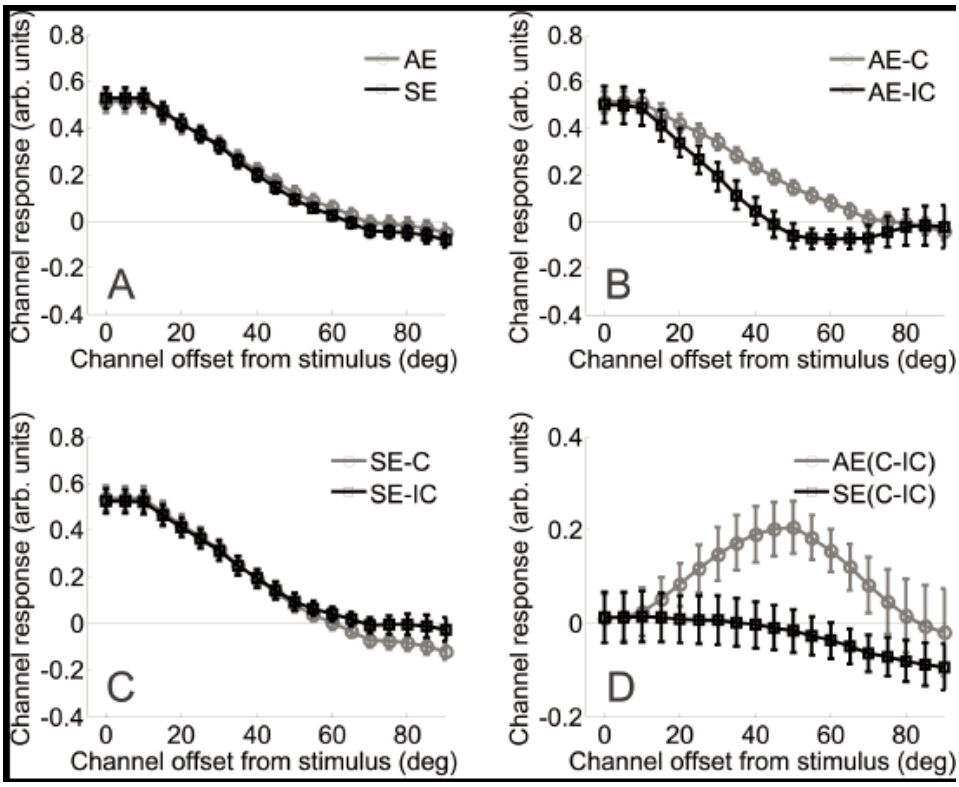

Figure 2.6 (**a**) Orientation-selective tuning functions for mismatch AE and SE trials. Channel tuning functions (TFs) for the AE data (gray) and SE data (black) are plotted as a function of each orientation-channel's offset from the stimulus presented on each trial (which is by convention always set 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 accuracy emphasized (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 ± 1S.E.M. (**c**) Tuning functions for mismatch speed emphasized (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 ± 1S.E.M. (**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 ± S.E.M. across subjects. All *p*-values were calculated based on a non-parametric randomization procedure (see *Materials and Methods* for more details).

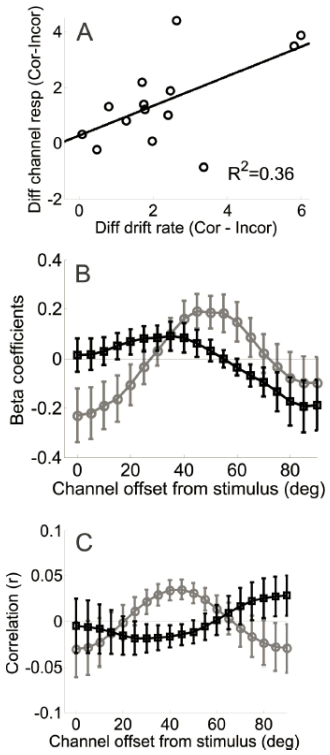

Figure 2.7 (**a**) Between-subject correlation between differences in off-target channel responses and differences in drift rates. The difference in the drift rates associated with correct and incorrect accumulators is positively correlated with the difference in off-target activation on correct and incorrect trials. (**b**) Logistic regression results. Beta 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 cross-over interaction. All error bars ± 1S.E.M. 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 where response threshold, non-decision time, and drift rate were allowed to vary between task conditions. All error bars ± 1S.E.M. across subjects. All *p*-values were calculated based on a non-parametric randomization procedure (see *Materials and Methods* for more details).

CHAPTER 3: Perceptual Consequences of Feature-based Attentional Enhancement and Suppression

Ho TC, Brown S, Abuyo NA, Ku EJ, Serences JT (in press). *Journal of Vision*.

Abstract

Feature-based attention has been shown to enhance the responses of neurons tuned to an attended feature while simultaneously suppressing responses of neurons tuned to unattended features. However, the influence of these suppressive neuronal-level modulations on perception is not well understood. Here, we investigated the perceptual consequences of feature-based attention by having subjects judge which of four random dot patterns (RDPs) contained a motion signal (Experiment 1) or which of four RDPs contained the most salient non-random motion signal (Experiment 2). Subjects viewed pre-cues which validly, invalidly, or neutrally cued the direction of the target RDP. Behavioral data were fit using the linear ballistic accumulator (LBA) model; the model design that best described the data revealed that the rate of sensory evidence accumulation (*drift rate*) was highest on valid trials and systematically decreased until the cued direction and the target direction were orthogonal. These results demonstrate behavioral correlates of both feature-based attentional enhancement and suppression.

Perceptual Consequences of Feature-based Attentional Enhancement and Suppression

Top-down visual attention is a highly adaptive mechanism that modulates sensory signals in order to facilitate the processing of behaviorally significant stimuli. Attention can be allocated based on prior knowledge of spatial locations (spatial attention) or of a target defining feature (feature-based attention). Numerous studies demonstrate that spatial attention improves behavioral performance on a wide array of tasks (Posner, 1980; reviewed in Carrasco, 2011) and that these improvements in performance are accompanied by corresponding increases in the gain of sensory neurons that have a receptive field at the attended location (e.g. McAdams and Maunsell, 1999; Reynolds et al., 1999, 2000; Williford and Maunsell, 2006; reviewed in Reynolds and Heeger, 2009). In contrast, feature-based attention enhances the gain of neurons that are tuned to an attended feature and suppresses the gain of neurons that are tuned away from the attended feature (Martinez-Trujillo and Treue, 2004; Treue and Martinez-Trujillo, 1999; Khayat et al., 2010; Cohen and Maunsell, 2011; Scolari, Byers, Serences, 2012).

Consistent with neurophysiology data, previous psychophysical studies suggest that feature-based attention selectively increases sensitivity to relevant visual features in a variety of perceptual tasks (Busse et al., 2008; Sàenz et al., 2003; Baldassi and Verghese, 2005; Felisberti and Zanker, 2005; Liu et al., 2007; Liu and Hou, 2011; Ling et al., 2009). However, the behavioral correlates of feature-based attentional suppression are less clear. One recent study suggests that feature-based attention suppresses neurons tuned away from an attended feature, consistent with evidence from single-unit recording (Ling et al., 2009), while another study reported only an enhancement of an attended feature value without concurrent suppression of dissimilar features (White and Carrasco, 2011).

However, neither study systematically varied the relationship between the target stimulus and the focus of feature-based attention, so the consequence of attentional suppression on the efficiency of processing unattended features was not directly evaluated.

In the present study, we employed a cueing paradigm to investigate both the facilitatory and suppressive effects of top-down feature-based attention on visual processing using a paradigm in which feature-based attention had to be deployed to stimuli that were distributed across the entire visual field. Using a quantitative model of perceptual decision making – the Linear Ballistic Accumulator (LBA; Brown and Heathcote, 2008) – we show that the rate of sensory evidence accumulation is highest for an attended feature and suppressed for stimulus directions rotated away from the attended direction. These findings provide behavioral evidence that feature-based attention can give rise to both improvements and impairments in perceptual processing.

Experiment 1

Methods

Subjects

11 right-handed subjects (5 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 in line with the guidelines of the local Institutional Review Board at UCSD and the Declaration of Helsinki and completed two 1-1.5 hr sessions in a climate and noise controlled subject room. Compensation for participation was $10.00/hr for the experiment. Data from 3 subjects

were discarded due to subjects' failure to return for the second session of the experiment (thus, data from 8 subjects were analyzed).

Stimuli and task

Subjects viewed the stimuli in a darkened room on a CRT monitor (MultiSync FP2141, refresh rate 85Hz) that was controlled by a PC running Windows XP. The luminance output of the monitor was measured with a Minolta LS-110 photometer and linearized in the stimulus presentation software. Subjects viewed the screen from a distance of approximately 60 cm. Visual stimuli were generated using the Psychophysics Toolbox stimulus presentation software (version 3; Brainard, 1997; Pelli, 1997) for Matlab (version 7.8.0; Mathworks, Natick, MA).

Subjects were presented with four random dot patterns (RDPs). One RDP was presented in each quadrant of the screen, and centered 7.82° from the horizontal midline and 6.98° from the vertical midline (see Figure 1). Each RDP was composed of small dots (0.15° x 0.15°) confined within a circular aperture 8° in diameter. Each dot was presented for a limited lifetime of 100 ms and moved at a speed of 5°/s. Dots that reached the edge of the aperture were moved to the opposite side of the aperture and redrawn. The target RDP contained 100% coherent motion (i.e., all dots moved in the same direction), while the distractors contained 0% coherent motion (the direction of each dot was drawn from a uniform distribution). Subjects had to indicate which RDP contained coherent motion using one of four keys on the number pad that corresponded to the spatial position of each of the four stimulus locations. At the beginning of each trial there was either a *valid* cue

(50% of the trials) that correctly indicated the impending direction of the target stimulus, an *invalid* cue (25% of the trials) that incorrectly indicated the impending direction of the target stimulus, or a *neutral* cue (25% of the trials) that gave no indication of the impending target direction (see Figure 1). Invalid cues indicated directions that were offset from the target direction by ±30° to ±180° in 30° steps, where the order of presentation was determined pseudo-randomly on each trial with the constraint that all offsets were equally represented. All cues were presented centrally for 1000 ms and were followed by a 1000 ms presentation of the four RDPs. Subjects were instructed to keep their eyes at fixation throughout the trial and were allowed to make a response any time after the onset of the stimulus array; each trial was self-paced and terminated once subjects provided a response. All subjects were encouraged to respond as quickly and as accurately as possible. Each block contained 48 trials in total, and the experiment consisted of 20 blocks. Subjects typically completed 4 to 8 blocks on the first day of the experiment (in addition to training, see below) and then returned a second day to complete the remaining blocks. Subjects were allowed to rest between blocks if they chose to do so, and all analyzed data came from subjects who completed all 20 blocks.

All participants were trained for a minimum of 160 trials immediately prior to the main experiment. All cues in the training session were neutral and provided no directional information. In order to equate performance and to ensure that subjects were not at ceiling, a staircasing procedure implemented in Psychtoolbox (QUEST, Watson and Pelli, 1983) was run to estimate the contrast threshold of the stimuli at which each individual could perform the task with approximately 75% accuracy. The contrast value estimated for each

subject was then fixed for the remaining blocks for the main experiment. The mean contrast level ± SEM across 8 subjects was 1.65% ± 0.05%.

Linear Ballistic Accumulator (LBA) model

The LBA model frames every decision as a race between *N* independent accumulators that correspond to each possible choice alternative, where $N=4$ in our experiments (see Figure 2 for a schematic of the model and Brown and Heathcote, 2008 for more details). The first accumulator to reach the *response threshold* (or *b*) determines the response choice and the response time. For every trial, each accumulator begins with a random activation level (the *starting point* or *k*) that is independently drawn from a uniform distribution on [0, *A*]. The starting points vary from trial to trial and from accumulator to accumulator, but the height of the distribution (*A*) was fixed for each of the four accumulators. Since "response caution" is defined as the distance between the response threshold and the starting point, we hereon use the response threshold parameter to represent "response caution," since the maximum of the starting point distribution here was fixed (in other situations where *A* is allowed to vary freely, "response caution" is sometimes defined as the response threshold minus the height of the starting point distribution; see Wolfe and Van Wert, 2010 for an example, although in that paper, "response caution" is referred to as "decision criterion"). During decision making, activity in each accumulator increases linearly and a response is deployed as soon as an accumulator crosses the response threshold. The predicted response time is thus the time taken to reach the threshold, in addition to a constant offset time (*non-decision time* or t_0). The stimulus display drives the rate at which sensory evidence is gathered for each

accumulator (*drift rate* or *d*). These drift rates vary from trial to trial according to independent normal distributions (with the standard deviation, *s*, of these distributions being arbitrarily fixed at 1), with means $v_1, v_2, \ldots, v_N$ for the *N* different response accumulators. The drift rate parameter estimated by the LBA model is thus the mean of this drift rate distribution, which reflects the quality of sensory information in favor of that particular response. For instance, if the upper right RDP contains 100% coherent motion while the other RDPs contain 0% coherent motion, there will be a large mean drift rate parameter for the accumulator corresponding to the upper right response, and small mean drift rates for the other three accumulators. All random values (i.e., the start points and drift rates) are drawn independently for each accumulator and are independent across decision trials.

