CHAPTER

15

Neurocognitive Modeling of Perceptual Decision Making

Thomas J. Palmeri, Jeffrey D. Schall, and Gordon D. Logan

Abstract

Mathematical psychology and systems neuroscience have converged on stochastic accumulator models to explain decision making. We examined saccade decisions in monkeys while neurophysiological recordings were made within their frontal eye field. Accumulator models were tested on how well they fit response probabilities and distributions of response times to make saccades. We connected these models with neurophysiology. To test the hypothesis that visually responsive neurons represented perceptual evidence driving accumulation, we replaced perceptual processing time and drift rate parameters with recorded neurophysiology from those neurons. To test the hypothesis that movement related neurons instantiated the accumulator, we compared measures of neural dynamics with predicted measures of accumulator dynamics. Thus, neurophysiology both provides a constraint on model assumptions and data for model selection. We highlight a gated accumulator model that accounts for saccade behavior during visual search, predicts neurophysiology during search, and provides insights into the locus of cognitive control over decisions.

Key Words: accumulator models, decision making, response time, visual search, stop task, countermanding, neurophysiology, computational modeling, neural modeling, frontal eye field, superior colliculus

Introduction

We make decisions all the time. Whom to marry? What car to buy? What to eat? Whether to turn left or right? Some are easy. Some are hard. Some involve uncertainty. Some involve risk or reward. Decision-making requires integrating our perceptions of the current environment with our knowledge and past experience and our assessments of uncertainty and risk in order to select a possible action from a set of alternatives. Behavioral research on decision-making has had a long and distinguished history in psychology (e.g., Kahneman & Tversky, 1984). We now have powerful computational and mathematical models of how decisions are made (e.g., Brown & Heathcote, 2008; Busemeyer & Townsend, 1993; Dayan & Daw, 2008; Ratcliff & Rouder, 1998). And we know more about the brain areas involved in a range of decision-making tasks (Glimcher & Rustichini, 2004; Heekeren, Marrett, & Ungerleider, 2008; Schall, 2001; Shadlen & Newsome, 2001). To develop an integrated understanding of decision-making mechanisms, new efforts aim to combine behavioral and neural measures with cognitive modeling (e.g., Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011; Gold & Shadlen, 2007; Palmeri, in press; Smith & Ratcliff, 2004), an approach we aim to illustrate in some detail here.

We focus on perceptual decisions. Perceptual decision-making involves perceptually representing the world with respect to current task goals and using perceptual evidence to inform the selection of an action. A broad class of accumulator models of perceptual decision-making assume that perceptual evidence accumulates over time to a response threshold (e.g., Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Brown & Heathcote, 2008; Link, 1992; Nosofsky & Palmeri, 1997; Palmeri, 1997; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff & Smith, in press; Smith & Van Zandt, 2000; Usher & McClelland, 2001; see also Nosofsky & Palmeri, 2015). These models have provided excellent accounts of observed behavior, including the choices people make and the time it takes them to decide. Moreover, the observation that the pattern of spiking activity of certain neurons resembles an accumulation to threshold (Hanes & Schall, 1996) has sparked exciting synergies of mathematical and computational modeling with systems neuroscience (e.g., Boucher, Palmeri, Logan, & Schall, 2007a; Churchland & Ditterich, 2012; Cisek, Puskas, & El-Murr, 2009; Ditterich, 2006, 2010; Mazurek, Roitman, Ditterich, & Shadlen, 2003; Purcell, Heitz, Cohen, Schall, Logan, & Palmeri, 2010; Purcell, Schall, Logan, & Palmeri, 2012; Ratcliff, Cherian, & Segraves, 2003; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007; Wong, Huk, Shadlen, & Wang, 2007; Wong & Wang, 2006). In this article, we provide a general review of our contributions to these efforts. We use variants of accumulator models to explain neural mechanisms, use neurophysiology to constrain model assumptions, and use neural and behavioral data as a tool for model section.

Our specific focus has been on perceptual decisions about where and when to make a saccadic eye movement to objects in the visual field. The first section of this article, *Perceptual Decisions by Saccades*, provides an overview of behavior, neuroanatomy, and neurophysiology of the primate saccade system, with an emphasis on the *frontal eye field* (FEF). There are numerous practical advantages to studying perceptual decisions made by saccades over perceptual decisions made by finger, hand, or limb movement and we can also capitalize on over two decades of careful systems neuroscience research with awake behaving monkeys characterizing the response properties of neurons in FEF and the interconnected network of other brain areas involved in saccadic eye movements (Figure 15.1). FEF itself provides physiologists and theoreticians a unique window on perceptual decision-making. FEF receives projections from a wide range of posterior brain areas involved in visual perception, projects to subcortical brain areas involved directly in the production of eye movements, and is modulated by prefrontal brain areas involved in cognitive control. Indeed, one class of *visually responsive neurons* in FEF represent task-relevant salience of objects in the visual field, whereas another class of *movement-related neurons* increase their activity in a manner consistent with accumulation of evidence models and modulate their activity according to changing task demands (e.g., see Schall, 2001, 2004).

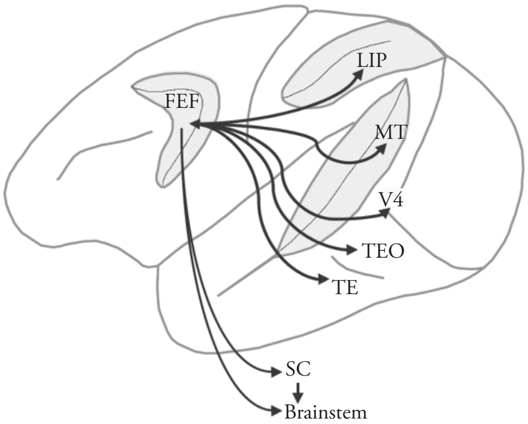

Fig. 15.1 Illustration of the macaque cerebral cortex. Frontal eye field (FEF) is a key brain area involved in the production of saccadic eye movements and the focus of our recent work. It receives projections from numerous posterior visual areas, including the middle temporal area (MT), visual area V4, inferotemporal areas TE and TEO, and the lateral intraparietal area (LIP). FEF projects to the superior colliculus (SC). Both FEF and SC project to the brainstem saccade generators that ultimately control the muscles of the eyes. Not shown are connections between FEF and prefrontal cortical areas and areas of the basal ganglia. (Adapted from Purcell et al., 2010.)

One form of an accumulator model is illustrated in Figure 15.2. Accumulator models assume that perceptual processing takes some amount of time. The product of perceptual processing is perceptual evidence that is accumulated over time to make a perceptual decision. The rate of accumulation is often called *drift rate*, and this drift rate can be variable within a trial, across trials, or both (e.g., Brown & Heathcote, 2008; Ratcliff & Rouder, 1998). Variability in the accumulation of perceptual evidence to a threshold is a major contributor to variability in predicted behavior.

In their most general form, accumulator models assume drift rates to be free parameters that can be optimized to fit a set of observed behavioral data. There has been concern that unrestricted assumptions about drift rate and its variability may imbue these models with too much flexibility (Jones & Dzhafarov, 2014; but see also Ratcliff,

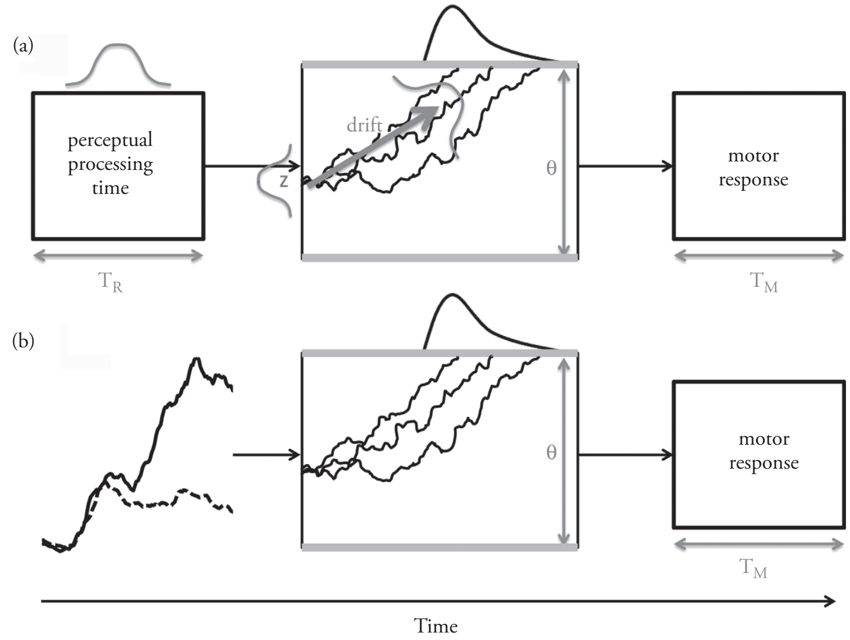

Fig. 15.2 (a) Illustration of a classic stochastic accumulator model of perceptual decision-making, highlighting some of the key free parameters. Perceptual processing of a visual stimulus takes some variable amount time with mean TR. The outcome of perceptual processing is noisy perceptual evidence in favor of competing decisions with some mean drift rate. Perceptual evidence is accumulated over time, originating at some variable starting point (z), and accumulating until some threshold is reached, determined by θ Illustrated here is a drift-diffusion model, but different architectures for the perceptual decision-making process can be assumed (see Figure 15.5). Variability in the accumulation of evidence to a threshold is a key constituent in predicting variability in RT. A motor response is made with some time TM, which for saccadic eye movements is on the order of 10-20ms. (b) Our recent work has tested whether many of the free parameters can be constrained by the observed physiological dynamics of one class of neurons in FEF (see Figure 15.5) and whether predicted model dynamics of the stochastic accumulator can predict observed physiological dynamics of another class of neurons in FEF (see Figure 15.8).

2013). One important step in theory development has been to significantly constrain these models by creating theories of the drift rates driving the accumulation of evidence, linking models of perceptual decision making with models of perceptual processing (e.g., Ashby, 2000; Logan & Gordon, 2001; Mack & Palmeri, 2010, 2011; Nosofsky & Palmeri, 1997; Palmeri, 1997; Palmeri & Cottrell, 2009; Palmeri & Tarr, 2008; Schneider & Logan, 2005, 2009; Smith & Ratcliff, 2009). As a first step toward a neural theory of drift rates, we hypothesized that activity of visually responsive neurons in FEF represent perceptual evidence driving the accumulation to threshold. To test this hypothesis, as described in the section titled A Neural Locus of Drift Rates, we replaced perceptual processing-time and drift-rate parameters directly with recorded neurophysiology from these neurons (see Figures 15.2 and 15.5), testing whether any model architecture for accumulation of perceptual evidence could then quantitatively account for observed saccade response probabilities and response time distributions.

