

Perceptual Decisions between Multiple Directions of Visual Motion

Mamiko Niwa¹ and Jochen Ditterich^{1,2}

¹Center for Neuroscience and ²Department of Neurobiology, Physiology, and Behavior, University of California, Davis, Davis, California 95618

Previous studies and models of perceptual decision making have largely focused on binary choices. However, we often have to choose from multiple alternatives. To study the neural mechanisms underlying multialternative decision making, we have asked human subjects to make perceptual decisions between multiple possible directions of visual motion. Using a multicomponent version of the random-dot stimulus, we were able to control experimentally how much sensory evidence we wanted to provide for each of the possible alternatives. We demonstrate that this task provides a rich quantitative dataset for multialternative decision making, spanning a wide range of accuracy levels and mean response times. We further present a computational model that can explain the structure of our behavioral dataset. It is based on the idea of a race between multiple integrators to a decision threshold. Each of these integrators accumulates net sensory evidence for a particular choice, provided by linear combinations of the activities of decision-relevant pools of sensory neurons.

Key words: choice; computational model; human; integration; psychophysics; response time

Introduction

Decision making is an essential higher-level neural mechanism for linking perception and action. Huge efforts are currently undertaken to understand neural mechanisms of decision making on various levels, ranging from the behavioral, cognitive, and computational levels to the neurophysiological implementation of the underlying mechanisms. Mainly for reasons of tractability, researchers have largely focused on studying choices between two alternatives [for recent reviews, see Ratcliff and Smith (2004) and Smith and Ratcliff (2004)]. In real life, however, we often face choices between multiple alternatives. Recently, some progress has been made on the theoretical aspects of multichoice decision making (Roe et al., 2001; Usher and McClelland, 2001; McMillen and Holmes, 2006; Bogacz et al., 2007), but experimental datasets that are quantitative enough to test computational models are desperately needed.

An experimental paradigm that has been very helpful in elucidating the neural mechanisms underlying perceptual decisions between two alternatives in both humans and monkeys is the random-dot motion direction discrimination task (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Palmer et al., 2005; Heekeren et al., 2006). We have asked human subjects to perform random-dot motion discrimination between three alternatives,

making use of a random-dot stimulus with multiple motion components. This new stimulus has the advantage of providing simultaneous experimental control over the sensory evidence for each of the three alternatives.

The collected behavioral data span a wide range of both accuracy levels, ranging from chance to perfect performance, and mean response times. We demonstrate that the complex data pattern can be captured by a relatively simple computational model with only five free parameters. This model is based on the idea of a race between three competing neural integrators toward a threshold. Each of these integrators accumulates the net sensory evidence for one of the alternatives. The net sensory evidence is computed as a linear combination of the neural activities of three task-relevant pools of motion-sensitive sensory neurons in extrastriate visual cortex. The model is a generalization and expansion of the computational model that we have used previously to explain both the behavior and the neural activity in the parietal cortex of monkeys performing a random-dot motion discrimination between two alternatives (Ditterich, 2006a,b).

Materials and Methods

Human subjects

Three young adults with normal vision participated in our experiment. Each of the subjects completed multiple experimental sessions (between 9 and 13) of 408 trials each, providing us with between ~3300 and 5100 valid (see below) experimental trials per subject (for details, see Table 1). Subject 1 (S1) had previous experience with visual psychophysics, S2 never had performed a similar experiment before, and S3 was experienced in visual psychophysics and one of the authors.

Experimental setup

The subjects sat in front of a 19 inch flat-screen cathode ray tube video monitor (PF790; viewing distance, 60 cm; ViewSonic, Walnut, CA) with their head on a chin and forehead rest. The visual stimuli were generated by a Macintosh G4 computer running Mac OS 9, Matlab (Mathworks, Natick, MA), and the Psychophysics Toolbox (Brainard, 1997; Pelli,

Received Dec. 16, 2007; revised Feb. 27, 2008; accepted March 20, 2008.

This work was supported by the Air Force Research Laboratory under agreement number FA9550-07-1-0205. The United States government is authorized to reproduce and distribute reprints for governmental purposes, notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the United States government. We are grateful to Ken Britten for providing feedback on a previous draft of this manuscript. We also thank QNX Software Systems for supporting our research through providing their operating system free of charge as part of their QNX-in-Education program.

Correspondence should be addressed to Jochen Ditterich, Center for Neuroscience, University of California, Davis, 1544 Newton Court, Davis, CA 95618. E-mail: jditterich@ucdavis.edu.