Since the starting point for the evidence accumulator is a random sample from a uniform distribution on [0,*A*], the amount of evidence that needs to be accumulated to reach the threshold *b* is a sample from the uniform distribution *U*[*b*-*A*,*b*], assuming $b \geq A$. Since the drift rate for the *i*-th accumulator is a random draw from $N(v_i,s)$, the distribution function for the time taken for the *i*-th accumulator to reach threshold is the given by the ratio of these two, which has the following CDF (at time $t>0$):

$$F_i(t) = 1 + \frac{b - A - tv_i}{A}\Phi\left(\frac{b - A - tv_i}{ts}\right) - \frac{b - tv_i}{A}\Phi\left(\frac{b - tv_i}{ts}\right) + \frac{ts}{A}\phi\left(\frac{b - A - tv_i}{ts}\right) - \frac{ts}{A}\phi\left(\frac{b - tv_i}{ts}\right)$$

(1)

The associated PDF is:

$$f_i(t) = \frac{1}{A}\left[-v_i\Phi\left(\frac{b-A-tv_i}{ts}\right) + s\phi\left(\frac{b-A-tv_i}{ts}\right) + v_i\Phi\left(\frac{b-tv_i}{ts}\right) - s\phi\left(\frac{b-tv_i}{ts}\right)\right] \quad (2)$$

Note that lower case Greek letters above refer to probability density functions (PDFs), while the upper case Greek letters refer to cumulative distribution functions (CDFs). For more details regarding these equations and their derivations, see Brown and Heathcote, 2008.

We used the LBA for several reasons. First, the parameter estimates are jointly constrained by both accuracy and response time (RT) data (for both correct and incorrect responses), as opposed to using only one dependent measure. Second, the joint use of accuracy and RT data allows the LBA to naturally handle any speed-accuracy tradeoffs that may arise in the data. Finally, the LBA (and other similar models) isolate specific aspects of cognitive processing that are influenced by various experimental conditions (as captured by the different parameters in the model). For instance, in our cueing experiment, it could be that the cue influences the rate or quality of information extracted from the stimulus (as captured by the drift rate parameter), how much response caution a participant displays (as captured by the distance between the start point and the response threshold parameter), or it may affect non-decision related processes (as captured by the non-decision time parameter).

To evaluate the plausibility of each of these possible models, we fit 8 different versions of the LBA model to the data (using all possible combinations of these three parameters to capture the effects of cueing). The parameters were estimated using the method of maximum likelihood (see Donkin et al., 2011 for full details of these methods,

as well as an extensive discussion of alternative approaches). Initial parameter values for searches were generated two ways: 1) heuristic calculations based on the data and 2) start points determined from the end points of searches for simpler, nested models. Different mean drift rates were estimated for accumulators corresponding to the correct responses (that is, responses that matched the actual location of the target RDP) and the same single value was estimated for all three accumulators corresponding to incorrect responses (responses which did not match the true location of the target RDP). These correct and incorrect drift rates were then averaged across all 4 accumulators for each subject, respectively. For the purposes of our experiments, we report only the correct drift rates.

The most parsimonious of the 8 models that we evaluated was selected using the Bayesian Information Criterion (BIC), a commonly used criterion that evaluates the trade-off between model complexity and goodness of fit (Schwarz, 1978; Raftery, 1995):

$$BIC = -2 \cdot (\log lik) + (\log N) \cdot d.$$

Where *lik*=likelihood, *N*=number of data and *d*=number of parameters.

We calculated the group BIC across all subjects for each model design, by summing log-likelihoods, sample sizes and parameter counts. The design with the lowest BIC value was considered the most parsimonious model. We then approximated posterior model probabilities based on the BIC by assuming a fixed effect for subjects. This approach assumed that every subject had an identical structure, such that when we compared model designs (e.g., the model where only drift rate varied versus the model where only response threshold varied), it was assumed that all subjects were described by

the same model design. These posterior model probabilities provide support regarding the likelihood of each model design (for more details, see Burnham and Anderson, 2002).

Results

A one-way repeated measures ANOVA revealed a significant effect of cue validity on accuracy ($F(7,49) = 5.140$, $p=0.0002$, Figure 3a). Mean accuracy rates in the neutral condition were slightly higher than in the invalid condition, but not different from the valid condition (see Figure 3a and the first column of Table 1). Most notably, accuracy was lowest when the target was offset by 90° or 120° from the attention cue, after which there was a gradual improvement in accuracy (i.e., *rebound effect*, see Discussion). Consistent with this observation, a one-way repeated measures ANOVA on the accuracy rates for only the invalid trials revealed a significant difference between accuracy levels across the different cue-target offsets ($F(5,35)=2.8717$, $p=0.02$).

Similarly, RTs varied significantly as a function of cue condition ($F(7,49) = 2.505$, $p=0.0278$, Figure 3b). RTs were significantly faster on valid trials compared to neutral trials, but RTs on valid and invalid trials were of a similar magnitude (see Figure 3b and the second column of Table 1). A one-way repeated measures ANOVA on the RTs for only the invalid trials revealed no significant differences between RTs across the different cue-target offsets ($F(5,35)= 0.7621$, $p=0.5832$).

The accuracy and RT data present a mixed picture: the accuracy data suggest that attention confers no facilitation for validly cued features, but instead operates primarily by suppressing features dissimilar from the cue. In contrast, the RT data suggest facilitation for the cued feature and little suppression of features dissimilar from the cue. However,

there is a speed-accuracy tradeoff occurring: subjects were equally accurate on both valid and neutral trials, but responded significantly faster on valid trials (see the first and second columns of Table 1 and Figure 3). By using both accuracy and RT data, the LBA model accounts for such tradeoffs and is able to estimate how the cue manipulation selectively influences various latent variables such as the rate of sensory evidence accumulation (i.e., the drift rate), the amount of information required to make a decision (i.e., the response threshold), and the time related to non-decisional processes.

Out of the 8 tested model designs (where either response caution, drift rate, or non-decision time could stay fixed or vary), the best fitting LBA model – as determined by the BIC – only allowed drift rate to vary between the 8 possible cue-target conditions (i.e. validly cued targets, neutrally cued targets, and the 6 possible invalidly cued targets, collapsing across clockwise and counterclockwise offsets). The parameter estimates for response caution and non-decision time for this model were 3.188 and 0.169, respectively (see Figure 4 for the drift rates across the cue conditions). The posterior probability of this model was almost 1, which was more than 10^{40} times more likely than the next best design, in which both drift rate and non-decision time varied (see *Methods* for more details). Extremely strong support for one alternative like this is characteristic of group-average analyses like ours since such analyses compare extreme hypotheses (i.e., every subject is better described by one particular model compared to another). To confirm that these analyses did not bias our results, we also calculated posterior model probabilities separately for each individual participant and averaged the resulting probabilities: the model where only drift rate varied still had the greatest average posterior probability.

Figure 4 shows the drift rates corresponding to the accumulators that matched the correct response (and all subsequent stats focus on these drift rates as well, see the *Linear Ballistic Accumulator (LBA) model* section under *Methods* for more details). A one-way repeated-measures ANOVA revealed a strong effect of cue condition on drift rates ($F(7,49) = 6.793$, $p<0.0001$), and an individual *t*-test revealed that drift rates were larger on valid trials compared to neutral trials (see Figure 4 and see third column of Table 1). In contrast, the drift rates on with invalidly-cued trials were not significantly different from the drift rates observed on neutrally-cued trials (see Figure 4 and the third column of Table 1). A one-way repeated measures ANOVA on the drift rates associated with just invalidly-cued targets revealed a significant difference across the possible offsets ($F(5,35)=5.092$; $p=0.0013$). Thus, drift rates on invalid trials did exhibit a pattern that is analogous to the rebound effect observed in the accuracy data: as the offset between the cued direction and target direction increased, the drift rates corresponding to the correct response decreased up until an offset of 90°, after which they began to rise again.

In Experiment 2, we attempted to amplify the suppressive effects of feature-based attention by creating a variant of the behavioral paradigm in which the distractors competed more effectively with the search target.

Experiment 2

Methods

Subjects

13 right-handed subjects (5 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 two 1-1.5 hr sessions in a climate and noise controlled subject room. Compensation for participation was $10.00/hr for the experiment. Data from 1 subject were discarded due to failure to return for the second testing session, so data from 12 subjects were analyzed.

Stimuli and task

All methods used in Experiment 2 are similar to those used in Experiment 1 unless noted. In Experiment 2, we also used a four alternative forced choice task, but the single target RDP contained 80% coherent motion, while the distractors contained 40% coherent motion (see Figure 5). The purpose of this manipulation was to magnify competition between distractors and the target, which we anticipated would lead to larger attentional effects compared to Experiment 1. The direction of each distractor RDP was offset from the target direction by either $\pm$15°, $\pm$45°, $\pm$75°, $\pm$105°, $\pm$135°, $\pm$165° (pseudo-randomly chosen on each trial, with the constraint that each distractor direction was unique). Subjects indicated which RDP contained the high coherence RDP using keys on the number pad to indicate their decision. Invalid cues indicated directions that were offset from the direction of motion of the target RDP anywhere from $\pm$30° to $\pm$180° (in 30° steps) on a trial by trial basis. All cues were presented centrally for 1000 ms and the stimulus array was presented for 2000 ms. Subjects were instructed to keep their eyes at fixation throughout the trial and were allowed to make a response any time after the

stimulus onset; each trial was self-paced and terminated once subjects provided a response. All subjects were encouraged to respond as quickly and as accurately as possible. Each block contained 48 trials and in total, the experiment consisted of 20 blocks. Subjects typically completed 4 to 8 blocks on the first day of the experiment (in addition to training) and then returned a second day to complete the remaining blocks. Subjects were allowed to rest between blocks if they chose to do so, and all analyzed data came from subjects who completed all 20 blocks.

We used the same training and staircasing procedures described in Experiment 1 to titrate performance to approximately 75% before data collection in the main task. For the main experiment, the mean contrast level $\pm$ SEM across the 12 subjects was 15.21% + 4.3%.

Results

A one-way repeated measures ANOVA revealed a significant effect of cue-validity on accuracy ($F(7,77) = 19.561$, $p<0.0001$; Figure 6a). Accuracy was higher on valid trials compared to neutral trials, and generally lower on invalid trials compared to neutral trials (see Figure 6a and the first column of Table 2a). A one-way repeated measures ANOVA for accuracy rates on only the invalid trials also revealed a highly significant difference ($F(5,55)=11.319$, $p<0.0001$). Most notably, accuracy was lowest at 90° and 120° (see Figure 6a and the first column of Table 2a), after which there was a gradual improvement (a rebound effect, see Discussion).

RTs also varied as a function of cue condition ($F(7,77) = 7.44$, $p<0.0001$; Figure 6b). Compared with neutral trials, RTs were significantly faster on valid trials and significantly slower on invalid trials (see Figure 6b and the second column of Table 2a). A one-way repeated measures ANOVA for RTs on only the invalid trials also revealed a significant difference ($F(5,55)=3.2106$, $p=0.013$).

A one-way repeated measures ANOVA also revealed a significant modulation in the error rate as a function of the directional offset between the cue and an incorrectly selected distractor ($F(5,55) = 36.164$, $p<0.0001$; see Table 2b and Figure 6c). In particular, subjects selected a distractor whose direction was only 15° away from the invalid cue significantly more often than distractors rotated more than 105° from the invalidly cued direction (see Table 2b and Figure 6c).