A number of different model architectures have been proposed that all involve some accumulation of perceptual evidence to a threshold (e.g., see Bogacz et al., 2006; Smith & Ratcliff, 2004). For example, as their name implies, independent race models assume that evidence for each alternative decision independently (Smith & Van Zandt, 2000; Vickers, 1970). Drift-diffusion models (Ratcliff, 1978; Ratcliff & Rouder, 1998) and random walk models (Laming, 1968; Link, 1992; Nosofsky & Palmeri, 1997; Palmeri, 1997) assume that perceptual evidence in favor of one alternative counts as evidence against competing alternatives. Competing accumulator models (Usher & McClelland, 2001) assume that support for various alternatives is mutually inhibitory, so as evidence in favor of one alternative grows, it inhibits the others, often in a winner-take-all fashion (Grossberg, 1976). Different models can vary in other respects as well, such as whether integration of evidence is perfect or leaky. We describe these alternative model architectures and how well they account for observed response probabilities and response time distributions in the section Architectures for Perceptual Decision Making.

We also tested the hypothesis that movement-related neurons in FEF instantiate an accumulator (Hanes & Schall, 1996). As described in the section Predicting Neural Dynamics, we quantitatively compared measured metrics of neural dynamics with predicted metrics of accumulator dynamics. Neurophysiology and modeling are synergistic in that we test quantitatively whether movement-related neurons have dynamics predicted by accumulator models, and we use the measured neural dynamics of movement-related neurons as an additional tool to select between competing model architectures. Finally, in a complementary way, in the section Control over Perceptual Decisions, we test whether competing hypotheses about cognitive control mechanisms can predict observed behavior as well as the observed modulation of movement-related neurons dynamics.

Perceptual Decisions by Saccades

Significant insights into the neurophysiological basis of perceptual decision-making have come from research on decisions about where and when to move the eyes (e.g., Gold & Shadlen, 2007; Schall, 2001, 2004; Smith & Ratcliff, 2004). Although the majority of human research on perceptual decisions has used manual key-press responses, a neurophysiological focus on saccadic eye movements is justified on several grounds: From the perspective of effect or dynamics and motor control, eye movements have relatively few degrees of freedom, far fewer than limb movements, allowing fairly direct links between neurophysiology and behavior to be established (Scudder, Kaneko, & Fuchs, 2002). Saccadic eye movements are also relatively ballistic, with movement dynamics quite stereotyped depending on the direction, starting point, and distance the eyes need to move (Gilchrist, 2011), unlike limb movement, which can reach the same endpoint using a multitude of different trajectories having vastly different temporal dynamics (Rosenbaum, 2009). Moreover, from the perspective of understanding the mechanisms by which perceptual evidence is used to produce a perceptual decision, the saccade system is also a choice candidate to study because of the Frontal Eye Field (FEF), an area where visual perception, motor production, and cognitive control come together in the primate brain (Schall & Cohen, 2011).

FEF has long been known to play a role in the production of saccadic eye movements (e.g., Bruce, Goldberg, Bushnell, & Stanton, 1985; Ferrier, 1874). This is reflected by its direct and indirect connectivity with the superior colliculus (SC) and brain stem nuclei necessary for the production of saccadic eye movement (e.g., Munoz & Schall, 2004; Scudder et al., 2002; Sparks, 2002), as illustrated in Figure 15.1. Also as illustrated, FEF is innervated by numerous dorsal and ventral stream areas of extrastriate visual cortex (Schall, Morel, King, & Bullier, 1995). Not illustrated are connections between FEF and brain areas implicated in cognitive control, such as medial frontal and dorsolateral prefrontal cortex (e.g., Stanton, Bruce, & Goldberg, 1995) and basal ganglia (Goldman-Rakic & Porrino, 1985; Hikosaka & Wurtz, 1983). Neuroanatomically, FEF lies at a juncture of perception, action, and control. This bears out functionally, as various neurons within FEF reflect the importance of objects in the visual field,signal the selection and timing of saccadic eye movements, and modulate in a controlled manner according to changing task conditions (e.g., Heitz & Schall, 2012; Murthy, Ray, Shorter, Schall, & Thompson, 2009; Thompson, Biscoe, & Sato, 2005).

At the start of each neurophysiological session, once a neuron in FEF has been isolated, a memory-guided saccade task is used to classify its response properties (Bruce & Goldberg, 1985). As illustrated in Figure 15.3, the monkey fixates a spot in the center of the screen while a target is flashed in the periphery. To earn reward, the monkey must maintain fixation for a variable amount of time after which the fixation spot disappears and then the monkey must make a single saccade to the remembered target location. When the target is in the receptive field of the FEF neuron, that neuron is classified as a *visually responsive neuron* (or visual neuron) if it shows a vigorous response to the appearance of the target, perhaps with a tonic response during the delay period, but with no significant saccade-related modulation. The neuron is classified as a *movement-related neuron* (or movement neuron, sometimes referred to as a buildup neuron) if it shows no or very weak modulation to the appearance of the target but pronounced growth of spike rate immediately

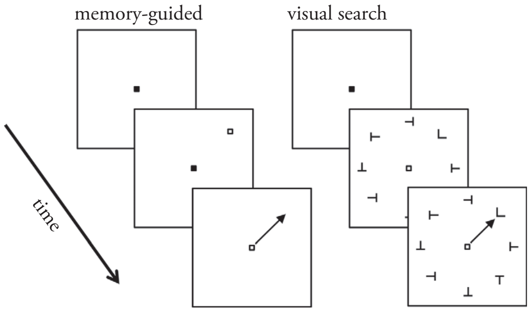

Fig. 15.3 Illustration of two saccade decision tasks discussed in this article. (a) In a memory-guided saccade task, the monkey fixates a central point while a peripheral target is quickly flashed; the location of the target is guided by the receptive field properties of the isolated neuron for a given experimental session. The monkey is required to maintain fixation for 400–1000ms, after which the fixation spot disappears. To earn reward, the monkey must make a single saccade to the remembered location of the peripheral target. (b) In a visual search task, the monkey first maintains fixation on a central point. An array of visual objects is then presented and to earn reward the monkey must make a single saccade to the target object and not one of the distractor objects. In this case, the reward target was an L and the distractors were variously rotated Ts, with the particular reward target changed from session to session. Various experiments manipulated the number of distractors (set size), the similarity between targets and distractors, and the particular dimensions on which targets and distractors differed (shape, color, or motion).

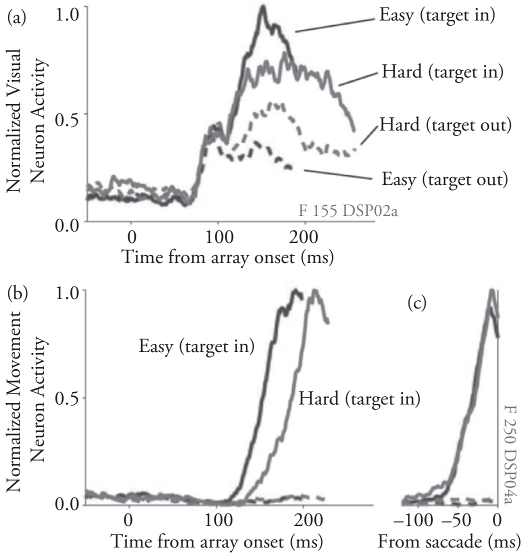

Fig. 15.4 Illustration of response properties of visually responsive and movement-related neurons in FEF (Hanes, Patterson, & Schall, 1998; Hanes & Schall, 1996; Purcell et al., 2010). Recordings were made while monkeys engaged in a visual search task where the target either appeared among dissimilar distractors (easy search) or among similar distractors (hard search). Plots display normalized spike rate as a function of time (ms). Visually responsive neuron activity aligned on visual search array onset time illustrated in panel (a), movement-related neuron activity aligned on visual search array onset time illustrated in panel (b), and movement-related neuron activity aligned on saccade time illustrated in panel (c). Solid lines are trials in which the target was in the visual neuron's receptive field or movement neuron's movement field (target in), and dashed lines are trials in which the target was outside the neurons' response fields (target out). (Adapted from Purcell et al., 2010.)

preceding saccade production. Other neurons in FEF show other response properties (e.g., Sato & Schall, 2003), but our recent work has focused primarily on visual and movement neurons, which we might loosely characterize as the incoming input signal and outgoing output signal from FEF (see also Pouget et al., 2009).

Once visually responsive neurons and movement-related neurons are identified, their response properties can be measured during a primary perceptual decision task. For example, in a visual search task, as illustrated in Figure 15.3, after the monkey fixates a central spot, a search array is shown containing a target (in this case an L) and several distractors (in this case rotated Ts) and the monkey must make a single saccade to the target in order to receive reward. During visual search, visually responsive and movement-related neurons display characteristic dynamics. Figure 15.4 shows the normalized spiking activity of representative neurons recorded during easy and hard visual search trials when the target (solid) or a distractor (dashed) was in the neuron's receptive field. For some time after the visual search array appears, visually responsive neurons (Figure 15.4a) show no discrimination between a target and a distractor. However, spiking activity eventually discriminates between target and distractor, with generally faster and more significant discrimination with easy compared to hard visual search trials (Bichot & Schall, 1999; Sato, Murthy, Thompson, & Schall, 2001) and small compared to large set sizes (Cohen, Heitz, Woodman, & Schall, 2009). We note that the particular shape of the trajectories taken to achieve this neural discrimination can be somewhat heterogeneous across different neurons, but virtually all visually responsive neurons discriminate target from distractor over time. We emphasize that this discrimination concerns the "targetness" of the object in the neuron's receptive field, not particular features or dimensions of the object like its color or shape, except under unique circumstances (Bichot, Schall, & Thompson, 1996). Visually responsive neurons display these same characteristic dynamics regardless of whether a saccade is made, such as when the monkey withholds or cancels an eye movement because of a stop signal (Hanes, Patterson, & Schall, 1998) or when the monkey is trained to maintain fixation and respond with a limb movement and not an eye movement (Thompson, Biscoe, & Sato, 2005).