DOI:10.1523/JNEUROSCI.5564-07.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/284435-11\$15.00/0

Table 1. Model parameters and data statistics for the pooled dataset and for the three individual subjects

		Subject		
	Pooled data	S1	S2	S3
Model parameters				
g	0.0126	0.0137	0.00769	0.0172
k_n	0.0989	0.0717	0.0907	0.0738
k_s	2.08	2.07	2.75	1.64
k_v	0.261	0.370	0.291	0.326
Residual time (ms)	366	340	413	326
Data statistics				
n_{valid}	12,247	3882	3291	5074
n_{exp}		10	9	13

n_{valid} , Number of valid trials; n_{exp} , number of experimental sessions.

1997) at a frame rate of 75 Hz. The experiment was controlled, and the data were collected by an Intel Pentium IV computer running QNX (QNX, Ottawa, Ontario, Canada) and a modified version of REX (Laboratory of Sensorimotor Research, National Eye Institute, Bethesda, MD).

Eye movements were monitored using an infrared video eye tracker (EyeLink; SR Research, Osgoode, Ontario, Canada). The eye position of the right eye was sampled at 250 Hz. Before each experimental session, the eye tracker was calibrated using the built-in fixation-based calibration routine.

Experimental task and visual stimulus

The experimental task is illustrated in Figure 1. Each trial started with the presentation of a central fixation mark (diameter, 0.4°). The measured fixation location had to remain within 2.0° of the center of the screen throughout the trial (up to the saccadic response). After 500 ms of stable fixation, three targets (diameter, 0.8°) appeared on the screen. They were all located on a virtual circle around the fixation mark with a radius of 8.0°. The target locations were chosen randomly (with an equal spacing) at the beginning of an experimental session and did not change throughout the session. After a random delay [minimum, 0.2 s; maximum (max), 5.0 s; mean, 0.9 s], a multicomponent random-dot pattern was presented at the center of the screen (diameter, 5.0°).

In the original version of the stimulus [as used, e.g., in Shadlen and Newsome (2001), Roitman and Shadlen (2002), and Palmer et al. (2005)], a certain fraction of the dots (defined as the coherence of the stimulus) was moving coherently in a particular direction, whereas the rest of the dots were flickering randomly. Our multicomponent random-dot pattern had up to three coherent motion components embedded. Thus, there were four subpopulations of dots: one of them was moving coherently in a particular direction θ (aligned with one of the choice targets; fraction of dots defined by the coherence of the first component), another one was moving coherently in the direction $\theta + 120^\circ$ (fraction defined by the coherence of the second component), a third one was moving coherently in the direction $\theta + 240^\circ$ (fraction defined by the coherence of the third component), and the rest of the dots were flickering randomly. The stimulus is therefore described by a set of three coherences. Which of the four subpopulations a particular dot belonged to changed randomly over time. As a consequence, the stimulus is not perceived as an overlay of several transparent layers of motion that could be easily separated, but as a mixture of different motion components. For a discussion of transparent random-dot motion stimuli, see, e.g., Treue et al. (2000). Corresponding pairs of dots, responsible for the percept of apparent motion, were presented with a temporal separation of 40 ms (three video frames). The coherently moving dots had a speed of 5°/s, the dot density was 16.7 dots/(deg² · s), and each dot was a little filled square with an edge length of 0.1°. On each trial, the set of coherences was randomly selected from the following list (given as percentages):

(0/0/0)
(5/0/0), (0/5/0), (0/0/5)

(10/0/0), (0/10/0), (0/0/10)
(20/0/0), (0/20/0), (0/0/20)
(40/0/0), (0/40/0), (0/0/40)
(10/10/10)
(20/10/10), (10/20/10), (10/10/20)
(30/10/10), (10/30/10), (10/10/30)
(20/15/5), (20/5/15), (15/20/5), (5/20/15), (15/5/20), (5/15/20)
(30/15/5), (30/5/15), (15/30/5), (5/30/15), (15/5/30), (5/15/30)
(20/20/20)
(30/20/20), (20/30/20), (20/20/30)
(40/20/20), (20/40/20), (20/20/40)
(30/25/15), (30/15/25), (25/30/15), (15/30/25), (25/15/30), (15/25/30)
(40/25/15), (40/15/25), (25/40/15), (15/40/25), (25/15/40), (15/25/40).

This provides a total of 51 different trial types. Each trial type was repeated eight times per experimental session (in random order).

The subjects were instructed to identify the direction of the strongest motion component and to make a saccadic eye movement to the associated choice target (aligned with the identified direction of motion). They were allowed to watch the stimulus for as long as they wanted to and to respond whenever they were ready. After each trial, they received auditory feedback as to whether they had picked the correct target. If the stimulus did not have one strongest motion component, the computer randomly identified one of the targets as being the correct one.

To complete a trial successfully (“valid trial”), the subject had to maintain accurate fixation until the random-dot pattern appeared. Once central fixation was broken, the eye position had to be within 3.0° of one of the three choice targets within 100 ms and had to stay on this target for at least 200 ms.