As in Experiment 1, we found that the best fitting model based on the lowest BIC value was the one where only drift rate varied between the 8 possible cue types (collapsing across clockwise and counterclockwise offsets). The parameter estimates for response caution and non-decision time were 20.215 and 0.139, respectively (see Figure 7 for the drift rates across the cue conditions).

The approximate posterior model probability (see *Methods* for more details) was close to 1, more than 10^6 more likely than the next best design in which both drift rate and non-decision time varied. A one-way ANOVA revealed a robust effect of cue-type on drift rates ($F(7,77) = 8.281$, $p<0.0001$). Consistent with the raw accuracy rates and RTs, drift rates estimated from valid trials were significantly larger than drift rates estimated from neutral trials. Drift rates were also significantly lower for all invalid offsets compared to

neutral trials, with the exception of the 180° offset. A one-way repeated measures ANOVA for drift rates on only the invalid trials revealed a highly significant difference ($F(5,55)=6.7937$, $p<0.0001$). Drift rates showed the same rebound effect evident in the accuracy data: drift rates continued to decrease as the offset between the cue and target increased, up until an offset of 90°, after which they returned to the level observed on neutral trials (see Figure 7 and the third column of Table 2a).

Discussion

We utilized two variants of a cueing paradigm to show that goal-directed feature-based attention influences visual performance through both facilitatory and suppressive mechanisms. In Experiment 1, when subjects searched for an oddball motion target, the LBA model revealed increased rates of sensory evidence accumulation for validly cued targets compared to neutral targets (facilitation). In addition, drift rates on invalidly cued trials varied as a function of the cue target offset, with highest drift rates observed for invalid targets that most closely matched the cue. However, while these patterns are suggestive, the drift rates on invalid trials were not significantly lower than drift rates on neutral trials. In Experiment 2, two key changes were made in an attempt to amplify attentional enhancement and suppression. First, we increased competition between the target and distractors in an effort to place additional demands on suppressive attentional mechanisms. Second, as a result of the increased difficulty associated with having more salient distractors, the RDPs in Experiment 2 had higher mean contrast levels compared to Experiment 1: 15.21% versus 1.65%. Given that both neurophysiological and psychophysical work suggest that higher contrast levels can yield greater feature-based

attentional effects (e.g. Martinez-Trujillo and Treue, 2002; Hermann, Heeger, and Carrasco, 2012), it is possible that this increase in contrast might also contribute to more pronounced attentional modulations. Indeed, these manipulations had the intended effect, as an analysis of drift rates revealed extremely robust evidence for both attention-related facilitation and suppression (Figure 7).

Rebound effect

In Experiment 1, drift rates corresponding to invalidly cued targets were suggestive of a suppressive effect of feature-based attention (see Figure 4 and the third column of Table 1). This pattern, however, was far more pronounced in Experiment 2, where drift rates corresponding to the invalidly cued targets were significantly lower than those for neutrally cued targets as the offset between the invalid cue and target approached 90° (see Figure 7 and the third column of Table 2a). However, as the offset between the invalid cue and target increased from 90° to 180°, the corresponding drift rates also gradually increased, giving rise to a "rebound effect" (see Figure 7 and the third column of Table 2a). A similar rebound pattern was also observed in Experiment 1, where a ANOVA conducted on drift rates on invalidly cued targets was significant (see Figure 4 and the third column of Table 1).

This non-monotonic change in drift rate as a function of cue-to-target separation is superficially inconsistent with neurophysiological data showing maximal neural suppression for directions opposite from the attended direction (Martinez-Trujillo and Treue, 2004). Recall that our task only required subjects to identify the quadrant that contained the RDP with the highest motion coherence (among distractors containing 0%

motion coherence in Experiment 1 and 40% motion coherence in Experiment 2). As a result, the observed rebound effect could be related to subjects monitoring orientation signals associated with the axis of motion rather than the direction of motion per se. Indeed, many neurons in areas such as primary visual cortex will respond robustly to the axis of motion in a manner analogous to a static bar rendered at the same orientation (see Albright, 1984; Livingstone, 1998; Conway and Livingstone, 2003; Livingstone and Conway, 2003). Note, however, that behavioral performance in Experiment 2 on invalid trials with a 180° offset was not as good as performance on valid trials (see Figures 6-7 and Table 2a). Therefore, signals from motion selective neurons tuned 180° from the cue might very well have been suppressed in motion selective visual areas such as MT, and the partial recovery of performance might be supported by attentional gain targeted towards orientation selective neurons in other areas that are tuned to the axis of motion. In either case, the data suggest that subjects were using the cue to guide search, and clear evidence supporting the suppressive effects of feature-based attention were observed.

Comparisons with the feature-similarity gain model

Compared to a condition in which only a fixation mark was attended, Martinez-Trujillo and Treue (2004) reported enhanced responses in MT neurons that were tuned to an attended direction of motion and suppressed responses in neurons that were tuned away from the attended direction of motion (see also: Cohen and Maunsell, 2011). These data predict enhanced perceptual sensitivity for attended features as well as a reduction in perceptual sensitivity for features that are maximally separated from the target in feature space (e.g. opposite directions of motion). Consistent with a suppressive component of

feature-based attention, Ling et al. (2009) used an equivalent-noise paradigm and a task in which subjects had to discern whether or not a centrally presented RDP contained motion that was clockwise or counterclockwise from one of four reference directions (see also Baldassi and Verghese, 2005). The coherence (proportion of dots moving in the same direction) of the RDP was systematically adjusted across multiple levels, and the directional offset between the target direction and the reference directions was adjusted to estimate a perceptual sensitivity threshold for each subject. On validly cued trials, attention reduced sensitivity thresholds across all motion coherence levels, suggesting that feature-based attention suppressed neural responses evoked by dots within the RPDs that were moving in uncued directions. This putative suppression of uncued motion directions is consistent with previous neurophysiology studies and with our psychophysical and modeling results. However, the present results further demonstrate that the suppressive effect of feature-based attention can operate across multiple locations in the visual field, as opposed to being restricted to a single spatial location that contains both the cued and the uncued features (as in the variable coherence RDPs used by Ling et al., 2009). In addition, our data reveal a systematic decline in processing efficiency as a function of the directional offset between the cue and the target. This systematic relationship between the cue/target offset and the efficiency of sensory processing is evident both in the pattern of drift rates on invalid trials (see Figure 7), and in the pattern of errors made on invalid trials, as subjects in Experiment 2 more often chose a distractor whose direction was similar to the cue compared to a distractor whose direction was far from the cue (see Figure 6c).

White and Carrasco (2011) also directly assessed the relationship between feature-based attentional suppression and behavior using a dual-task paradigm. Subjects were asked to indicate whether there was a speed change in a primary RDP (that was shown in one hemifield) and then asked to discern which of two secondary RDPs shown in the other hemifield contained coherent motion. The direction of coherent motion in the secondary RDP either matched or mismatched (by 180°) the direction of motion of the primary stimulus. As in the present study, performance was enhanced when subjects were given a valid cue that indicated the direction of the primary RDP compared to when they were given a neutral cue. In addition, subjects were more sensitive to coherent motion in the secondary stimulus when the motion in the secondary stimulus matched the cued direction. However, when the direction of the secondary stimulus mismatched the cued direction, performance was not impaired compared to a neutral cue condition. Thus, their overall pattern of data is consistent with attentional facilitation in the absence of suppression. In contrast, the results from our experiments – particularly Experiment 2 – provide strong evidence for both attentional facilitation and suppression. However, our results are not necessarily inconsistent with White and Carrasco's findings, as motion in their secondary stimulus either matched the cued direction or was offset by 180° from the cued direction. In both of our experiments, suppression was minimized for directions that were 180° from an invalid cue (the "rebound effect," see Figures 3-4 and 6-7). Thus, suppression might simply have been minimized at the 180° offset in their study.

Finally, Tombu and Tsotsos (2008) evaluated a feature-based version of their *selective-tuning model*, which posits that attention facilitates the processing of attended

features, suppresses the processing of immediately adjacent features, and has little impact on the processing of distant features (see also Tsotsos, 1995). In their study, subjects had to decide whether the stripes on a grating stimulus were jagged or straight, and the attended orientation was cued on a block-by-block basis. The orientation of the stripes either matched the attended orientation, or was offset by 45° (the "similar" condition) or by 90° (the "dissimilar" condition). Accuracy rates were highest when the stripes matched the cued orientation, lowest in the similar condition (45° offset) and intermediate in the dissimilar condition (90° offset). Although this pattern was only observed with "jagged" stripes (and not with straight stripes), their data are consistent with the pattern of suppression and rebound that we report in Experiment 1 (see Figures 3-4) and Experiment 2 (see Figures 6-7). Namely, the pattern of suppression was non-monotonic as the offset between the cued orientation and the presented orientation increased, and offsets farther away from the cued feature were not suppressed as strongly as offsets that were at an intermediate distance from the cued feature value. Thus, it is possible that our results – particularly in Experiment 2 – tap into a similar mechanism proposed in Tsotsos' selective-tuning model. The present study, however, complements and extends their results by demonstrating enhancement and suppression relative to a neutral cue baseline, by sampling more feature values, and by implementing a quantitative decision model that can better isolate the different latent cognitive factors that are involved in perceptual decision making.

In sum, our results reveal both facilitatory and suppressive effects of feature-based attention that systematically depend on the directional offset of the stimulus and the currently attended direction. These graded attention effects – that primarily influence the

rate of sensory evidence accumulation – are generally consistent with the feature-similarity gain model, which predicts increasingly impaired performance for features that are tuned progressively farther away from the cued direction (but see our discussion of the "rebound effect" and the selective tuning model above). To extend these findings in future studies, we can couple our general experimental and analytical approach with paradigms developed by other investigators to examine interactions between suppression and the spatial extent of feature-based attention (e.g. Liu et al., 2011), as well as the impact of feature-based attention on the simultaneous processing of multiple relevant stimuli (e.g. White and Carrasco, 2011; Sàenz et al., 2003).

Mechanisms of attentional enhancement and suppression: sensory gain versus selective weighting

Given the large and growing amount of neurophysiological evidence, it is tempting to ascribe the facilitatory and suppressive behavioral effects observed in our study to modulations of neurons at relatively early stages of sensory processing (Treue and Maunsell, 1996; 1999; Martinez-Trujillo and Treue, 2004; Maunsell and Treue, 2006; Sàenz et al., 2002; Serences et al., 2009; Scolari et al., 2012; Liu et al., 2007). However, the attentional enhancement and suppression effects that we observe may not reflect changes in early sensory gain, but instead a selective weighting of sensory responses by downstream decision mechanisms (Law and Gold, 2008, 2009; Dosher and Lu, 1999; Eckstein et al., 2000; Palmer, 1995; Palmer et al., 2002; Baldassi and Verghese, 2002; Baldassi and Verghese, 2005). On this account, areas involved in integrating sensory evidence during decision making might overweight the output of direction-selective neural

populations tuned to the cued feature, and underweight input from neural populations tuned away from the cued direction. This selective weighting would give rise to enhanced performance in the valid-cue condition, and systematically worse performance for targets rotated farther from the cue (as in Figures 3-4 from Experiment 1, and Figures 6-7 from Experiment 2). Moreover, this type of selective weighting also predicts the systematic pattern of errors observed in Experiment 2, as neurons that respond to distractor directions that are close to the cued direction might more strongly influence decision mechanisms and trigger more frequent incorrect responses (Figure 6c). Finally, the selective weighting account might explain the larger suppressive effects that we observed in Experiment 2 compared to Experiment 1, as the 40% coherent RDP distractors in Experiment 2 are more likely to influence decision mechanisms that differentially weight responses from neural populations that respond to directions that are adjacent to the cued direction.

Thus, based on the current behavioral data alone, we cannot unambiguously determine if the behavioral facilitation on valid trials and the impaired performance on invalid trials is due to sensory enhancement, selective weighting during decision making, or – more likely – to some combination of the two mechanisms. One potential method to disentangle these alternatives would be to re-run a version of Experiment 2 with fewer offsets and more sessions in order to dramatically increase the amount of data that is collected. With more power, it may be possible to use a modified LBA model to estimate response caution for every possible cue-distractor offset on trials in which a distractor is incorrectly selected in place of the target. Thus, one could test whether or not response caution is lower when distractors are similar to the invalid cue, which might support the

idea that the decision criteria are adjusted based on cue-stimulus similarity. However, an alternate, and perhaps more fruitful, approach would be to adapt a version of Experiment 2 to either single unit recording or functional magnetic resonance imaging environments. This would enable measurements of attention-mediated changes in sensory gain throughout visual cortex, and allow one to examine if most of the variability in behavior can be accounted for by only considering the observed changes in sensory gain or if positing an additional mechanism that selectively reweights sensory responses during decision making is needed to fully explain the nature and magnitude of the observed changes in behavior (see Pestilli et al., 2011).