Normalized activity of a representative movement-related neuron is shown aligned on the onset time of the visual search array (Figure 15.4b) and aligned on the time of the saccade (Figure 15.4c). When the monkey makes a saccade to the object in the receptive field (movement field) of the neuron, there is a characteristic buildup of activity some time after array onset; there is far less activity when the nonselected object is in the receptive field, although the precise nature of those dynamics varies somewhat from neuron to neuron. We see clearly that, when aligned on saccade initiation time, activity reaches a relatively constant threshold level immediately prior to the eye movement (Hanes & Schall, 1996), and this pattern of activity holds across search difficulty and set size (Woodman, Kang, Thompson, & Schall, 2008). Movement-related neuron activity does not reach threshold if the monkey withholds or cancels an eye movement because of a stop signal (Hanes et al., 1998; Murthy et al., 2009) or makes a response to the target using a limb movement and not an eye movement (Thompson, Biscoe, & Sato, 2005).We discuss more detailed aspects of the temporal dynamics of movement-related neurons later in this article. One of our primary goals has been to develop models that both predict the saccade behavior of the monkey and predict the temporal dynamics of movement-related neurons in FEF.

A Neural Locus of Drift Rates

Movement-related neurons increase in spike rate over time and reach a constant level of activity immediately prior to a saccade being initiated (Figure 15.4). The dynamics of movement-related neurons appear consistent with the dynamics of models that assume a stochastic accumulation of perceptual evidence to a threshold (Hanes & Schall, 1996; Ratcliff et al., 2003; Schall, 2001; Smith & Ratcliff, 2004). This insight raises several questions that we have begun to address in our recent work: If movement-related neurons instantiate an accumulator model, what kind of accumulator model do they instantiate? What kind of an accumulator model can predict the fine-grained dynamics of movement-related neurons? What drives the accumulator model? We begin with the last question.

A broad class of models of perceptual decision-making assumes that perceptual evidence is accumulated over time to a threshold (Figure 15.2; see also Ratcliff & Smith, this volume). The rate at which perceptual evidence is accumulated, the drift rate, can vary across objects, conditions, and experience. When accumulator models are tested by fitting them to observed behavior, it is not uncommon to assume that different drift rates across different experimental conditions are free parameters that are optimized to maximize or minimize some fit statistic (e.g., Brown & Heathcote, 2008; Boucher et al., 2007a; Ratcliff & Rouder, 1998; Usher & McClelland, 2001). But other theoretical work has aimed to connect models of perceptual decision-making to models of perceptual processing by developing a theory of the drift rates.

For example, Nosofsky and Palmeri (1997; Palmeri, 1997) proposed an exemplar-based random walk model (EBRW) that combined the generalized context model of categorization (Nosofsky, 1986) with the instance theory of automaticity (Logan, 1988) to develop a theory of the drift rates driving a stochastic accumulation of evidence. Briefly, EBRW assumes that a perceived object activates previously stored exemplars in visual memory, the probability and speed of exemplar retrieval is governed by similarity, and repeated exemplar retrievals determine the direction and rate of accumulation to a response threshold. EBRW predicts the effects of similarity, experience, and expertise on response probabilities and response times for perceptual decisions about visual categorization and recognition (see Nosofsky & Palmeri, 2015; Palmeri & Cottrell, 2009; Palmeri, Wong, & Gauthier, 2004). Other theorists have similarly connected visual perception and visual attention mechanisms to accumulator models of perceptual decision making by creating theories of drift rate (e.g., Ashby, 2000; Logan, 2002; Mack & Palmeri, 2010; Schneider & Logan, 2005; Smith & Ratcliff, 2009).

As a first step toward a neural theory of drift rates, we recently proposed a neural locus of drift rates when decisions are made by saccades (Purcell et al., 2010, 2012). We hypothesize that the accumulation of evidence is reflected in the firing rate of FEF movement-related neurons and the perceptual evidence driving this accumulation is reflected in the firing rate of FEF visually responsive

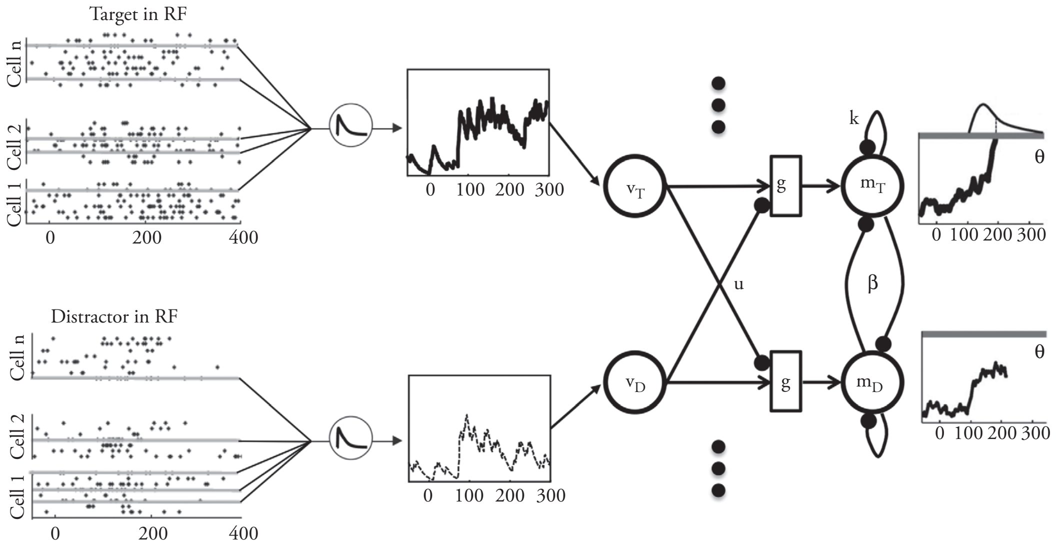

Fig. 15.5 Illustration of simulation model architectures tested in Purcell et al. (2010, 2012). Spike trains were recorded from FEF visually-responsive neurons during a saccade visual search task. Trials were sorted into two populations according to whether the target or a distractor was within the neuron's response field. Spike trains were randomly sampled from each population to generate a normalized activation function that served as the dynamic model input associated with a target (v_T) and a distractor (v_D) on a given simulated trial, as illustrated. Different architectures for perceptual decision-making were systematically tested. Decision units (m_T) could integrate evidence or not, and they could be leaky (*k*) or not. Decision units could integrate a difference between the inputs (*u*) or not, the stochastic input could be gated (*g*) or not, and the units could compete with one another (β) or not. Here, only two decision units are shown, one for a target and one for a distractor. In Purcell et al. (2012) there were eight accumulators, one for each possible stimulus location in the visual search array.

neurons. One way to test this hypothesis would be to develop a model of the dynamics of visually responsive neurons, a model of how those dynamics are translated into drift rates, and then use those drift rates to drive a model of the accumulation of perceptual evidence. We chose a different approach. Rather than model the dynamics of visually responsive neurons, we used the observed firing rates of those neurons directly as a dynamic neural representation of the perceptual evidence that was accumulated over time.

Figure 15.5 illustrates our general approach. Activity of visually responsive neurons was recorded from FEF of monkeys performing a visual search task. In Figure 15.4, we illustrate spike density functions of a representative neuron when a target or distractor appeared in its receptive field during easy or hard visual search. For our modeling, we did not use the mean activity of neurons as input but, instead, generated thousands of simulated spike-density functions by subsampling from the full set of individually recorded trials of visually responsive neurons. Specifically, on each simulated trial, we first randomly sampled, with replacement, a set of spike trains recorded from individual neurons. We subsampled from trials when the target was in the receptive fields of the neurons to simulate perceptual evidence in favor of the target location and trials when a distractor was in the receptive field to simulate perceptual evidence in favor of each of the distractor locations. Along its far left, Figure 15.5 illustrates raster plots for example neurons, with individual trials arranged sequentially along the *y* axis, time along the *x* axis, and each black dot indicating the incidence of a recorded spike on a given trial for that neuron. The gray thick bars illustrate a random sampling from those recorded neurons. These sampled spike trains were convolved with a temporally asymmetric doubly exponential function (Thompson, Hanes, Bichot, & Schall, 1996), averaged together, and normalized to create dynamic drift rates associated with target and distractor locations (Purcell et al., 2010, 2012), as illustrated in the middle of Figure 15.5; the resulting input functions are mathematically similar to a Poisson shot noise process (Smith, 2010). Different inputs were defined according to the experimental condition under which the visually responsive neurons were recorded on each trial, such as easy versus hard search or small versus large set sizes.

Arguably, this approach allows the most direct test of whether the dynamics of visually responsive neurons provide a sufficient representation of perceptual evidence to predict where and when the monkey moves its eyes. If no model can predict saccade behavior using visually responsive neurons as input, then some other neural signal must be significantly modulating behavior of the monkey. Furthermore, as illustrated by contrasting Figures 2a and 2b, this novel approach imposes significant constraints on possible models by replacing free parameters governing the mean and variability of perceptual processing time, starting point of accumulation, and drift with observed neurophysiology. Finally, because the neurophysiological signal from visually responsive neurons is continuous in time, the models cannot merely assume that perceptual processing and perceptual decisions constitute discrete stages, as typical for many accumulator models.

Architectures for Perceptual Decision-Making

Within the broad class of perceptual decision-making models assuming an accumulation of perceptual evidence to a threshold, a variety of different model architectures have been proposed (e.g., see Ratcliff & Smith, 2004; Smith & Ratcliff, 2004). We instantiated several of these competing architectures, and using drift rates defined by the recorded spiking activity of visually responsive neurons as inputs, evaluated how well each could fit observed response probabilities and response times of monkeys making saccades during a visual search task (Purcell et al., 2010, 2012).

Figure 15.5 illustrates the common architectural framework. Drift rates defined by neurophysiology constitute the input nodes labeled v_T (target) and v_D (distractor). We assume an accumulator associated with the target location (m_T) and distractor locations (m_D). Figure 15.5 shows only one target and one distractor accumulator (Purcell et al., 2010) but we have extended this framework to multiple accumulators, one for every possible target location in the visual field (Purcell et al., 2012). Each accumulator is governed by the following stochastic differential equation

$$dm_i(t) = \frac{dt}{\tau}\left[\left(v_i(t) - \sum_{j\neq i} uv_{j(t)} - g\right)^{+} - \sum_{k\neq i}\beta m_k(t) - km_i(t)\right] + \sqrt{\frac{dt}{\tau}}\xi.$$

The $m_i(t)$ are rectified to be greater than or equal to zero because we later compare the dynamics of these accumulators to the observed spike rates of movement-related neurons, and those spike rates are greater than zero by definition. ξ represents Gaussian noise intrinsic to each accumulator with mean 0 and standard deviation σ; in all of our simulations, this intrinsic accumulator variability could be assumed to be quite small relative to the variability of the visual inputs $v_i(t)$. All accumulators, $m_i(t)$, are assumed to race against one another to be the first to reach their threshold θ. The winner of that race between accumulators determines which saccade response is made on that simulated trial and the response time is given by the time to reach threshold plus a small ballistic time of 10–20ms.