Linear ballistic accumulator versus signal detection models

The linear ballistic accumulator (LBA) model is one of many models that utilizes RT and accuracy to examine how experimental manipulations impact latent cognitive factors such as drift rate, decision boundaries, and non-decision times (e.g. Ratcliff, 1978; Brown and Heathcote, 2008; Link and Heath, 1975; Usher and McClelland, 2001; van Zandt et al. 2000; Wagenmakers et al., 2007). Other models, such as those based on SDT, also do an excellent job explaining accuracy data that is obtained in experiments similar to ours (see Palmer, 1995; Eckstein et al., 2000; Palmer et al, 2000; Baldassi and Verghese, 2002; Baldassi and Verghese, 2005). In these SDT models, each element in a visual display elicits a noisy representation within each relevant feature dimension. The observer then combines these representations across all features to obtain a single decision variable, and the stimulus with the maximum value is deemed the most likely to be the target. This approach is particularly adept at explaining behavior in situations where the target is

exposed briefly in a data-limited manner, and accuracy is the primary dependent measures (as opposed to RT). However, even though SDT models account for performance on an impressive array of search tasks, sequential sampling models often have an advantage. This is particularly true when a decision is based on noisy sensory inputs that are continuously present so that the reliability of the final decision variable continuously increases as more evidence is unveiled (Smith and Ratcliff, 2004). In addition, the conception of a perceptual decision as an accumulation process is now well supported based on studies that examine the buildup of firing rates of sensorimotor neurons that are thought to play a key role in mediating basic perceptual decisions (Mazurek et al., 2003; Gold and Shadlen, 2007; Leon and Shadlen, 1999; Heekeren et al., 2004).

Feature-based attention and changes in neural variability

As discussed above, changes in the rate of sensory evidence accumulation (depicted in Figures 4 and 7) can be intuitively linked to changes either in the firing rate of sensory neurons or to changes in the relative weighting of sensory signals by downstream decision mechanisms. However, there is growing evidence that top-down attention can also influence neural variability (Mitchell et al., 2007; 2009; Cohen and Maunsell, 2009; 2010; 2011; Kohn and Cohen, 2011). For example, space and feature-based attention have been shown to reduce the ratio of the variance of individual sensory neurons to their mean firing rate (or the *fano factor*; Mitchell et al., 2009; Cohen and Maunsell, 2011). In turn, a reduction in neural variability should lead to more stable sensory responses across time, thus speeding the process of evidence accumulation during decision making. Recently, Rahnev and colleagues (2011) used a signal detection theory (SDT) model to demonstrate

that such a reduction in neural variability can increase behavioral performance by increasing the signal available for supporting perceptual decisions. Interestingly, this reduction in variability improves performance on average, but also reduces the probability of a high signal response on any given trial, thereby leading to a more conservative bias in behavioral responses (Rahnev et al., 2011). While these results are not directly comparable to the present findings due to differences in the task and the type of model that was employed, they suggest an important role for attention-mediated reductions in neural noise during perceptual decision making. An advantage of sequential sampling models, like the LBA, is that such explanations can be directly investigated. For example, one could explore LBA models in which the parameter governing the variance of the drift rate distributions was free to vary with cueing manipulations. Such models can be difficult to estimate due to the highly correlated posterior distributions for mean drift rate and variance parameters, so such investigations will require future studies that have many more data points per participant (see Smith et al., 2004; Donkin et al., 2009).

Conclusions

Here, we provide evidence for both facilitatory and suppressive effects of top-down feature-based attention on human performance that can best be explained by changes in the rate of sensory evidence accumulation, particularly when competition between the target and distractors was high. One critical outstanding issue concerns developing a mechanistic explanation for the observed effects and determining whether feature-based attention mediates behavior primarily via changes in sensory gain, changes in the weighting of sensory evidence by downstream decision mechanisms, or through

some combination of the two. Future studies employing neurophysiological or neuroimaging approaches will hopefully dissociate these possibilities and isolate the relative contributions of each mechanism.

Chapter 3, in full, is a reprint of the material as it appears in Perceptual Consequences of Feature-based Attentional Enhancement and Suppression in *Journal of Vision*. Ho, T.C., Brown, S. Abuyo, N.A., Ku, E.J., Serences, J.T. (in press).

Table 3.1. Pair-wise comparisons of mean accuracy rates, mean response times, and mean drift rates between different cued conditions for Experiment 1 (*n*=8). Offsets were collapsed across clockwise and counterclockwise directions. Please refer to Figure 3.3 and Figure 3.4. Note: the reported *p*-values here have been corrected for multiple comparisons using a False Discovery Rate procedure. **p*<0.05, ***p*<0.01, ****p*<0.001.

Comparison	**Accuracy rate**	**Reaction time**	**Drift rate**
Neutral v. Valid	*p*=0.9948	*p*=0.0364 *	*p*=0.0168 *
Neutral v. 30°	*p*=0.995	*p*=0.085	*p*=0.9392
Neutral v. 60°	*p*=0.1367	*p*=0.1962	*p*=0.8191
Neutral v. 90°	*p*=0.1003	*p*=0.7059	*p*=0.7433
Neutral v. 120°	*p*=0.1	*p*=0.2545	*p*=0.744
Neutral v. 150°	*p*=0.1699	*p*=0.254	*p*=0.8631
Neutral v. 180°	*p*=0.1	*p*=0.1895	*p*=0.1157
Valid v. 30°	*p*=0.17	*p*=0.25	*p*=0.0566
Valid v. 60°	*p*=0.1005	*p*=0.0847	*p*=0.0642
Valid v. 90°	*p*=0.1168	*p*=0.254	*p*=0.0616
Valid v. 120°	*p*=0.1005	*p*=0.19	*p*=0.062
Valid v. 150°	*p*=0.17	*p*=0.1895	*p*=0.057
Valid v. 180°	*p*=0.101	*p*=0.0770	*p*=0.1814
30° v. 60°	*p*=0.1862	*p*=0.019	*p*=0.5965
30° v. 90°	*p*=0.14	*p*=0.3811	*p*=0.3192
30° v. 120°	*p*=0.1367	*p*=0.1962	*p*=0.345
30° v. 150°	*p*=0.2625	*p*=0.0196	*p*=0.056
30° v. 180°	*p*=0.1699	*p*=0.1305	*p*=0.062

Table 3.1. Pair-wise comparisons of mean accuracy rates, mean response times, and mean drift rates between different cued conditions for Experiment 1 (*n*=8), Continued. Offsets were collapsed across clockwise and counterclockwise directions. Please refer to Figure 3.3 and Figure 3.4. Note: the reported *p*-values here have been corrected for multiple comparisons using a False Discovery Rate procedure. *$p<0.05$, **$p<0.01$, ***$p<0.001$.

Comparison	**Accuracy**	**Reaction time**	**Drift rate**
60° v. 90°	$p=0.2123$	$p=0.784$	$p=0.8635$
60° v. 120°	$p=0.1368$	$p=0.9225$	$p=0.8651$
60° v. 150°	$p=0.5676$	$p=0.7840$	$p=0.345$
60° v. 180°	$p=0.2625$	$p=0.92$	$p=0.0617$
90° v. 120°	$p=0.3746$	$p=0.7908$	$p=0.9392$
90° v. 150°	$p=0.995$	$p=0.791$	$p=0.25$
90° v. 180°	$p=0.4597$	$p=0.784$	$p=0.062$
120° v. 150°	$p=0.2787$	$p=0.79$	$p=0.249$
120° v. 180°	$p=0.3084$	$p=0.9667$	$p=0.059$
150° v. 180°	$p=0.9949$	$p=0.785$	$p=0.056$

Table 3.2a. Pair-wise comparisons of mean accuracy rates, mean response times, and drift rates between different cued conditions for Experiment 2 (*n*=12). Offsets were collapsed across clockwise and counterclockwise directions. Please refer to Figures 3.6a,b and Figure 3.7. Note: the reported *p*-values here have been corrected for multiple comparisons using a False Discovery Rate procedure. *$p<0.05$, **$p<0.01$, ***$p<0.001$.

Comparison	**Accuracy rate**	**Response time**	**Drift rate**
Neutral v. Valid	p=0.000364 ***	p=0.0021 **	p=0.0056 **
Neutral v. 30°	p=0.3960	p=0.5851	p=0.0409 *
Neutral v. 60°	p=0.0022 **	p=0.052	p=0.0491 *
Neutral v. 90°	p=0.0004 ***	p=0.0491 *	p=0.0409 *
Neutral v. 120°	p=0.00035 ***	p=0.0037 **	p=0.0436 *
Neutral v. 150°	p=0.0019 **	p=0.0031 **	p=0.044 *
Neutral v. 180°	p=0.1656	p=0.01 *	p=0.4928
Valid v. 30°	p=0.00043 ***	p=0.052	p=0.0261 *
Valid v. 60°	p=0.00035 ***	p=0.0018 **	p=0.0325 *
Valid v. 90°	p=0.0003 ***	p=0.01 *	p=0.026 *
Valid v. 120°	p=0.0003 ***	p=0.0019 **	p=0.027 *
Valid v. 150°	p=0.0004 ***	p=0.0037 **	p=0.026 *
Valid v. 180°	p=0.0065 **	p=0.01 *	p=0.0437 *
30° v. 60°	p=0.00037 ***	p=0.046 *	p=0.3794
30° v. 90°	p=0.00036 ***	p=0.052	p=0.0659
30° v. 120°	p=0.00036 ***	p=0.0112 *	p=0.1376
30° v. 150°	p=0.0019 **	p=0.0328 *	p=0.5381
30° v. 180°	p=0.1387	p=0.0835	p=0.0491 *

Table 3.2a. Pair-wise comparisons of mean accuracy rates, mean response times, and drift rates between different cued conditions for Experiment 2 (*n*=12), Continued. Offsets were collapsed across clockwise and counterclockwise directions. Please refer to Figures 3.6a,b and Figure 3.7. Note: the reported *p*-values here have been corrected for multiple comparisons using a False Discovery Rate procedure. *$p<0.05$, **$p<0.01$, ***$p<0.001$.

Comparison	**Accuracy**	**Reaction time**	**Drift rate**
60° v. 90°	p=0.0239 *	p=0.3645	p=0.02
60° v. 120°	p=0.0575	p=0.1426	p=0.041
60° v. 150°	p=0.7924	p=0.3948	p=0.7489
60° v. 180°	p=0.3326	p=0.6406	p=0.044 *
90° v. 120°	p=0.6875	p=0.5851	p=0.4029
90° v. 150°	p=0.057	p=0.7351	p=0.0473 *
90° v. 180°	p=0.0275 *	p=0.6283	p=0.03 *
120° v. 150°	p=0.0565	p=0.3171	p=0.0436 *
120° v. 180°	p=0.0145 *	p=0.1737	p=0.026 *
150° v. 180°	p=0.0413*	p=0.7499	p=0.027 *

Table 3.2b. Pair-wise comparisons of the mean error rate on invalid trials according to the different offsets between the invalid cue and chosen distractor direction (*n*=12). Offsets were collapsed across clockwise and counterclockwise directions. Please refer for Figure 3.6c. Note: the reported *p*-values here have been corrected for multiple comparisons using a False Discovery Rate procedure. *$p<0.05$, **$p<0.01$, ***$p<0.001$.

Comparison	
15° v. 45°	$p=0.4164$
15° v. 75°	$p=0.2786$
15° v. 105°	$p=0.0015$ **
15° v. 135°	$p=0.006$ **
15° v. 165°	$p=0.0321$ *
45° v. 75°	$p=0.8433$
45° v. 105°	$p=0.1436$
45° v. 135°	$p=0.0071$ **
45° v. 165°	$p=0.084$
75° v. 105°	$p=0.1579$
75° v. 135°	$p=0.006$ **
75° v. 165°	$p=0.2217$
105° v. 135°	$p=0.2456$
105° v. 135°	$p=0.8051$
135° v. 165°	$p=0.169$

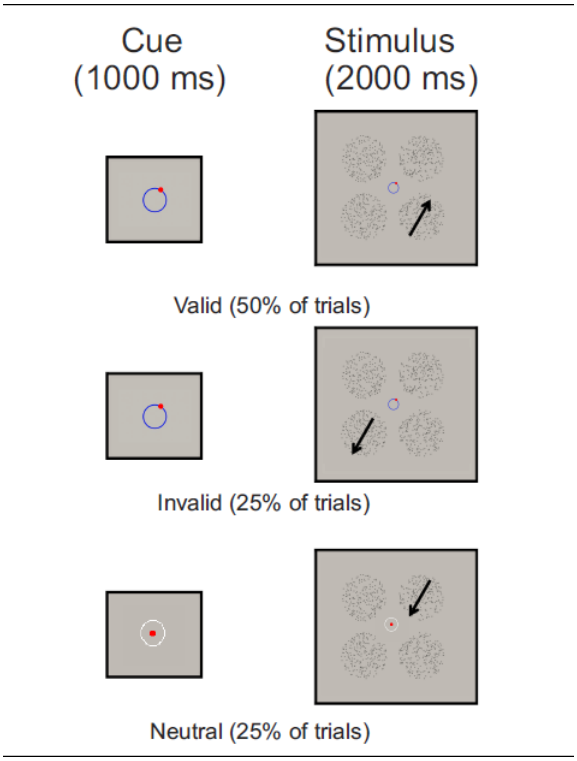

Figure 3.1. Behavioral paradigm for Experiment 1. At the beginning of each trial there was either a *valid* cue (50% of the trials) correctly indicating the impending target direction of motion, an *invalid* cue (25% of the trials) incorrectly indicating the impending target direction of motion, or a *neutral* cue (25% of the trials) which contained no directional information whatsoever. Invalid cues indicated directions that were offset from the direction of motion of the target RDP anywhere from $\pm$30° to $\pm$180° on a trial by trial basis (determined pseudo-randomly within a given block of trials). All cues were presented centrally for 1000 ms and were followed by 1000 ms of the four RDPs. Note: cues in the first panel are exaggerated in size for the purposes of clarity).

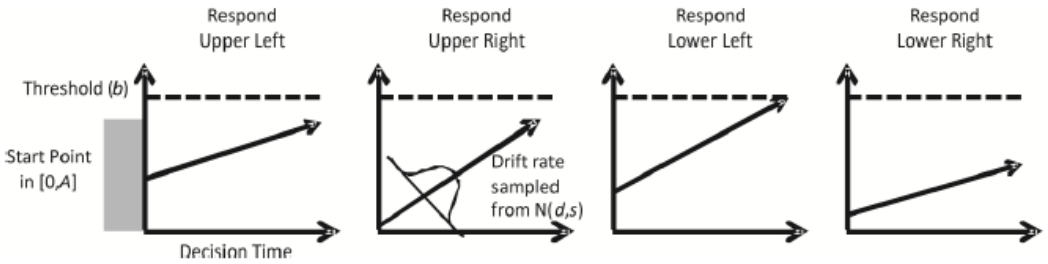

Figure 3.2. Schematic of the linear ballistic accumulator (LBA) model. The stimulus provides information to four racing accumulators (each corresponding to one of four spatially distinct RDPs); the first accumulator to reach the response threshold determines the response, and thus the decision time. One accumulator corresponds to each possible response and their average rates of increase (drift rates) are assumed to be determined by the stimulus properties. Response caution determines how much sensory evidence needs to be accumulated before a response is made and is captured by the relative distance between the response threshold and the start point. The final response time is the decision processing time (i.e., the time taken for the first accumulator to reach the response threshold) plus a constant offset (non-decision time).

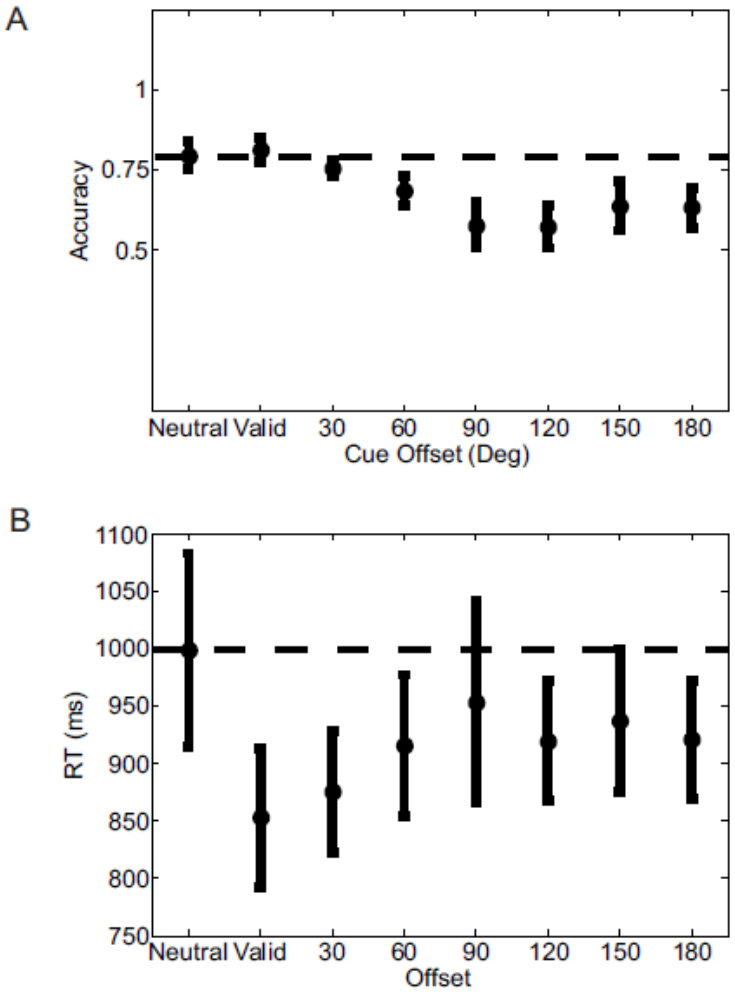

Figure 3.3 (a) Mean acccuracy rates across the different cue conditions for Experiment 1. Mean accuracy rates computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). **(b) Mean Response times on correct trials across the different cued conditions for Experiment 1.** Mean response times computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). Offsets were collapsed across clockwise and counterclockwise directions. All error bars are $\pm$1 SEM, computed after subtracting the mean from each subject.

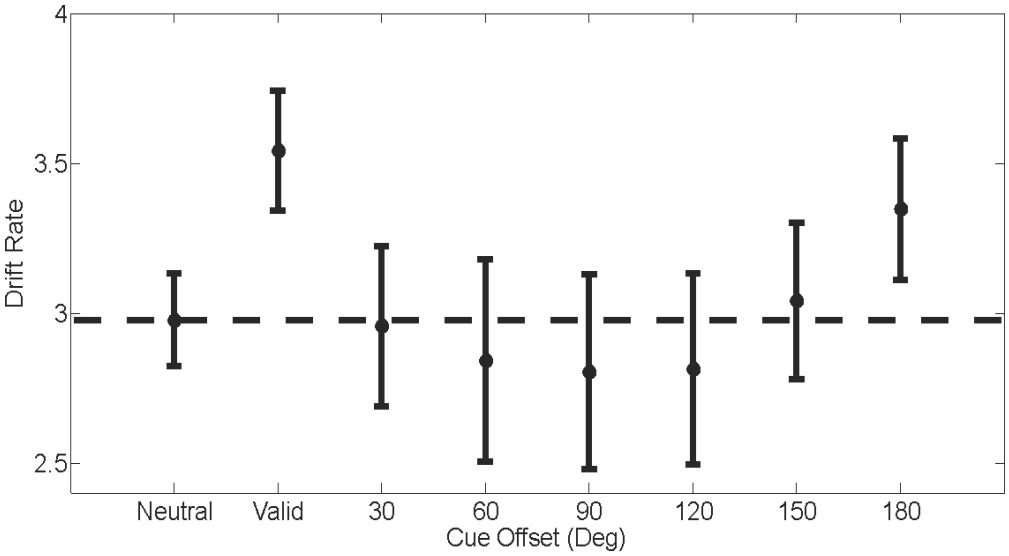

Figure 3.4. Drift rates for the accumulators corresponding to the correct response on the different cue conditions for Experiment 1. Mean drift rates corresponding to the correct response on each trial computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). Offsets were collapsed across clockwise and counterclockwise directions.All error bars are $\pm$1 SEM, computed after subtracting the mean from each subject.

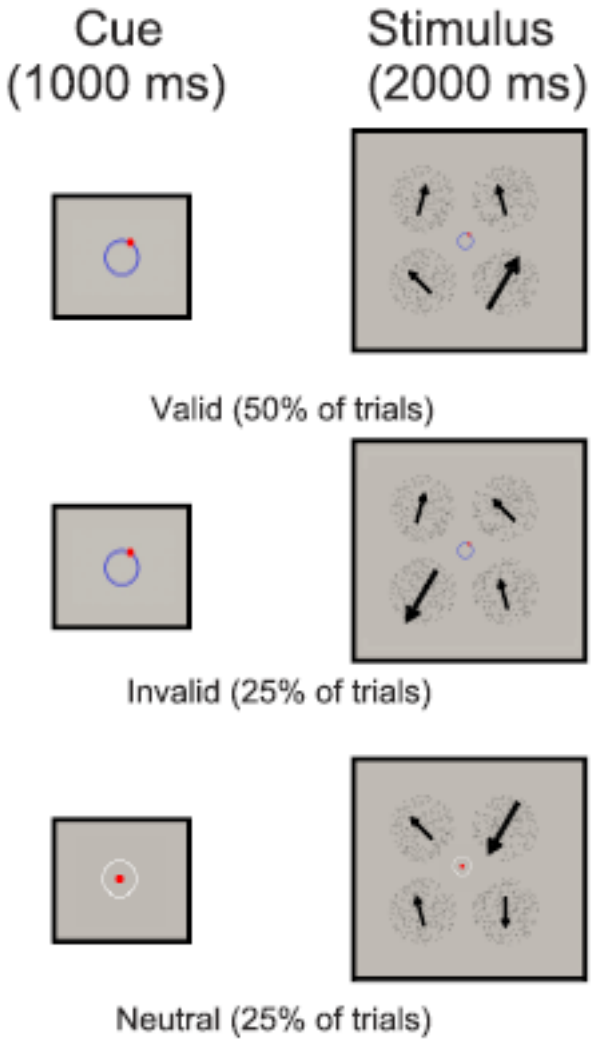

Figure 3.5. Behavioral paradigm for Experiment 2. At the beginning of each trial there was either a *valid* cue (50% of the trials) correctly indicating the impending direction of motion, an *invalid* cue (25% of the trials) incorrectly indicating the impending direction of motion, or a *neutral* cue (25% of the trials) which contained no directional information whatsoever. Invalid cues indicated directions that were offset from the direction of motion of the target RDP anywhere from ±30° to ±180° on a trial by trial basis (determined pseudo-randomly within a given block of trials). The target RDP contained 80% motion coherence, while the distractors contained 40% motion coherence, where each distractor direction was offset from the target direction by either ±15°, ±45°, ±75°, ±105°, ±135°, ±165° (randomly chosen per trial, where each distractor direction was unique) for 8 of the subjects or by ±18°, ±36°, or ±72° for the remaining 4 subjects. Note: cues in the first panel are exaggerated in size for the purposes of clarity.