If $k > 0$, these are leaky accumulators, otherwise they are perfect integrators. If $\beta = 0$ and $u = 0$, we have a version of a simple horse race model. If $\beta > 0$, these are competing accumulators, and combined with leakage, $k > 0$, we have the leaky competing accumulator model (Usher & McClelland, 2001). If $u > 0$, then weighted differences are accumulated by each $m_i(t)$. In the case of only two accumulators, one for a target and the other for a distractor, and assuming $u = 1$, both $m_i(t)$ accumulates the difference between evidence for a target versus evidence for a distractor, which is quite similar to a standard drift-diffusion model (see Bogacz et al., 2006; Ratcliff et al., 2007; Usher & McClelland, 2001), and when assuming positive leakage ($k > 0$) is quite similar to an Ornstein-Uhlenbeck process (Smith, 2010); this similarity can become mathematical identity with some added assumptions (Bogacz et al., 2006; Usher & McClelland, 2001).

Finally, we also proposed a novel aspect to this general architecture, which we called a *gated accumulator* (Purcell et al., 2010, 2012). When $g > 0$ and the input is positive-rectified, as indicated by the $^{+}$ subscript in the equation, then only inputs that are sufficiently large can enter into the accumulation. For example, consider a gated accumulator assuming $u > 0$; this would mean that the differences in the evidence in favor of the target over the distractors must be sufficiently large before that differences will accumulate. Recall that we assumed that the inputs are defined by neurophysiology, which has no beginning or ending, apart from the birth or death of the organism. Intuitively, the gate forces the accumulators to accumulate signal, not merely noise, and noise is all that is present before

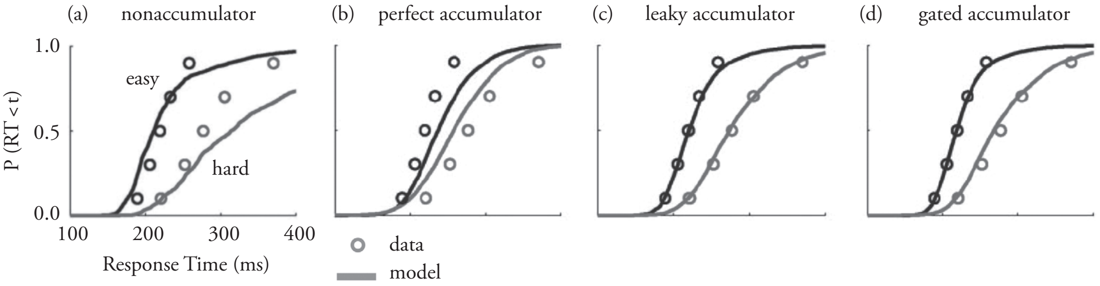

Fig. 15.6 In Purcell et al. (2010), models (Figure 15.5) were tested on how well they could account for observed RT distributions of the onset of saccades in an easy visual search where the target and distractors were dissimilar or where the target and distractors were similar hard. Each panel shows observed cumulative RT distributions (symbols) for easy and hard search. Best-fitting model predictions for a subset of the models tested in Purcell et al. (2010) are shown for illustration, ranging left-to-right from a nonaccumulator model that does not integrate perceptual evidence over time, a perfect integrator model with no leakage, a leaky accumulator model, and a gated accumulator model. (Adapted from Purcell et al., 2010.)

perceptual processing has begun to discriminate targets from distractors.

We evaluated the fits of competing model architectures to observed response probabilities and distributions of response times using standard model fitting techniques (e.g., Ratcliff & Tuerlinckx, 2002; Van Zandt, 2000). We systematically compared models assuming a horse race, a diffusion-like difference accumulation process, or competition via lateral inhibition, factorially combined with various leaky, nonleaky, or gated accumulators. For example, Figure 15.6 displays observed response time distributions for easy versus hard visual search along with a sample of predictions from some of the model architectures evaluated by Purcell et al. (2010); for these particular data (Bichot, Thompson, Rao, & Schall, 2001; Cohen et al., 2009), there were very few errors. As shown in the left two panels, models assuming nointegrationat all, meaning that the current value of $m_i(t)$ simply reflects the current inputs at time t, and models assuming perfect integration without leakage, provided a relatively poor fit to the observed behavioral data. Although these particular behavioral data were fairly limited, with only a response-time distribution for easy and hard visual search, we could rule out some potential model architectures. However, other competing models, including those with leakage or gate, assuming a competition or an accumulation of differences, all provided reasonable quantitative accounts of the behavioral data, a couple of examples of which are shown in the two right panels of Figure 15.6.

Purcell et al. (2012) evaluated fits of these models to a more comprehensive dataset where set size was systematically manipulated and where the search was difficult enough to produce significant errors (Cohen et al., 2009). Models were required to fit correct- and error-response probabilities as well as distributions of correct- and error-response times. These data are shown in Figure 15.7. Also shown are the predictions of the best fitting model, which was a gated accumulator model that assumed both significant leakage and competition via lateral inhibition. Likely because this dataset was larger, it also provided a greater challenge to other models, since many horse-race models and diffusion-like models failed to provide adequate fits to the observed data, whether they included leakage or gating (see Purcell et al., 2012).

Just based on the quality of fits to observed data, models with leakage and competition via lateral inhibition provided comparable fits whether those models included gating or not in both Purcell et al. (2010) and Purcell et al. (2012). So based on parsimony, a nongated version, which is essentially a leaky competing accumulator model (Usher & McClelland, 2001), would win the theoretical competition. But our goal was also to test whether the accumulators in the competing models could provide a theoretical account of the movement-related neurons in FEF. To do that, we also tested whether the dynamics measured in the accumulators could predict the dynamics measured in movement-related neurons (see also Boucher et al., 2007a; Ratcliff et al., 2003, 2007).

Predicting Neural Dynamics

Until now, the work we have described follows a long tradition of developing and testing computational and mathematical models of cognition.

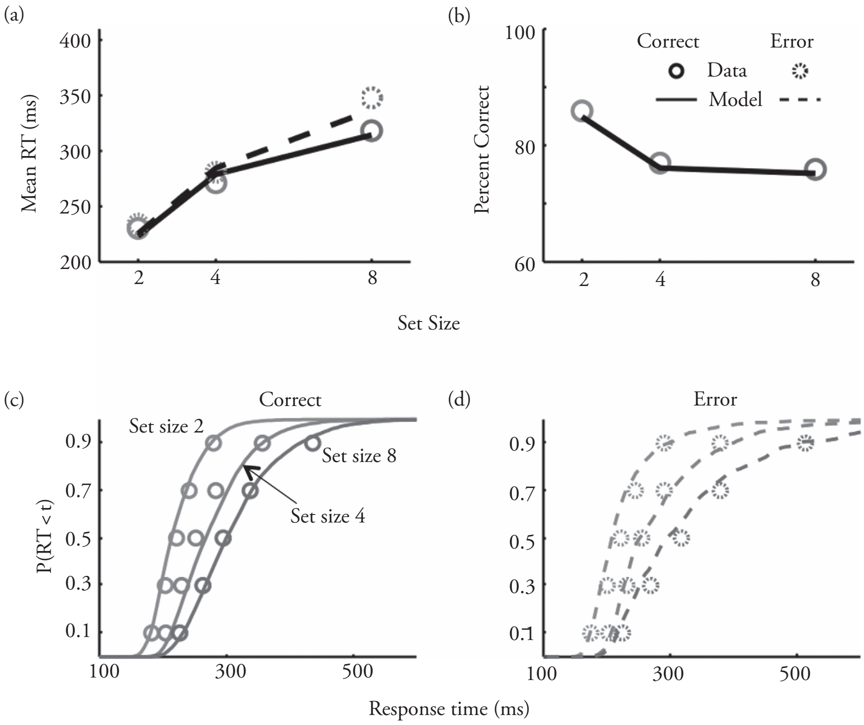

Fig. 15.7 In Purcell et al. (2012), models (Figure 15.5) were tested on how well they could account for correct- and error-response probabilities and correct- and error-response time distributions of saccades in a visual search task with three levels of set size: 2 (blue), 4 (green), or 8 (red) objects in the visual array. Predictions from the best-fitting gated accumulator model are shown. (a) Mean observed (symbols) and predicted (lines) correct- (solid) and error- (dashed) response times as a function of set size. (b) Mean observed (symbols) and predicted (lines) probability correct as a function of set size. (c) Observed (symbols) and predicted (lines) cumulative RT distributions of correct responses at each set size. (d) Observed (symbols) and predicted (lines) cumulative RT distributions of error responses at each set size. (Adapted from Purcell et al., 2012.)

Competing models are evaluated on their ability to predict behavioral data by optimizing parameters in order to maximize or minimize the fit of each model to the observed data, and then statistical tests are performed for nested or nonnested model comparison (e.g., see Busemeyer & Diederich, 2010; Lewandowsky & Farrell, 2010). We go beyond this approach to evaluate linking propositions (Schall, 2004; Teller, 1984) that aim to map particular cognitive model mechanisms onto observable neural dynamics. Specifically, we evaluate the linking proposition that movement-related neurons in FEF instantiate an accumulation of evidence to a threshold. We do this by testing how well the simulated dynamics of accumulators in the various model architectures described in the previous section predict the observed dynamics in movement-related neurons. Although the qualitative relationship between accumulator dynamics and movement neuron dynamics has long been recognized (e.g., Hanes & Schall, 1996; Ratcliff et al., 2003; Smith & Ratcliff, 2004), we go beyond noting qualitative relationships to test quantitative predictions.

Following the approach used by Woodman et al. (2008), we evaluated how several key measures of neural dynamics varied according to the measured response time of a saccade. The top row of Figure 15.8 illustrates several hypotheses for how variability in response time is related to variability in the underlying neural dynamics. Fast responses could be associated with an early initial onset of the neural activity from baseline, whereas slow responses could be associated with a delayed onset. Alternatively, fast responses could be associated with high growth rate in spiking activity to threshold, whereas slow responses could be associated with low growth rate. Fast responses could be associated with an increased baseline firing rate or decreased threshold, whereas slow responses could be associated with a decreased baseline firing rate or increased threshold. To evaluate these proposals, the onset time, growth rate, baseline, and threshold of neural activity were all measured within bins of trials defined by response times from fastest to slowest, both within conditions and across conditions (see Purcell et al., 2010, 2012, for details). The middle row shows the relationship between onset time, growth rate, baseline, and threshold of neural activity and mean response time for each bin of an RT distribution for a representative neuron in a representative condition.The bottom row shows the mean correlation of neural measures with RT as a function of set size from Purcell et al. (2012), with a significant relationship between onset time and response time observed in neural activity in movement-related neurons in FEF.