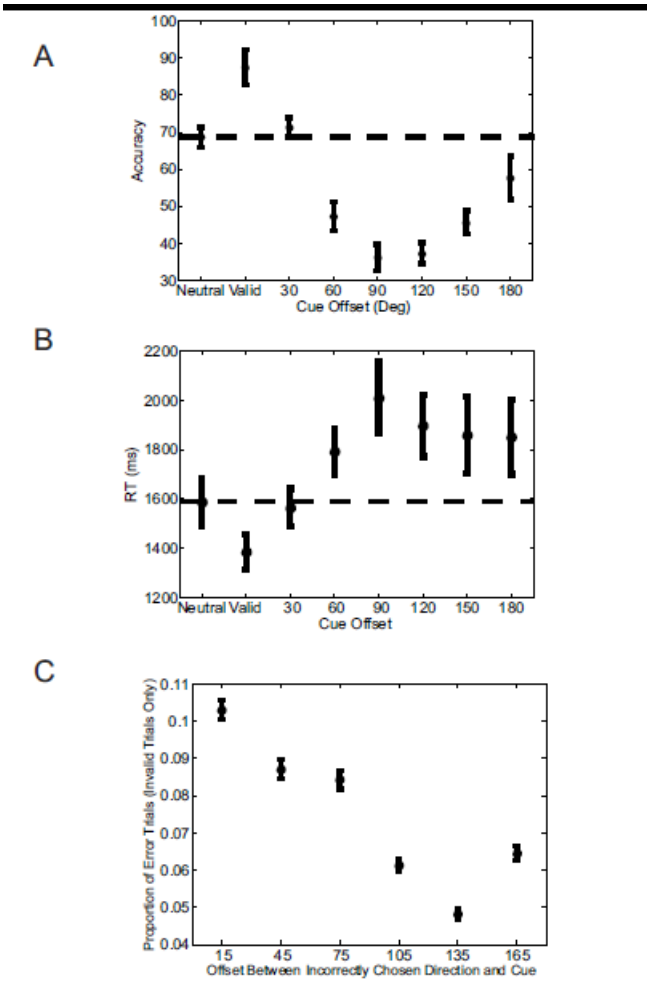

Figure 3.6 (a) Accuracy rates across the different cue conditions for Experiment 2. Mean accuracy rates computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). **(b) Mean response times on correct trials across the different cued conditions for Experiment 2.** Mean accuracy rates computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). **(c) Mean proportion of error trials for every offset between the chosen distractor and invalid cue for Experiment 2.** Mean proportion of errors as a function of the offset between the direction of the chosen distractor (ordinate) and the invalid cue for that trial (abscissa). Only incorrect trials with invalid cues were used in this analysis. Offsets were collapsed across clockwise and counterclockwise directions. All error bars are $\pm$ SEM, computed after subtracting the mean from each subject.

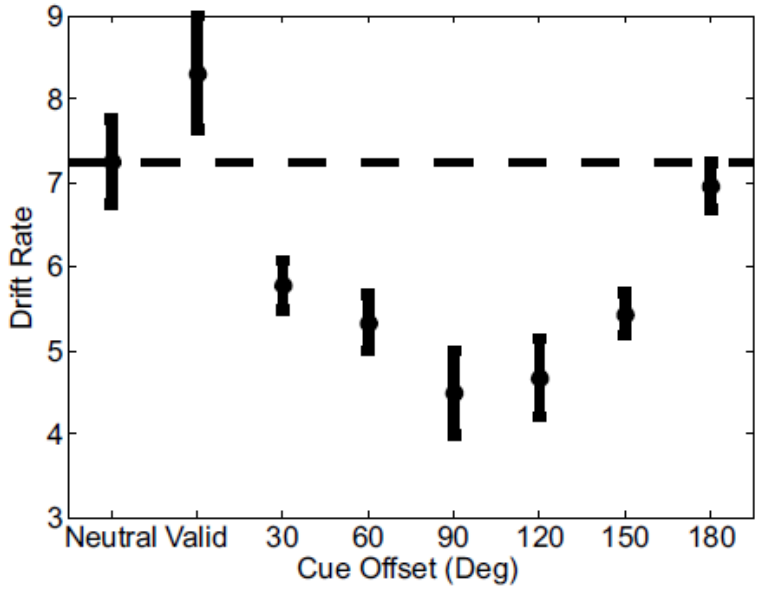

Figure 3.7. Drift rates for the accumulators corresponding to the correct response on the different cue conditions for Experiment 2. Mean drift rates corresponding to the correct response on each trial computed across subjects (ordinate) as a function of the 8 possible cue conditions (abscissa). Offsets were collapsed across clockwise and counterclockwise directions. All error bars are $\pm$1 SEM, computed after subtracting the mean from each subject.

GENERAL DISCUSSION

Sequential sampling models have been used to describe a variety of cognitive processes, including perceptual decision making (PDM; Luce, 1986; Usher and McClelland, 2001; Ratcliff and Smith, 2004). According to these models, sensory data is repeatedly sampled so that confidence or evidence for a particular choice option grows. Once enough evidence for an option exceeds a set threshold or criterion, the decision process terminates and the option with the most evidence determines the decision (and the time it takes for this process to finish, the response time). Such models possess parameters corresponding to latent cognitive processes involved in PDM, such as the response threshold (the amount of evidence needed before terminating the decision process), the drift rate (the rate of evidence accumulation), and non-decision time (time unrelated to the decision process itself, such as motor execution). Thus, these models are able to provide powerful explanations based on the response time (RT) and accuracy data from PDM tasks that extends beyond simply reporting mean RT and accuracy rate.

More recently, neurophysiological studies recording from individual neurons found that the firing rates of neurons involved with evidence accumulation can be fit with a sequential sampling model: these firing rates steadily increase until reaching a fixed threshold near the time of a behavioral response (Gold and Shadlen 2001; 2002; 2007; Shadlen and Newsome, 2001). The majority of these neurons were found in areas related to the planning of oculomotor commands (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Gold and Shadlen, 2000; 2003; Bichot et al., 2000; Horwitz et al., 2004; Ratcliff et

al., 2003; Kim and Shadlen, 1999), lending credence to the idea that the neurons underlying motor responses also compute perceptual decisions. Such findings were unsurprising given that making a perceptual decision was designed to be tantamount to selecting a particular motor response in these PDM tasks. However, in many everyday situations, a combination of motor responses is issued in response to a single stimulus (e.g., pressing on the brake while looking at a red light). While these single-unit recordings made major advancements in understanding the neural basis of PDM, the mechanisms underlying more general and abstract decisions have yet to be clarified. Moreover, understanding how the activity of individual neurons translates to population-level responses or to actual perception still remains unclear.

With the use of a simplified sequential sampling model (Brown and Heathcote, 2005), Ho et al. (2009) investigated the possibility of a general mechanism of PDM by measuring BOLD activation in human subjects performing PDM task with two different response modalities (Chapter 1). The right insula was the only region whose activation profile matched our simulations for an evidence accumulator across different response modalities. Two fMRI studies since then have also employed similar mathematical models to fit their data in order to predict the neural signature of an evidence accumulator (Kayser et al., 2010; Liu and Pleskac, 2011). Using similar criteria as Ho et al. (2009) to adjudicate the expected BOLD profile (i.e., responses should be higher on more low coherence trials), Kayser et al. (2010) found that the mean BOLD signal in the medial intraparietal sulcus (mIPS) was the only region to match their predictions. These results are not entirely inconsistent with the ones Ho et al. report, as the latter also observed bilateral regions of IPS with greater activation on more difficult trials. However, Ho et al. predicted latency

onsets in addition to larger BOLD responses on difficult trials, yet no parts of IPS exhibited the exact profile across both response modalities. Kayser and colleagues, on the other hand, required subjects to respond with only manual button presses, thereby possibly explaining the differences in these results. Liu and Pleskac (2011) adopted a task paradigm similar to Ho et al. (2009) and required subjects to respond either by making saccades or manual button presses. In addition, subjects were cued about response modality either in advance of the stimulus or after a delay. The authors isolated BOLD responses in independently defined sensorimotor areas, as well as task-defined non-sensorimotor areas – an approach that differed from Ho et al. since the latter conducted a whole brain exploratory analysis. To look for sensory evidence accumulators, Liu and Pleskac explored only areas where BOLD activation was greater on difficult trials (similar to what Ho et al. and Kayser et al. had done). This analysis revealed FEF, IPS, anterior insula, and inferior frontal sulcus as potential sensory evidence accumulators not dependent on response modality or foreknowledge. While Liu and Pleskac's criterion for an evidence accumulator was less stringent than Ho and colleagues – the former did not require differences in latency offsets between easy and hard trials – the results from both of these studies are nevertheless consistent. Ho et al. also found regions of FEF, IPS, and insula as being more active on hard trials. All together, these findings strengthen the account that modality-independent signals are also formed during perceptual decision making.

In order to map the activity of neural populations with perceptual performance, Ho et al. (2012) conducted a fMRI experiment while human subjects performed a fine discrimination between oriented gratings. Previous studies have demonstrated that such

fine discriminations utilize neurons that are tuned away from the target (termed *off-target neurons*) as opposed to those tuned to the target (*on-target neurons*). By using a forward encoding model that maps stimuli to population responses (Brower and Heeger, 2009; 2011), the authors found that the off-target neurons in V1 predicted trial-by-trial performance when subjects emphasized accuracy over response speed. While it is possible that attentional-feedback originating from regions outside of visual cortex (Higo et al., 2010; Purcell et al., 2010; Gold and Shadlen, 2007; Heekeren et al., 2004; de Lafuente and Romo, 2005; 2006; Lemus et al., 2010; Hernandez et al., 2010; Ho et al., 2009; Kayser et al., 2010) could explain this observed relationship between behavior and off-target modulations in V1, Ho et al. (2012) did not explicitly manipulate attention in this study. Thus, it was difficult to dissociate whether the changes observed were indeed solely or partly due to top-down biasing signals. Future studies could test these possibilities by adopting valid and neutral attentional cues in a similar paradigm in order to determine if speed pressure selectively impairs the ability to utilize prior information in optimally biasing population response profiles in visual cortex. Nevertheless, these findings provide support for how the integration and readout of sensory information impacts perceptual performance and are consistent with a set of other recent studies. In 2011, Kahnt and colleagues found that training-related improvements in performance on a difficult perceptual discrimination task could be explained by a greater efficiency in the readout of sensory information, thereby sharpening the representations of the decision variables leading up to the ultimate choice. Likewise, Rahnev et al. (2011) observed that manipulating prior expectation increased functional connectivity between posterior and frontal areas, consistent with an account where greater rates of sensory evidence transfer

from earlier visual areas to putative decision mechanisms. Pestilli and colleagues (2011) also jointly modeled psychophysical and BOLD data to determine whether sensory gain, noise reduction, or selective readout models could explain attentional modulations throughout visual cortices. The data were best fit by a model that included both additive response shifts in early visual areas and a selective readout rule, where inputs from the most responsive neurons are preferentially weighted. Thus, the present results complement other recent studies that emphasize the importance of efficient sensory readout in perceptual decision making.

In Chapter 3, Ho et al. (in press) tests the perceptual predictions of the feature-similarity gain model of attention, which postulates that feature-based attention selectively increases the gain of neuronal populations that are tuned to a relevant feature value while simultaneously suppressing the gain of neurons that are tuned to features dissimilar from the target (McAdams and Maunsell, 1999; Treue and Martinez-Trujillo, 1999; Martinez-Trujillo and Treue, 2004; Maunsell and Treue, 2006). Ling et al. (2009) reported that feature-based attention suppresses neurons tuned away from an attended feature, while White and Carrasco (2011) observed only attentional enhancement without suppression. However, neither study systematically varied the degree of similarity between the attended feature value and the target (as was done in Martinez-Trujillo and Treue, 2004), so the effect of attentional enhancement and suppression on the processing efficiency of unattended features was not directly evaluated. Ho and colleagues (in press) had subjects discern which of four random dot patterns contained the maximal amount of coherent motion (100% in Experiment 1 and 80% in Experiment 2). At the beginning of each trial,

subjects were neutrally, validly, or invalidly cued to the impending motion direction of the target (with 6 different offsets between the invalid cue and target). An analysis of drift rates revealed evidence for facilitation in both experiments and evidence of suppression in Experiment 2 (where distractors contained a low level of motion coherence so as to make the task more attentionally demanding). Changes in the drift rate can be intuitively linked to changes in the firing rate of sensory neurons or to changes in the relative weighting of sensory signals by downstream decision mechanisms. The current study proposed by Ho et al. (2012), however, is unable to differentiate between these two possibilities. Future investigations could implement a variant of the LBA model to estimate response caution for every possible offset between the invalid cue and distractor (analyzing specifically the trials where a distractor is incorrectly chosen). Lower response caution on trials when the distractors are similar to the invalid cue would provide support that the decision criteria are adjusted based on the similarity between the cue and the stimulus. Neurophysiological or neuroimaging approaches could also investigate whether or not attention-mediated changes in visual cortex are explained by changes in sensory gain – or if the addition of a mechanism that selectively reweights sensory responses during perceptual decision making is needed (see Pestilli et al., 2011).

The picture that emerges from these set of fMRI and psychophysical studies of perceptual decision making is consistent with the neurophysiological literature in some respects: off-target neurons provide the most sensory information for a fine discrimination task (Chapter 2) and attentional enhancement and suppression at the neuronal level appears to translate to improvement and impairment, respectively, in perceptual

performance (Chapter 3). However, evidence of a domain general mechanism of perceptual decision making is also clear (Chapter 1) and challenges much of the previous neurophysiological work. Furthermore, these studies demonstrate the utility of modeling psychophysical and neural data. Precisely quantifying psychological processes of interest for each individual and then relating these inter-subject differences with corresponding individual differences in brain measurements (e.g., BOLD fMRI, EEG, etc) provides a powerful tool to understanding the neural basis of various cognitive functions. Future progress in the understanding perceptual decision making will come from quantitative models that link reaction time distributions and accuracy data in the behavioral domain to firing rate data and neural patterns in the neuroimaging domain, forming a unified theory of the psychology and neurobiology of decisions.

REFERENCES

Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*, 52, 1106-1130.

Baldassi S, Verghese P (2002) Comparing integration rules in visual search. *Journal of Vision*, 2, 559-570.

Baldassi S, Verghese P (2005) Attention to locations and features: Different top-down modulation of detector weights. *Journal of Vision*, 5, 556-570.

Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. *Neuron*, 60, 1142-1152.

Bichot NP, Rao SC, Schall JD (2001) Continuous processing in macaque frontal cortex during visual search. *Neuropsychologia*, 39, 972-982.

Binder JR, Liebenthal E, Possing ET, Medler DA, Ward BD (2004) Neural correlates of sensory and decision processes in auditory object identification. *Nature Neuroscience*, 7, 295-301.

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. *Psychological Review*, 113, 700-765.

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

Boynton GM (2005) Imaging orientation selectivity: decoding conscious perception in V1. *Nature Neuroscience*, 8, 541-542.

Brainard DH (1997) The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.

Britten K, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12, 4745-4767.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, 10, 1157-1169.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, 13, 87-100.

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

Brown SD, Heathcote A (2008) The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153-178.

Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, 29, 3992-4003.

Brouwer GJ, Heeger DJ (2011) Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, 106, 2108-2119.

Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain's default network: anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124, 1-38.

Burnham KP, Anderson DR (2002) *Model selection and multimodel inference: A practical information theoretic approach* (2nd ed.) Springer: New York, NY.

Busemeyer JR, Townsend JT (1992) Fundamental derivations from decision field theory. *Mathematical Social Sciences*, 23, 255-282.

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

Busse L, Katzner S, Tillmann C, Treue S (2008) Effects of attention on perceptual Direction tuning curves in the human visual system. *Journal of Vision*, 8, 1-13.

Butts DA, Goldman MS (2006) Tuning curves, neuronal variability, and sensory coding. *PLoS Biology*, 4:e92. doi:10.1371/journal.pbio.0040092.

Buxton RB (2002) *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques*. Cambridge University Press: Cambridge, UK.

Carrasco M (2011) Visual attention: the past 25 years. *Vision Research*, 51, 1484-1525.

Churchland AK, Ditterich J (2012) New advances in understanding decisions among multiple alternatives. *Current Opinion in Neurobiology*, http://dx.doi.org/10.1016/j.conb.2012.04.009.

Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision making. *Neuron*, 69, 818-831.

Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nature Neuroscience*, 11, 693-702.

Clark HH, Chase WG (1972) On the process of comparing sentences against pictures. *Cognitive Psychology*, 3, 472-517.

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nature Neuroscience*, 14, 811-819.

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, 12, 1594-1600.

Cohen MR, Maunsell JHR (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*, 70, 1192-1204.

Conway BR, Livingstone MS (2003) Space-time maps and two-bar interactions of different classes of direction-selective cells in macaque V1. *Journal of Neurophysiology*, 89, 2726-2742.

Cook EP, Maunsell JH (2002) Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience*, 5, 985-994.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201-15.

Dale AM, Buckner RL (1997) Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, 5, 329-340.

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 Lafuent V, Romo R (2005) Neuronal correlates of subjective sensory experience. *Nature Neuroscience*, 8, 1698-1703.

de Lafuente V, Romo R (2006) Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proceedings of the National Academy of Sciences, USA*, 103, 14266-14271.

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193-222.

Dickman SJ, Meyer DE (1988) Impulsivity and speed-accuracy tradeoffs in information processing, *Journal of Personality and Social Psychology*, 54, 274-90.

Ditterich J, Mazurek ME, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, 6, 891-898.

Donkin C, Brown SD, Heathcote A (2009) The overconstraint of response time models: rethinking the scaling problem. *Psychonomic Bulletin & Review*, 16, 1129-1135.

Donkin C, Brown SD, Heathcote A (2011a) Drawing conclusions from choice response time models: a tutorial. *Journal of Mathematical Psychology*, 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? *Psychonomic Bulletin & Review*, 18, 61-69.

Douglas RJ, Martin KA (2007) Recurrent neuronal circuits in the neocortex. *Current Biology*, 17, R496-500.

Dosher BA, Lu ZL (1999) Mechanisms of perceptual learning.*Vision Research*, 39, 197-221.

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

Eckstein MP, Thomas JP, Palmer J, Shimozaki SS (2000) A signal detection model predicts the effects of set size on visual search for feature, conjunction, triple conjunction, and disjunction displays. *Attention, Perception, and Psychophysics*, 62, 425-451.

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

Felisberti FM, Zanker JM (2005) Attention modulates perception of transparent motion. *Vision Research*, 45, 2587-2599.

Fitts PM (1966) Cognitive aspects of information processing: III. Set for speed versus accuracy. *Journal of Experimental Psychology*, 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. *Proceedings of the National Academy of Sciences, USA*, 105, 17538-17542.

Forstmann BU, Anwander A, Schafer 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. *Proceedings of the National Academy of Sciences, USA*, 107, 15916-15920.

Freeman WT, Adelson EH (1991) The design and use of steerable filters. *IEEE Transactions on Pattern Analysis*, 13, 891-906.

Freeman J, Brouwer GJ, Heeger DJ, Merriam EP (2011) Orientation decoding depends on maps, not columns. *Journal of Neuroscience*, 31, 4792-4804.

Gitelman DR (2002) ILAB: a program for postexperimental eye movement analysis. *Behavior Research Methods, Instruments, and Computers*, 34, 605-612.

Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404, 390-394.

Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5, 10-16.

Gold JI, Shadlen MN (2001) Banburismus and the Brain: Decoding the Relationship between Sensory Stimuli, Decision, and Reward. *Neuron*, 36, 299-308.

Gold JI, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands *Journal of Neuroscience*, 23, 632-651.

Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535-574.

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

Greicius MD, Krasnow B, Reiss AL, Menon V (2003) Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences, USA*, 100, 253-258.

Grinband J, Hirsch J, Ferrera VP (2006) A neural representation of categorization uncertainty in the human brain. *Neuron*, 49, 757-763.

Gur M, Kagan I, Snodderly DM (2005) Orientation and direction selectivity of neurons in V1 of alert monkeys: functional relationships and laminar distributions. *Cerebral 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.

Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. *Science*, 274, 427-430.

Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, 9, 682-689.

Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, 8, 686-91.

Heathcote A, Brown S (2004) Beyond curve fitting? Comment on Liu, Mayer-Kress, and Newell (2003). *Journal of Motor Behvior*, 36, 225-238.

Heathcote A, Brown S, Mewhort DJ (2002) Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, 9, 394-401.

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

Heeger DJ, Ress D (2002) What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience*, 3, 142-151.

Heekeren HR, Marrett S, Bandettini, PA, Ungerleider LG (2004) A general mechanism for perceptual decision-making in the human brain. *Nature*, 14, 859-862.

Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG (2006) Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proceedings of the National Academy of Sciences, USA*, 103, 10023-10028.

Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. *Nature Review Neurosciences*, 9, 467-479.

Herrmann K, Heeger DJ, Carrasco M (2012) Feature-based attention enhances performance by increasing response gain. *Vision Research*, http://dx.doi.org/10.1016/j.visres.2012.04.016.

Hernandez A, Nacher V, Luna R, Zainos A, Lemus L, Alvarez M, Vazquez 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. *Proceedings of the National Academy of Sciences, USA*, 108, 4230-4235.

Ho TC, Brown S, Serences JT (2009) Domain general mechanisms of perceptual decision making in human cortex. *Journal of Neuroscience*, 29, 8675-8687.

Ho T, Brown SD, Van Maanen L, Forstmann BU, Wagenmakers EJ, Serences JT (2012). *Journal of Neuroscience*, 32, 7992-8003.

Ho TC, Brown S, Abuyo NA, Ku EJ, Serences JT (in press). *Journal of Vision*.

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

Horwitz GD, Newsome WT (1999) Separate signals for target selection and movement specification in the superior colliculus. *Science*, 284, 1158-1161.

Horwitz GD, Newsome WT (2001) Target selection for saccadic eye movements: Prelude activity in the superior colliculus during a direction discrimination task. *Journal of Neurophysiology*, 86, 2543-2558.

Horwitz GD, Batista AP, Newsome WT (2004) Representation of an abstract perceptual decision in macaque superior colliculus. *Journal of Neurophysiology*, 91, 2281-2296.

Hubner R, Steinhauser M, Lehle C (2010) A dual-stage two-phase model of selective attention. *Psychological Review*, 117, 759-784.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, 25, 10420-10436.

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. *Nature Neuroscience*, 9, 690-696.

Jia X, Smith MA, Kohn A (2011) Stimulus selectivity and spatial coherence of gamma

components of the local field potential. *Journal of Neuroscience*, 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 TF (2005) Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8, 679-685.

Kay KN, Gallant JL (2009) I can see what you see. *Nature Neuroscience*, 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. *Journal of Neurophysiology*, 103, 1179-94.

Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. Annu Rev Neurosci. 23, 315-41.

Khayat, P.S., Neibergall, R., Martinez-Trujillo, J.C. (2010) Attention differentially modulates similar neuronal responses evoked by varying contrast and direction stimuli in area MT. *Journal of Neuroscience*, 30, 2188-2197.

Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *Journal of Neuroscience*, 28, 3017-3029.

Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2, 176-185.

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

Kriegeskorte N, Simmons WK, Bellgowan PS, Baker CI (2009) Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, 12, 535-540.

LaBerge DA (1962) A recruitment theory of simple behavior. *Psychometrika*, 27, 275-396.

Laming DRJ (1968) *Information theory of choice reaction time*. New York: Wiley.

Law CT Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, 11, 505-13.

Law CT Gold JI (2009) Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, 12, 655-63.

Lemus L, Hernandez 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 receptive field position of neurons in cat striate cortex. *Journal of Comparative Neurology*, 220, 476-483.

Leon MI, Shadlen MN (1999) Effect of reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron*, 24, 415-425.

Ling S, Liu T, Carrasco M (2009) How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research*, 49, 1194-120.

Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrikia*, 40, 77-105.

Liu T, Stevens ST, Carrsaco M (2007) Comparing the time course and efficacy of spatial and feature-based attention. *Vision Research*, 47, 108-113.

Liu T, Hou Y (2011) Global feature-based attention to orientation. *Journal of Vision*, 11, 1-8.

Liu T, Pleskac TJ (2011) Neural correlates of evidence accumulation in a perceptual decision task. *Journal of Neurophysiology*, 106, 2383-2398.

Livingstong MS (1998) Mechanisms of direction selectivity in macaque V1. *Neuron*, 20, 509-526.

Livingstone MS, Conway BR (2003) Substructure of direction-selective receptive fields in macaque V1. *Journal of Neurophysiology*, 89, 2743-2759.

Lo CC, Wang XJ (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience*, 9, 956-963.

Loftus GR, Masson MEJ (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, 1, 476-490.

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

Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. *Annual Review of Physiology,* 66, 735-69.

Luce RD (1986) *Response times*. New York: Oxford University Press.

Luck SJ, Massimo G, McDermott MT, Ford MA (1997) Bridging the gap between monkey neurophysiology and human perception: an ambiguity resolution theory of visual selective attention. *Cognitive Psychology*, 33, 64-87.

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

Martinez-Trujillo JC, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*, 35, 365-370.

Martinez-Trujillo, JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology,* 14, 744-751.

Mazurek ME, Roitman JD, Ditterich, J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cerebral Cortex,* 13, 1257-1269.

McAdams CJ, Maunsell JHR (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, 19, 431-441.

Meier JD, Aflalo TN, Kastner S, Graziano MS (2008) Complex organization of human primary motor cortex: a high-resolution fMRI study. *Journal of Neurophysiology*, 100, 1800-1812.

Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulations across cell classes in macaque visual area V4. *Neuron*, 55, 131-141.

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron,* 63, 879-888.