Using analogous methods, we also measured the relationship between onset time, growth rate, baseline, and threshold of accumulator dynamics and response time predicted by each of the competing model architectures that we simulated. Shown in Figure 15.8 are the predictions of the gated accumulator model from Purcell et al. (2012), illustrating a good match between model and neurons. These are true model predictions, not model fits. After the model was fitted to behavioral data, the accumulator dynamics using the best-fitting model parameters were measured and compared directly with the observed neural dynamics. All other models failed to predict the observed neural dynamics. For example, models without gate typically predicted a significant negative correlation between baseline and response time that was completely absent in the observed data. Part of the reason for this is that, with nongated models, the accumulators are allowed to accumulate noise in the input defined by visually responsive neurons. Although a leakage term may be sufficient to keep a weak noise signal from leading to a premature accumulation to threshold, it cannot prevent significant differences in baseline activity from being correlated with differences in predicted response time when the accumulators reach threshold, at least without significantly compromising fits to the observed behavior.

Control over Perceptual Decisions

We have also considered the neurophysiological basis of cognitive control over perceptual decisions. Mirroring our other research, we used cognitive models to better understand neural mechanisms and used neural data to constrain competing cognitive models.

Perhaps the most widely used task for studying normal and dysfunctional cognitive control is the stop-signal task (Lappin & Eriksen, 1966; Logan & Cowan, 1984). Saccade variants of this task have been used with monkeys, and neurophysiological activity has been recorded from neurons in FEF (Hanes et al., 1998). The basic stop-signal task with saccades is in certain ways a converse of the memory-guided saccade task illustrated in Figure 15.4. Monkeys initially fixate the center of the screen. After a variable amount of time, the fixation spot disappears and a peripheral target appears somewhere in the visual field, and the monkey must make a single saccade to the target in order to earn reward. This is the primary task, or go signal. On a fraction of trials, some time after the peripheral target appears, the fixation spot is reilluminated, and the monkey is rewarded for cancelling its saccade, maintaining fixation. This is the stop signal. The interval between the appearance of the go signal, the peripheral target, and the stop signal, the fixation point, is called stop signal delay (SSD). Monkeys' ability to inhibit their saccade is probabilistic due to the stochastic variability of go and stop processes and depends on SSD.

Figure 15.9 displays the key behavioral data observed in the saccade stop-signal paradigm (Hanes et al., 1998). Figure 15.9a displays the probability of responding to the go signal (*y* axis), despite the presence of a stop signal at a particular SSD (*x* axis). When the stop signal illuminates shortly after the appearance of the target, at a short SSD, the probability of responding to the go signal is quite small. Control over the saccade as a consequence of the stop signal has been successful. In contrast, for a long SSD, the probability of successfully inhibiting the saccade is rather small. Figure 15.9b displays distributions of response times for primary go trials with a stop signal (signal response trials), in which a saccade was erroneously made, shaded by gray according to SSD (see figure caption). These response times are significantly faster than response times without any stop signal (no-stop-signal trials) in black.

Behavioral data in the stop-signal paradigm has long been accounted for by an independent race model (Logan & Cowan, 1984), which assumes that performance is the outcome of a race between a go process, responsible for initiating the movement, and a stop process, responsible for inhibiting the movement (see also Becker & Jürgens, 1979; Boucher, Stuphorn, Logan, Schall, & Palmeri, 2007b; Camalier et al., 2007;

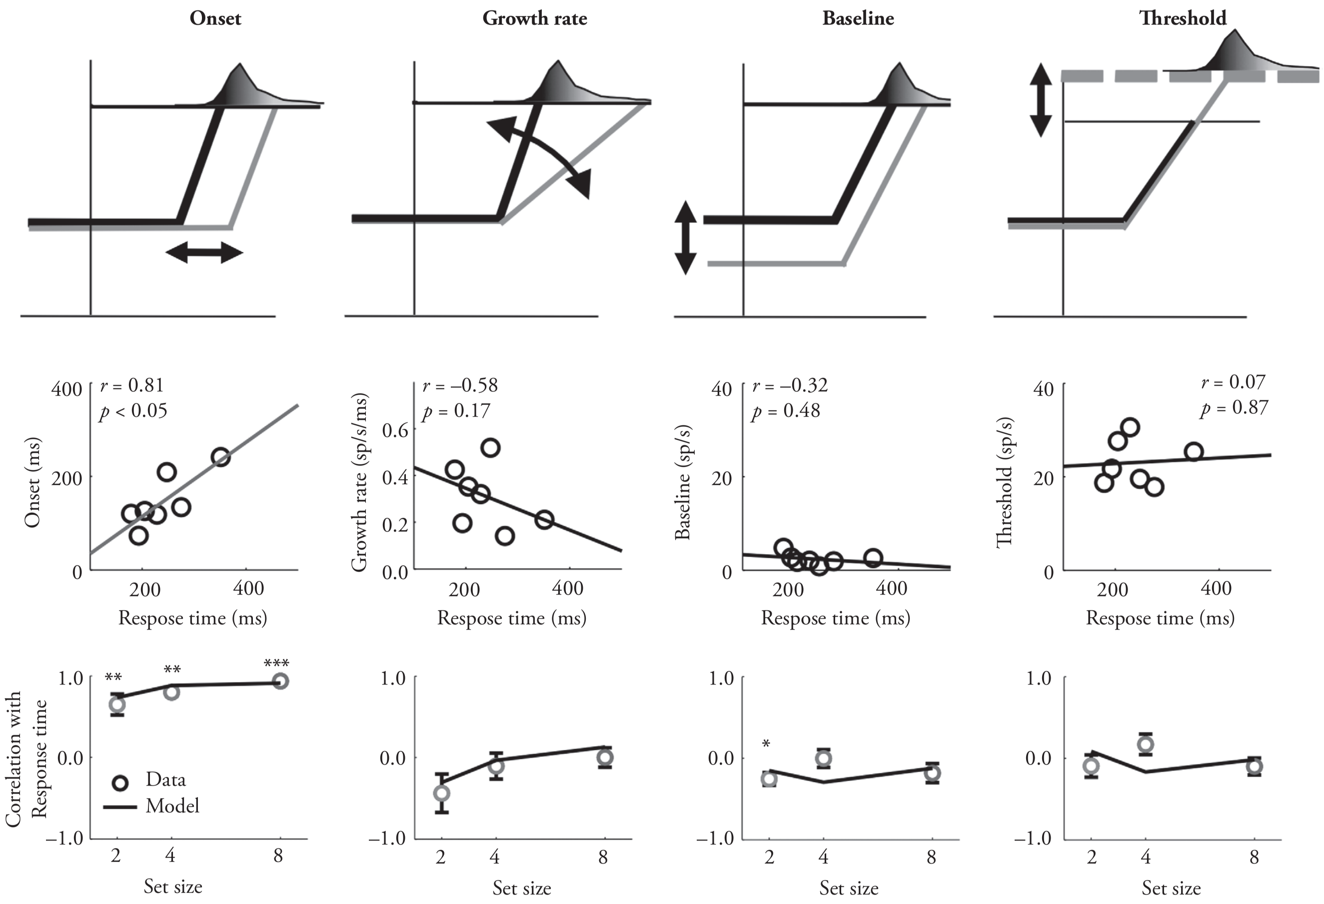

Fig. 15.8 Comparing observed neural dynamics and predicted model dynamics. Top row: Four possible hypotheses for how variability in RT is related to variability in neural or accumulator dynamics: from left to right, variability in RT could be correlated with variability in the onset time, growth rate, baseline, or threshold. Middle row: Following Woodman et al. (2008), correct RTs were binned in groups from fastest to slowest and within each bin the onset time, growth rate, baseline, and threshold of the spike density functions were calculated. The relationship between RT and neural measure (left to right: onset time, growth rate, baseline, and threshold) are shown for one representative neuron in set size 4 for one of the monkeys tested; the correlation between RT and neural measure and its associate p-value are also shown. Bottom row: Average correlation between RT and neural measure (left-to-right: onset time, growth rate, baseline, and threshold) as a function of set size observed in neural dynamics and predicted in model dynamics for the gated accumulator model. (Adapted from Purcell et al., 2012.)

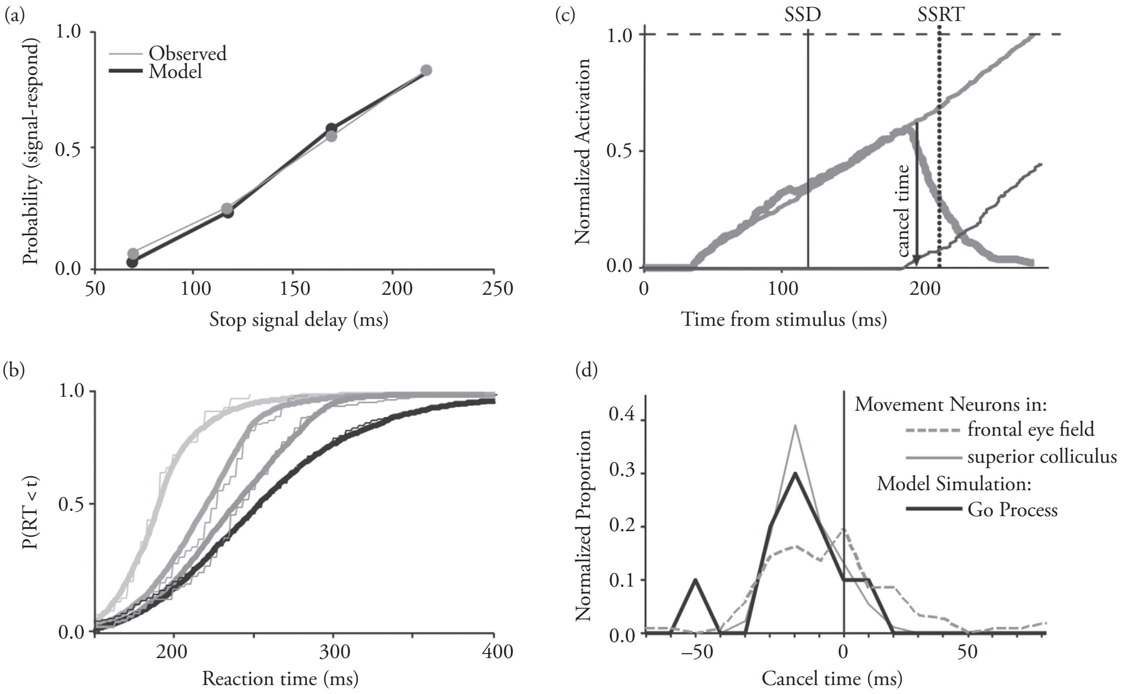

Fig. 15.9 (a) Observed inhibition function(gray line) and simulated inhibition function from the interactive race model (black line). (b) Observed (thin lines) and simulated (thick lines) cumulative RT distributions from no stop signal (black line) and signal-response trials with progressively longer stop signal delays (progressively darker gray lines). (c) Illustration of simulated activity in the interactive race model of the go unit and stop unit activation on signal-inhibit (thick solid line) and latency-matched no-stop-signal trials (thin solid lines) with stop-signal delay (SSD) and stop-signal reaction time (SSRT) indicated. Cancel time is indicated by the downward arrow. (d) Histogram of cancel times of the go unit predicted by the interactive race model compared with the histogram of cancel times measured for movement-related neurons in FEF and SC.(Adapted from Boucher et al., 2007a.)

Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014; Olman, 1973). Boucher et al. (2007a) addressed an apparent paradox of how seemingly interacting neurons in the brain could produce behavior that appears to be the outcome of independent processes. Mirroring the general model architectures described earlier and illustrated in the right half of Figure 15.5, they instantiated and tested models that assumed stochastic accumulators for the go process and for the stop process that were either an independent race or that assumed competitive, lateral interactions between stop and go. Outstanding fits to observed behavioral data for both the independent race model and the interactive race model were observed. Figures 9a and 9b show fits of the interactive race model, but fits of the independent race model were virtually identical. Parsimony would favor the independent race. But neural data favored the interactive race.

In the absence of a stop signal, visually responsive neurons in FEF select the target, and movement-related neurons in FEF increase their activity until a threshold level is reached, shortly after which a saccade is made (Hanes & Schall, 1996), just as they do on memory-guided saccade tasks or visual search tasks. On trials with a stop signal, the dynamics of visually responsive neurons are unaffected (Hanes et al., 1998). For movement-related neurons, we can distinguish between activity when a stop was successful, signal-inhibit trials, from activity when a stop was unsuccessful, that is signal-respond trials. On signal-respond trials, the activity of movement-related neurons is qualitatively the same as the activity on no-signal trials, with neurons reaching a threshold level before a saccade is made. Even more striking, the activity on signal-respond trials is quantitatively indistinguishable from activity on no-signal trials that are equated for response time (latency-matched trials). On signal-inhibit trials, the activity increases in a manner indistinguishable from latency-matched no-signal trials until some time after the SSD, at which point the activity of movement-related neurons is reduced back to baseline without reaching the threshold. The saccade has been inhibited.

Figure 15.9c displays the predicted accumulator dynamics of the interactive race model (Boucher et al., 2007a). The dynamics of the go accumulator in the interactive race precisely mirrors the description of the dynamics of movement-related neurons provided earlier, with dynamics not observed in the independent race model. For signal-inhibit trials and latency-matched no-signal trials, activity increases for some time after SSD, after which activity on signal-inhibit trials returns to baseline while activity on latency-matched no-signal trials continues to threshold. The accumulator dynamics in the interactive race model qualitatively captures the neural dynamics of movement-related neurons. But we could go further than that. We also calculated a metric called *cancel time* (Hanes et al., 1998), which is a function of the time at which the dynamics statistically diverge between signal-inhibit trials and latency-matched no-signal trials. This time can be calculated from movement-related neurons. It can also be calculated from accumulator dynamics. And as shown in Figure 15.9b, these measures from neurons and the model nicely converge. We emphasize that, as was the case for Purcell et al. (2010, 2012), these are true model predictions. Boucher et al. (2007a) fitted models to behavioral data, then calculated the cancel time predicted by the models, and compared that to the observed cancel time in neurons. Parameters were not adjusted to maximize the correspondence.

The hypothesized locus of control in Boucher et al. (2007a) is inhibition of a stop process on the go process, with the stop process identified as activity of fixation-related neurons and the go process identified as activity of movement-related neurons. The gate in the gated accumulator model (Purcell et al., 2010, 2012) could be another hypothesized locus of control over perceptual decisions. In recent work, we have suggested that blocking the input to the go unit, rather than actively inhibiting it via a stop unit, could be an alternative mechanism for stopping. Indeed, a blocked input model predicted observed data and distributions of cancel times at least as well as the interactive race model (Logan, Schall, & Palmeri, 2015; Logan, Yamaguchi, Schall, & Palmeri, 2013). One suggestion we made was that the stop process could raise a gate between visual neurons that select the target and movement neurons that generate a movement to it, blocking input to the movement neurons and thereby preventing them from reaching threshold.

As another example, in a stop-signal task, both humans and monkeys adapt their performance from trial to trial, for example, producing longer RTs after successfully inhibiting a planned movement (e.g., Bissett & Logan, 2011; Nelson, Boucher, Logan, Palmeri, & Schall, 2010; Verbruggen & Logan, 2008). For monkeys, within FEF, activity of visually responsive neurons are unaffected by these trial-to-trial adjustments, but the onset time of activity of movement-related neurons is significantly delayed (Pouget et al., 2011). Purcell et al. (2012) suggested that strategic adjustment in the level of the gate could explain the delayed onset of movement-related neurons in the absence of any modulation of visually responsive neurons. Moreover, they demonstrated that this strategic adjustment of gate could be couched in terms of optimality. It has been previously suggested that strategic modulation of accumulator threshold could maximize reward rate, which is defined as the proportion of correct responses per unit time (e.g., Gold & Shadlen, 2002; Lo & Wang, 2006). We observed that strategic modulation of the level of the gate could maximize reward rate in much the same way (Purcell et al., 2012).

Summary and Conclusions

Here we reviewed some of our contributions to a growing synergy of mathematical psychology and systems neuroscience. Our starting point has been a class of successful cognitive models of perceptual decision-making that assume a stochastic accumulation of perceptual evidence to a threshold over time (Figure 15.2). Models of this sort have long provided excellent accounts of response probabilities and distributions of response times in a wide range of perceptual decision-making tasks and manipulations (e.g., see Nosofsky & Palmeri, 2015; Ratcliff & Smith, 2015). We have extended these models to account for response probabilities and distributions of response times for awake behaving monkeys to make saccades to target objects in their visual field (Boucher et al., 2007a; Pouget et al., 2011; Purcell et al., 2010, 2012). Applying techniques common to mathematical psychology, we instantiated different model architectures and ruled out models that provided poor fits to observed data.

These models have free parameters that govern theoretical quantities like perceptual processing time, the starting point of accumulation, the drift rate of accumulation, and the response threshold. We constrained many of these parameters using neurophysiology. Unlike some approaches that constrain parameters *values* based on neurophysiology, often based on neural findings with rather large confidence intervals, we replaced parameterized model assumptions directly with recorded neurophysiology. Specifically, we sampled from neural activity recorded from visually responsive neurons in FEF, feeding these spike trains directly into stochastic accumulator models, thereby creating a largely nonparametric neural theory of perceptual processing time and the drift rate of accumulation. Not only did this approach constrain computational modeling, it also provided a direct test of the hypothesis that the activity of visually responsive neurons in FEF encodes perceptual evidence: This neural code can be accumulated over time to predict where and when the monkey moves its eyes (Purcell et al., 2010, 2012).

We also tested the hypothesis that movement-related neurons in FEF instantiate a stochastic accumulation of evidence. Although it has long been acknowledged that these neurons behave in a way consistent with accumulator models (e.g., Hanes & Schall, 1996; Schall, 2001), we went beyond qualitative description to test whether movement neuron dynamics can be quantitatively predicted by accumulator model dynamics. We measured how the onset of activity, baseline activity, rate of growth, and threshold varies with behavioral response time in both movement-related neurons and model accumulators, and we found close correspondences for some models.

Not only does this test an hypothesis about the theoretical role of FEF movement-related neurons in perceptual decision-making, it also provides a powerful means of contrasting models that otherwise make indistinguishable behavioral predictions. Our gated accumulator model, which enforces accumulation of discriminative neural signals from visually responsive neurons, not only accounted for the detailed saccade behavior of monkeys, but also predicted quantitatively the dynamics observed in movement-related neurons in FEF, whereas other models could not (Purcell et al., 2010, 2012; see also Boucher et al., 2007a). This gated accumulator model also suggests a potential locus of cognitive control over perceptual decisions. Increasing the gate may account for speed-accuracy tradeoffs (Purcell et al., 2012) as well as stopping behavior and trial history effects described by Boucher et al. (2007a) and Pouget et al. (2011), respectively.

Turning to more general issues, our work has confronted a common challenge in the development of mathematical and computational models of cognition where competing models reach a point where they make very similar predictions, examples of which are discussed in other chapters in this volume (Busemeyer, Townsend, Wang, & Eidels, 2015). This could be a consequence of true mimicry, where models assuming vastly different mechanisms nonetheless produce mathematically identical predictions that cannot be distinguished behaviorally. Often, however, it is that the current corpus of experimental manipulations and measures are insufficient to discriminate between competing models. Cognitive modelers have long turned to predicting additional complexity in behavioral data to resolve mimicry, going from predicting accuracy alone to predicting response probabilities as well as response times, and from predicting mean response-times to predicting response, time distributions, including those for correct and error responses. Indeed, in our work reviewed here, predicting jointly response probabilities and response time distributions yielded considerable traction in discriminating between competing models. Unfortunately, outside the mathematical psychology community, it is not uncommon to hear researchers state with complete confidence that response time distributions yield no more useful information than response time means, sadly unknowledgeable about the state of reality (e.g., see Townsend, 1990). That said, recognition is emerging, for example, that response time distributions are key aspects of data that theories of visual cognition needs to account for (e.g., Palmer, Horowitz, Torralba, & Wolfe, 2011; Wolfe, Palmer, & Horowitz, 2010), that response time distributions provide challenging constraints for low-level spiking neural models (e.g., Lo, Boucher, Paré, Schall, & Wang, 2009), and more generally that considerations of behavioral variability can yield insights into neural processes (e.g., Churchland et al., 2011; Purcell, Heitz, Cohen, & Schall, 2012). But even joint modeling of response probabilities and response-time distributions may be insufficient to contrast competing models.