Moore BCJ (2008) Basic auditory processes involved in the analysis of speech sounds. *Philosophical Transactions of the Royal Society B: Biological Sciences,* 363, 947-963.

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

Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL (2009) Bayesian reconstruction

of natural images from human brain activity. *Neuron*, 63, 902-915.

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. *Computer Journal*, 7, 308-313.

Newsome WT, Paré EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience*, 8, 2201-2211.

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature*, 341, 52-54.

Newsome WT, Shadlen MN, Zohary E, Britten KH, Movshon JA (1995) Visual motion: linking neuronal activity to psychophysical perofmrance. *The Cognitive Sciences*, M.S. Gazzaniga, Ed. MIT Press.

Noudoost B, Chang MH, Steinmetz NA, Moore T (2010) Top-down control of visual attention. *Current Opinion in Neurobiology*, 20, 183-90.

Otter T, Johnson J, Rieskamp J, Allenby GM, Brazell JD, Diederich A, Hutchinson JW, MacEachern, Ruan S, Townsend J (2008) Sequential sampling models of choice: some recent advances. *Marketing Letters*, 19, 255-267.

Palmer J (1995) Attention in visual search: Distinguishing four causes of set-size effects. *Current Directions in Psychological Science*, 4, 118-123.

Palmer J, Alexander CH, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5, 376-404.

Pelli DG (1997) The Video Toolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-442.

Pessoa L, Padmala S (2005) Quantitative prediction of perceptual decisions during near-threshold fear detection. *Proceedings of the National Academy of Sciences, USA*, 102, 5612-5617.

Pessoa L, Padmala S (2007) Decoding near-threshold perception of fear from distributed single-trial brain activation. *Cerebral Cortex*, 17, 691-701.

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.

Philiastides MG, Sajda P (2007) EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. *Journal of Neuroscience,* 27, 13082-13091.

Philiastides MG, Ratcliff R, Sajda P (2006) Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. J Neurosci 26, 8965-8975.

Pleger B, Ruff CC, Blankenburg F, Bestmann S, Wiech K, Stephan KE, Capilla A, Friston KJ, Dolan RJ (2006) Neural coding of tactile decisions in the human prefrontal cortex. *Journal of Neuroscience*, 26, 12596-12601.

Ploran EJ, Nelson SM, Velanova K, Donaldson DI, Petersen SE, Wheeler ME (2007) Evidence accumulation and the moment of recognition: dissociating perceptual recognition processes using fMRI. *Journal of Neuroscience*, 27, 11912-11924.

Preuschhof C, Heekeren HR, Taskin B, Schubert T, Villringer A (2006) Neural correlates of vibrotactile working memory in the human brain. *Journal of Neuroscience*, 26, 13231-13239.

Purcell, BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ (2010) Neurally constrained modeling of perceptual decision making. *Psychological Review*, 117, 1114-43.

Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. *Nature Neuroscience*, 8, 99-106.

Nadler JW, DeAngelis GC (2005) Precision pooling predicts primate perceptual performance. *Nature Neuroscience*, 8, 12-13.

Raftery AE (1995) Bayesian Model Selection in Social Research. *Sociological Methodology*, 25, 111-163.

Rahnev D, Lau H, de Lange FP (2011a) Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *Journal of Neuroscience*, 31, 10741-10748.

Rahnev D, Maniscalco B, Graves T, Huang E, de Lange FP, Lau H (2011b) Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, 14, 1513-1515.

Raichle ME, Snyder AZ (2007) A default mode of brain function: a brief history of an evolving idea. *Neuroimage*, 37, 1083-1090; discussion 1097-1099.

Ratcliff R (1978) A theory of memory retrieval. *Psychological Review*, 85, 59-108.

Ratcliff R (1985) Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, 92, 212-225.

Ratcliff R, Cherian A, Segraves M (2003) A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two choice decisions. *Journal of Neurophysiology*, 20, 1392-1407.

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

Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction times. *Psychological Review*, 111, 333-367.

Ratcliff R, Gomez P, McKoon G (2004) A diffusion model account of the lexical decision task. *Psychological Review*, 111, 159-182.

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

Ratcliff R, Tuerlinckx F (2002) Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438.481.

Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nature Neuroscience*, 3, 827-830.

Reeves A, Santhi N, Decaro S (2005) A random-ray model for speed and accuracy in perceptual experiments. *Spatial Vision*, 18, 73-83.

Regan D, Beverley KI (1985) Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, 2, 147-155.

Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron*, 61, 168-185.

Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19, 1736-1753.

Reynolds, J.H., Pasternak, T., Desimone, R. (2000) Attention increases sensitivity of V4 neurons. *Neuron*, 26, 703-714.

Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *Journal of Neuroscience*, 22, 5639-5651.

Ringach DL, Bredfeldt CE, Shapley RM, Hawken MJ (2002) Suppression of neural responses to nonoptimal stimuli correlates with tuning selectivity in macaque V1. *Journal of Neurophysiology*, 87, 1018-1027.

Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22, 9475-9489.

Romo R, Salinas E (1999) Sensing and deciding in the somatosensory system. *Current Opinion in Neurobiology*, 9, 487-493.

Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nature Review Neuroscience*, 4, 203-218.

Romo R, Hernandez A, Zainos A, Brody C, Salinas E (2002) Exploring the cortical evidence of a sensory-discrimination process. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357, 1039-1051.

Sàenz M, Buraĉas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, 5, 631-632.

Sàenz M, Buraĉas GT, Boynton GM (2003) Global feature-based attention for motion and color. *Vision Research*, 43, 629-637.

Salzman CD, Britten KH, Newsome WT (1990) Cortical microstimulation influences perceptual judgments of motion direction. *Nature*, 346, 174-177.

Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *Journal of Neuroscience*, 12, 2331-2355.

Salzman CD, Newsome WT (1994) Neural mechanisms for forming a perceptual decision. *Science*, 264, 231-237.

Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RBH (2006) The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*, 51, 661.670.

Schall JD (2001) Neural basis of deciding, choosing and acting. *Nature Review Neuroscience*, 2, 33-42.

Schall JD (2003) Neural correlates of decision processes: neural and mental chronometry. *Current Opinion Neurobiology*, 13, 182-186.

Schall JD, Vitek DJ, Leventhal AG (1986) Retinal constraints on orientation specificity in cat visual cortex. *Journal of Neuroscience*, 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. *Journal of Neurophysiology*, 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.

Schwarz G (1978) Estimating the dimension of a model. *Annuals of Statistics*, 6, 461-464.

Scolari M, Byers A, Serences JT (2012) Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience*, 32, 7723-7733.

Scolari M, Serences, JT (2009) Adaptive allocation of attentional gain. *Journal of Neuroscience*, 29, 11933-11942.

Scolari M, Serences JT (2010) Basing perceptual decisions on the most informative sensory neurons. *Journal of Neurophysiology*, 104, 2266-2273.

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.

Serences JT, Yantis S (2006) Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, 10, 38-45.

Serences JT, Boynton GM. (2007a) The representation of behavioral choice for motion in human visual cortex. *Journal of Neuroscience*, 27, 12893-12899.

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

Serences JT, Saproo S (2011) Computational advances towards linking BOLD and behavior. *Neuropsychologia*, 50, 435-446.

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 RR, Tootel RBH (1995) Borders of multiple visual areas in human revealed by functional magnetic resonance imaging. *Science*, 268, 889-893.

Shadlen MN, Newsome WT (1996) Seeing and deciding. *Proceedings of the National Academy of Sciences, USA*, 93, 628-633.

Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916-1936.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience*, 16, 1486-1510.

Shulman GL, McAvoy MP, Cowan MC, Astafiev SV, Tansy AP, d'Avossa G, Corbetta M (2003) Quantitative analysis of attention and detection signals during visual search. *Journal of Neurophysiology*, 90, 3384-3397.

Simen P (2012) Evidence accumulator or decision threshold – which cortical mechanism are we observing? *Frontiers in Psychology*, 3, 183, doi: 10.3389/fpsyg.2012.00183.

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

Smith PL (1995) Psychophysically principled models of visual simple reaction time. *Psychological Review*, 102, 567-591.

Smith PL, Ratcliff R, Wolfgang BJ. (2004) Attention orienting and the time course of perceptual response time distributions with masked and unmasked displays. *Vision Research*, 44, 1297-1320.

Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27, 161-168.

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

Sternberg S (1969) The discovery of processing stages: Extensions of Donder's method. In W.G. Koster (Ed.), *Attention and performance II* (pp. 276-315). Amsterdam: North Holland.

Swendale NV(1998) Orientation tuning curves: empirical description and estimation of parameters. *Biological Cybernertics*, 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. *Journal of Neuroscience*, 30, 325-330.

Talairach J, Tournoux P (1988) *Co-planar sterotaxic atlas of the human brain*. New York: Thieme Medical Publishers.

Tegenthoff M, Ragert P, Pleger B, Schwenkreis P, Forster AF, Nicolas V, Dinse HR (2005) Improvement of tactile discrimination performance and enlargement of cortical somatosensory maps after 5 Hz rTMS. *PLoS Biology*, 3:e362.

Thielscher A, Pessoa L (2007) Neural correlates of perceptual choice and decision making during fear-disgust discrimination. *Journal of Neuroscience*, 27, 2908-2917.

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.

Thompson KG, Hanes DP, Bichot NP, Schall JD (1996) Perceptual and motor processing stages identified in frontal eye field neurons during visual search. *Journal of Neurophysiology*, 76, 4040-4055.

Tombu M, Tsotsos JK (2008) Attending to orientation results in an inhibitory surround in orientation space. *Attention, Perception, and Psychophysics*, 70, 30-35.

Tosoni A, Galati G, Romani GL, Corbetta M (2008) Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. *Nature Neuroscience*, 11, 1446-1453.

Townsend JT, Ashby FG (1983) *Stochastic modeling of elementary psychological processes*. Cambridge, England: Cambridge University Press.

Treue S, Martinez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399, 575-579.

Tsotsos JK, Culhane SM, Wai WYK, Lai Y, Davis N, Nuflo F (1995) Modeling visual attention via selective tuning. *Artificial Intelligence*, 78, 507-545.

Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108, 550-592.

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. *Journal of Neuroscience*, 31, 17488-17495.

Van Veen V, Krug MK Carter CS (2008) The neural and computational basis of

controlled speed-accuracy tradeoff during task performance. *Journal of Cognitive Neuroscience*, 20, 1952-1965.

Vandekerckhove J, Tuerlinckx F (2007) Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, 14, 1011-1026.

Vandekerckhove J, Tuerlinckx F (2008) Diffusion model analysis with MATLAB: A DMAT primer. *Behavior Research Methods*, 40, 61-72.

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

Voss A, Rothermund K, Voss J (2004) Interpreting the Parameters of the diffusion model: an empirical validation. *Memory & Cognition*, 32, 1206-1220.

Van Zandt T, Colonius H, Proctor RW (2000) A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7, 208-256.

Vickers D, Caudrey D, Willson RJ (1971) Discriminating between the frequency of occurrence of two alternative events. *Acta Psychologica*, 35, 151-172.

Vickers D (1979) *Decision processes in visual perception*. New York: Academic Press.

Vul E, Kanwisher N (2010) Begging the question: The non-independence error in fMRI data analysis. In: *Foundational Issues for Human Brain Mapping* (Hanson S, Bunzl M, eds). Cambridge, MA: MIT Press.

Vul E, Harris C, Winkielman P, Pashler H (2009) Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition. *Perspectives on Psychological Science*, 4, 274-290.

Wagenmakers, E.J., van der Maas, H.L., Grasman, R.P. (2007) An EZ-diffusion model for response time and accuracy. *Psychonomic Bulletin & Review*, 14, 3-22.

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

Watson AB, Pelli DG (1983) QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33, 113-120.

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

White AL, Carrasco M (2011) Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of Vision*, 11, 1-10.

Williford T, Maunsell JH (2006) Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, 96, 40-54.

Wickelgren W (1977) Speed-accuracy tradeoff and information-processing dynamics. *Acta Psychologica*, 41, 67-85.

Wolfe JM, Van Wert MJ (2010) Varying target prevalence reveals two, dissocial decision criteria in visual search. *Current Biology*, 20, 121-124.

Woodworth RS (1899) The accuracy of voluntary movement. *Psychological Review*, 3, 1-114.

Yantis S (2008) The Neural Basis of Selective Attention: Cortical Sources and Targets of Attentional Modulation. *Current Directions in Psychological Science*, 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. *Society for Neuroscience Abstract.*