Our work illustrates how neurophysiological data can also help distinguish between models. We have described cases in which two models fit behavioral data equally well (Boucher et al., 2007a;

Purcell et al., 2010, 2012) but one model is more complex than the other. With only behavioral data and an appeal to parsimony, we would have demanded the exclusion of the more complex model in favor of the simpler one. However, in order to successfully mapobserved neural dynamics onto predicted model dynamics, the assumptions of the more complex model were required. Key here is that we believe that it is the important to map between neural dynamics and model dynamics, not between neural dynamics and model parameters (see also e.g., Davis, Love, & Preston, 2012). Variation in model parameters need not uniquely map onto variation in neural dynamics, but predicted variation in model dynamics must. And while we have demonstrated the theoretical usefulness of neural data in adjudicating between competing models, we do not believe that neural data has any particular empirical primacy. Just as mimicry issues can emerge when examining behavioral measures like accuracy and response time, analogous mimicry issues may be found at the level of neurophysiology and neural dynamics. Neural data are not necessarily more intrinsically informative than behavioral data, but more data provides additional constraints for distinguishing between competing models.

More generally, our work allies with a growing body of research supporting accumulator models of perceptual decision making (e.g., Nosofsky & Palmeri, 1997; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Usher & McClelland, 2001), not just as models that explain behavior but also as models that explain brain activity measured using neurophysiology (e.g., Boucher et al., 2007; Churchland & Ditterich, 2012; Purcell et al., 2010, 2012; Ratcliff et al., 2003; but see Heitz & Schall, 2012, 2013), EEG (e.g., Philiastides, Ratcliff, & Sajda, 2006), and fMRI (e.g., Turner et al., 2013; van Maanen et al., 2011; White, Mumford, & Poldrack, 2012). The relative simplicity of cognitive models like accumulator models is a virtue in that they are computationally tractable, making them easily applicable across a wide range of phenomena and levels of analysis.

Making explicit links to brain mechanisms does expose complexities. Our focus here has been largely on FEF, but other brain areas have neurons with dynamics that are visually responsive or movement-related, including SC (Hanes & Wurtz, 2001; Paré & Hanes, 2003) and LIP (Gold & Shadlen, 2007; Mazurek et al., 2003; Shadlen & Newsome, 2001). Compared to the relative simplicity of most

Box 1 Top-down versus Bottom-up Theoretical Approaches

Computational cognitive neuroscience aims to understand the relationship between brain and behavior using computational and mathematical models of cognition. One approach is bottom up. Theorists begin with fairly detailed mathematical models of neurons based on current understanding of cellular and molecular neurobiology. A common approach is to develop and test a single model of a neural network built up from these detailed models of neurons along with hypotheses about their excitatory and inhibitory connectivity. Although these neural models provide excellent accounts of spiking and receptor dynamics of individual neurons and may also account well for emergent network activity, they may provide only fairly coarse accounts of observed behavior, have somewhat limited generalizability, and be impractical to rigorously simulate and evaluate quantitatively.

Another approach is top down (e.g., Forstman et al., 2011; Palmeri, 2014). Cognitive models account for details of behavior across multiple conditions, have significant generalizability across tasks and subject populations, and are often relatively easy to simulate and evaluate. It is common to evaluate multiple competing models and to test the necessity and sufficiency of model assumption with nested model comparison techniques. Although these models do not provide the same level of detailed predictions of spiking and receptor dynamics, they can provide predictions about the temporal dynamics of neural activity at the same level of precision as commonly summarized in neurophysiological investigations, as we illustrated in our review. In fact, Carandini (2012) suggested that bridging between brain and behavior can only be done by considering intermediate-level theories, that the gap between low-level neural models and behavior is simply a "bridge too far." Although he considered linear filtering and divisive normalization as example computations that may be carried out across cortex (Carandini & Heeger, 2011), we consider accumulation of evidence as a similar computation that may be carried out in various brain areas, included FEF. These computations can simultaneously explain behavioral and neural dynamics.

stochastic accumulator models, there isa network of brain areas involved in evidence accumulations for perceptual decision making (Gold & Shadlen, 2007; Heekeren et al., 2008; Schall, 2001; 2004). Such mechanisms involving accumulation of evidence for perceptual decision-making may be replicated across different sensory and effector systems in the brain, such as those for visually guided saccades, but there may also be domain-general mechanisms as well (e.g., Ho, Brown, & Serences, 2009). Although the dynamics of specific individual neurons within particular brain areas mirror the dynamics of accumulators in models, we also know that, within any given brain area, ensembles of tens of thousands of neurons are involved in the generation of any perceptual decision. We need to understand the scaling relations from simple accumulator models to complex ensembles of thousands of neural accumulators (Zandbelt, Purcell, Palmeri, Logan, & Schall, 2014) and how to map the relatively few parameters that define simple accumulator models onto the great number of parameters that define complex neural dynamics (Umakantha, Purcell, & Palmeri, 2014).

Acknowledgments

This work was supported by NIH R01-EY021833, NSF Temporal Dynamics of Learning Center SMA-1041755, NIH R01-MH55806, NIH R01-EY008890, NIH P30-EY08126, NIH P30-HD015052, and by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience. Address correspondences to Thomas J. Palmeri, Department of Psychology, Vanderbilt University, Nashville TN 37203. Electronic mail may be addressed to thomas.j.palmeri@vanderbilt.edu.

Glossary

drift rate: The mean rate of perceptual evidence accumulation in a stochastic accumulator model of perceptual decision-making.

frontal eye field: An area of prefrontal cortex that governs whether, where, and when the eyes moves to a new location in the visual field.

gated accumulator: A stochastic accumulator model that includes a gate that enforces accumulation of discriminative neural signals, a model which quantitatively accounts for both behavioral and neural dynamics of saccadic eye movement.

leakage: A weighted self-inhibition on the accumulation of perceptual evidence, turning a perfect integrator of perceptual evidence into a leaky integrator of perceptual evidence.

movement-related neurons: Neurons in FEF that show little or no modulation to the appearance of the target in the visual field but pronounced growth of spike rate immediately preceding the production of a saccade.

perceptual decision-making: Perceptual decision-making requires representing the world with respect to current task goals and using perceptual evidence to inform the selection of a particular action.

saccade: A ballistic eye movement of some angle and velocity to a particular location in the visual field.

stochastic accumulator model: A class of computational models that assume that noisy perceptual evidence is accumulated over time from a starting point to a threshold, allowing predictions of both response probabilities and distributions of response times.

stop-signal task: A classic cognitive control paradigm in which a primary go task is occasionally interrupted with a stop signal.

visually responsive neurons: Visually responsive neurons are neurons in FEF that respond to the appearance of an object in their receptive field relative to that object's salience with respect to current task goals but show little or no change in activity prior to the onset of a saccade

References

Ashby, F. G. (2000). A stochastic version of general recognition theory. *Journal of Mathematical Psychology, 44,* 310–329.

Becker, W., & Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vision Research, 19,* 976–983.

Bichot, N. P., & Schall. J. D. (1999). Effects of similarity and history on neural mechanisms of visual selection. *Nature Neuroscience, 2,* 549–554.

Bichot, N. P., Schall, J. D., & Thompson, K. G. (1996). Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature, 381,* 697–699.

Bichot, N. P., Thompson, K. G., Rao, S. C., & Schall, J. D. (2001). Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *Journal of Neuroscience, 21,* 713–725.

Bissett, P. G., & Logan, G. D. (2011). Balancing cognitive demands: Control adjustments in the stop-signal paradigm. *Journal of Experimental Psychology: Learning, Memory and Cognition, 37,* 392–404.

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

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007a). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review, 114,* 376–397.

Boucher, L., Stuphorn, V., Logan, G. D., Schall, J. D., & Palmeri, T. J. (2007b). Stopping eye and hand movements: Are the processes independent? *Perception & Psychophysics, 69,* 785–801.

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

Bruce, C. J., & Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *Journal of Neurophysiology, 53,* 603–635.

Bruce, C. J., Goldberg, M. E., Bushnell, M. C., & Stanton, G. B. (1985). Primate frontal eye fields: II. Physiological and anatomical correlates of electrically evoked eye movements. *Journal of Neurophysiology, 54,* 714–734.

Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Thousand Oaks, CA: Sage Publications.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review,* 100(3), 432–459.

Busemeyer, J. R., Townsend, J. T., Wang, Z. J. & Eidels A. (2015). *Mathematical and computational models of cognition*. Oxford, UK: Oxford University Press.

Camalier, C. R., Gotler, A., Murthy, A., Thompson, K. G., Logan, G. D., Palmeri, T. J., & Schall, J. D. (2007). Dynamics of saccade target selection: Race model analyses of double step and search step saccade production in human and macaque. *Vision Research, 47,* 2187–2211.

Carandini, M. (2012). From circuits to behavior: A bridge too far? *Nature Neuroscience, 15,* 507–509.

Carandini, A. K., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience, 13,* 51–62.

Churchland, A. K., & Ditterich, J. (2012). New advances in understanding decisions among multiple alternatives. *Current Opinion in Neurobiology,* 22(6), 920–926.

Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X. J., Pouget, A., & Shadlen, M. N. (2011). Variance as a signature of neural computations during decision making. *Neuron,* 69(4), 818–831.

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience,* 29(3), 11560–11571.

Cohen, J. Y., Heitz, R. P., Woodman, G. F., Schall, J. D. (2009). Neural basis of the set-size effect in frontal eye field: Timing of attention during visual search. *Journal of Neurophysiology, 101,* 1699–1704.

Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience, 8*(4), 429–453.

Davis, T., Love, B. C., & Preston, A. R. (2012). Learning the exception to the rule: Model-based fMRI reveals specialized representations for surprising category members. *Cerebral Cortex, 22,* 260–273.

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

Ditterich, J. (2010). A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Decision Neuroscience, 4*.

Ferrier, D. (1874). The localization of function in brain. *Proceedings of the Royal Society of London, 22,* 229–232.

Forstmann, B. U., Wagenmakers, E. J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and formal cognitive models: opposites attract? *Trends in Cognitive Sciences, 15*(6), 272–279.

Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The consilience of brain and decision. *Science, 306*(5695), 447–452.

Gilchrist, I. D. (2011). Saccades. In S. P. Liversedge, I. P. Gilchrist, & S. Everling (Eds.), *Oxford Handbook on Eye Movements*, (pp. 85–94). Oxford, UK: Oxford University Press.

Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron, 36,* 299–308.

Gold, J. I., & Shadlen, M.N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–560.

Goldman-Rakic, P. S., & Porrino, L. J. (1985). The primate mediodorsal (MD) nucleus and its projection to the frontal lobe. *Journal of Comparative Neurology, 242,* 535–560.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding: II. Feedback, expectation, olfaction, illusions. *Biological Cybernetics, 23,* 187–202.

Hanes, D. P., Patterson, W. F., II, & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology, 79,* 817–834.

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science, 274,* 427–430.

Hanes, D. P., & Wurtz, R. H. (2001). Interaction of the frontal eye field and superior colliculus for saccade generation. *Journal of Neurophysiology, 85*(2), 804–815.

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, *9*(6), 467–479.

Heitz, R. P., & Schall, J. D. (2012) Neural mechanisms of speed-accuracy tradeoff. *Neuron*, *76,* 616–628.

Heitz, R. P., & Schall, J. D. (2013). Neural chronometry and coherency across speed-accuracy demands reveal lack of homomorphism between computational and neural mechanisms of evidence accumulation. *Philosophical Transactions of the Royal Society of London B, 368,* 20130071.

Hikosaka, O., & Wurtz, R. H. (1983). Visual and oculomotor functions of monkey substantia nigra pars reticulata: IV. Relation of substantia nigra to superior colliculus. *Journal of Neurophysiology, 49,* 1285–1301.

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

Jones, M., & Dzhafarov, E. N. (2014). Unfalsifiability of major modeling schemes for choice reaction time. *Psychological Review, 121,* 1–32.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist, 39*(4), 341–350.

Laming, D. R. J. (1968). *Information theory of choice-reaction times*. New York, NY: Academic.

Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction–time response. *Journal of Experimental Psychology, 72,* 805–811.

Lewandowsky, S., & Farrell, S. (2010). *Computational modeling in cognition: principles and practice*. Thousand Oaks, CA: Sage.

Link, S. W. (1992). *The wave theory of difference and similarity*. Hillsdale, NJ: Erlbaum.

Lo, C.-C., Boucher, L., Paré, M., Schall, J. D., & Wang, X.-J. (2009). Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *Journal of Neuroscience, 29,* 9059–9071.

Lo, C.-C., & Wang, X. J. (2006). Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience, 9,* 956–963.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review, 109,* 376.

Logan, G. D. & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91,* 295–327.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review, 108*, 393–434.

Logan, G. D., Schall, J. D., & Palmeri, T. J. (2015). Neural models of stopping and going. Manuscript in preparation. To appear in B. Forstmann & E. J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience*. Springer Neuroscience.

Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review, 121*(1), 66–95.

Logan, G. D., Yamaguchi, M., Schall, G. D., & Palmeri, T. J. (2013). Inhibitory control in mind and brain 2.0: A blocked-input model of saccadic countermanding.

Mack, M. L., & Palmeri, T. J. (2010). Modeling categorization of scenes containing consistent versus inconsistent objects. *Journal of Vision, 10*(3):11, 1–11.

Mack, M. L., & Palmeri, T. J. (2011). The timing of visual object categorization. *Frontiers in Perception Science*.

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

Munoz, D. P., & Schall, J. D. (2004). Concurrent, distributed control of saccade initiation in the frontal eye field and superior colliculus. In W.C. Hall & A. Moschovakis, (Eds.), *The superior colliculus: New approaches for studying sensorimotor integration* (pp. 55–82). Boca Raton, FL: CRC Press.

Murthy, A., Ray, S., Shorter, S. M., Schall, J. D., & Thompson, K. G. (2009). Neural control of visual search by frontal eye field: effects of unexpected target displacement on visual selection and saccade preparation. *Journal of Neurophysiology, 101*(5), 2485–2506.

Nelson, M. J., Boucher, L., Logan, G. D., Palmeri, T. J., Schall, J. D. (2010). Impact of nonstationary response time in stopping and stepping saccade tasks. *Attention, Perception, & Performance, 72,* 1913–1929.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–57.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266–299.

Nosofsky, R. M., & Palmeri, T. J. (2015). Exemplar-based random walk model. In J. R. Busemeyer, J. T. Townsend, Z. J. Wang, & A. Eidels (Eds.), *Mathematical and computational models of cognition*. Oxford University Press.

Palmer, E. M., Horowitz, T. S., Torralba, A., & Wolfe, J. M. (2011). What are the shapes of response time distributions in visual search? *Journal of Experimental Psychology: Human Perception and Performance, 37,* 58.

Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 324–354.

Palmeri, T. J. (2014). An exemplar of model-based cognitive neuroscience. *Trends in Cognitive Science, 18(2),* 67–69.

Palmeri, T. J., & Cottrell, G. (2009). Modeling perceptual expertise. In I. Gauthier, M. Tarr, & D. Bub (Eds.), *Perceptual expertise: bridging brain and behavior*. Oxford, UK: Oxford University Press.

Palmeri, T. J., & Tarr, M. (2008). Visual object perception and long-term memory. In S. Luck & A. Hollingworth (Eds.), *Visual Memory*, (pp. 163–207). Oxford, UK: Oxford University Press.

Palmeri, T. J., Wong, A. C.-N., & Gauthier, I. (2004). Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences, 8,* 378–386.

Paré, M., & Hanes, D. P. (2003). Controlled movement processing: superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience, 23*(16), 6480–6489.

Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. *Journal of Neuroscience, 26*(35), 8965–8975.

Pouget, P., Logan, G. D., Palmeri, T. J., Boucher, L., Paré, M., & Schall, J. D. (2011). Neural basis of adaptive response time adjustment during saccade countermanding. *Journal of Neuroscience, 31*(35), 12604–12612.

Pouget, P., Stepniewska, I., Crowder, E. A., Leslie, M. W., Emeric, E. E., Nelson, M. J., & Schall, J. D. (2009). Visual and motor connectivity and the distribution of calcium-binding proteins in macaque frontal eye field: Implications for saccade target selection. *Frontiers in Neuroanatomy, 3, 2.*

Purcell, B. A., Heitz, R. P., Cohen, J. Y., & Schall, J. D. (2012). Response variability of frontal eye field neurons modulates with sensory input and saccade preparation but not visual search salience. *Journal of Neurophysiology. 108,* 2737–2750

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review, 117,* 1113–1143.

Purcell, B. A., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2012). From salience to saccades: Multiple-alternative gated stochastic accumulator model of visual search. *Journal of Neuroscience, 32,* 3433–3446.

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

Ratcliff, R. (2013). Parameter variability and distributional assumptions in the diffusion model. *Psychological Review, 120,* 281–292.

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

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

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

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

Ratcliff, R., & Smith, P.L. (in press). Diffusion model. In J. R. Busemeyer, J. T. Townsend, Z. J. Wang, & A. Eidels (Eds.), *Mathematical and computational models of cognition*, Oxford University Press.

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.

Rosenbaum, D. A. (2009). *Human motor control.* (2nd ed.). New York, NY: Academic.

Sato, T., Murthy, A., Thompson, K. G., & Schall, J. D. (2001). Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron, 30,* 583–591.

Sato, T., & Schall, J. D. (2003). Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron, 38*(4), 637–648.

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience, 2,* 33–42.

Schall, J. D. (2004). On building a bridge between brain and behavior. *Annual Review of Psychology, 55,* 23–50.

Schall, J. D., & Cohen, J. Y. (2011). The neural basis of saccade target selection. In S. P. Liversedge, I. P. Gilchrist, & S. Everling (Eds.). *Oxford handbook on eye movements*. Oxford, UK: Oxford University Press.

Schall, J. D., Morel, A., King, D., & Bullier, J. (1995). Topography of visual cortex connections with frontal eye field in macaque: Convergence and segregation of processing streams. *Journal of Neuroscience, 15,* 4464–4487.

Schneider, D. W., & Logan, G. D. (2005). Modeling task switching without switching tasks: A short-term priming account of explicitly cued performance. *Journal of Experimental Psychology: General, 134,* 343–367.

Schneider, D. W., & Logan, G. D. (2009). Selecting a response in task switching: Testing a model of compound cue retrieval. *Journal of Experimental Psychology: Learning, Memory and Cognition, 35,* 122–136.

Scudder, C. A., Kaneko, C. S., & Fuchs, A. F. (2002). The brainstem burst generator for saccadic eye movements: A modern synthesis. *Experimental Brain Research, 142,* 439–462.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology, 86*(4), 1916–1936.

Smith, P. L. (2010). From Poisson shot noise to the integrated Ornstein-Uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. *Journal of Mathematical Psychology, 54,* 266–283.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neuroscience, 27,* 161–168.

Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review, 116,* 283.

Smith, P. L., & Van Zandt, T. (2000). Time-dependent Poisson counter models of response latency in simple judgment. *British Journal of Mathematical* & *Statistical Psychology, 53,* 293–315.

Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews Neuroscience, 3,* 952–964.

Stanton, G. B., Bruce, C. J., & Goldberg, M. E. (1995). Topography of projections to posterior cortical areas from the macaque frontal eye fields. *Journal of Comparative Neurology, 353,* 291–305.

Teller, D. Y. (1984). Linking propositions. *Vision Research, 24,* 1233–1246.

Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience, 25,* 9479–9487.

Thompson, K. G., Hanes, D. P., Bichot, N. P., & Schall, J. D. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *Journal of Neurophysiology, 76,* 4040–4055.

Townsend, J. T. (1990). The truth and consequences of ordinal differences in statistical distributions: Toward a theory of hierarchical inference. *Psychological Bulletin, 108,* 551–567.

Turner, B. M., Forstmann, B. U., Wagenmakers, E. J., Brown, S. D., Sederberg, P. B., and Steyvers, M. (2013). A Bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage, 72,* 193–206.

Umakantha, A., Purcell, B. A., & Palmeri, T. J. (2014). Mapping between a spiking neural network model and the diffusion model of perceptual decision making.

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

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.

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

Verbruggen, F., & Logan, G. D. (2008). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences, 12,* 418–424.

White, C. N., Mumford, J. A., & Poldrack, R. A. (2012). Perceptual criteria in the human brain. *Journal of Neuroscience, 32*(47), 16716–16724.

Wolfe, J. M., Palmer, E. M., & Horowitz, T. S. (2010). Reaction time distributions constrain models of visual search. *Vision Research, 50,* 1304–1311.

Wong, K. F., Huk, A. C., Shadlen, M. N., & Wang, X. J. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Frontiers in Computational Neuroscience, 1,* 1–11.

Wong, K. F., & Wang, X. J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience, 26,* 1314–1328.

Woodman, G. F., Kang, M. S., Thompson, K., & Schall, J. D. (2008). The effect of visual search efficiency on response preparation: Neurophysiological evidence for discrete flow. *Psychological Science, 19,* 128–136.

Zandbelt, B. B., Purcell, B. A., Palmeri, T. J., Logan, G. D., Schall, J. D. (2014). Response times from ensembles of accumulators. *Proceedings of the National Academy of Sciences, 111,* 2848–2853.