Data analysis

When analyzing the data, we collapsed across different target locations. Thus, we only cared about the set of coherences and whether the subject picked the target associated with the strongest motion component, the one associated with the intermediate component, or the one associated with the weakest component.

Because in the two-alternative forced-choice (2AFC) version of the task the psychometric function is usually well described by a logistic function of the form $p(\text{correct choice}) = e^{\alpha \cdot C^*} / (1 + e^{\alpha \cdot C^*})$ (Roitman and Shadlen, 2002; Palmer et al., 2005), we used a function of the form

$$p(\text{correct choice}) = \frac{e^{\alpha \cdot C^*}}{2 + e^{\alpha \cdot C^*}}$$

for quantification of the slope of our psychometric functions, with C^* being the difference between the coherence of the strongest motion component and the coherence of the other two components. (We performed this analysis only for trial types with the two lower coherence levels being the same. The probability of a correct choice has to be 1/3 for three identical coherence levels.) The slope parameter α was determined using maximum likelihood estimation. The standard error of the parameter was calculated from the second partial derivative of the log likelihood with respect to the parameter (Meeker and Escobar, 1998).

The response time (RT) was measured as the time between the appearance of the random-dot stimulus and the breaking of central fixation. For obtaining the RT distributions, we eliminated any RTs > 4 s, divided the remaining range of RTs into 10 bins of equal width, and counted the number of RTs within the range of each bin.

Computational model

A general description of the ideas behind our model can be found in the Results section.

Model of the neural representation of the sensory stimulus. Specifically, the mean response of a population of motion-sensitive neurons to a three-component random-dot stimulus with coherences c_1 (in the preferred direction of the pool), c_2 , and c_3 was modeled to be of the form

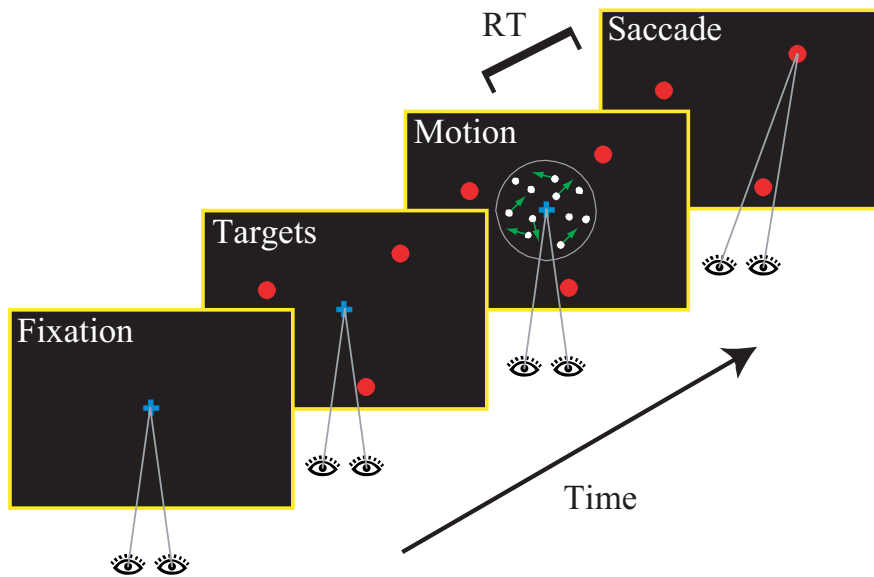


Figure 1. Experimental paradigm. Human observers were asked to make a judgment about the strongest direction of motion in a random-dot pattern with multiple motion components. They were free to watch the stimulus as long as they wanted to and responded with a goal-directed eye movement to one of three choice targets. Choices and RTs were measured.

$$\bar{s}_1 = \frac{g \cdot [c_1 + k_n \cdot (1 - \sum_{i=1}^3 c_i)]}{1 + k_s \cdot (c_2 + c_3)},$$

where g is the overall gain of the sensory response (relationship between neural activity and motion strength). The two additive terms in the brackets reflect the two linear response components: the first one describes the response to the coherent motion in the preferred direction, and the second one describes the response to the noise dots. The term in parentheses reflects the proportion of noise dots in the stimulus. k_n is the relative gain of the response to the noise dots compared with the response to an identical fraction of dots moving coherently in the preferred direction. The term in the denominator reflects the divisive normalization. Because the term in the numerator accurately describes the response to a single-component stimulus, only the coherences of the motion components away from the preferred direction should show up in the denominator. For simplicity, we have chosen a linear term with k_s describing the gain/strength of the divisive normalization. There is probably room for improvement in the mathematical description of the normalization process, but, as we will see, this simple approximation will allow us to capture the structure of the behavioral data quite nicely.

In general, the mean responses of each of the three task-relevant sensory pools can be written as

$$\bar{s}_j = \frac{g \cdot [c_j + k_n \cdot (1 - \sum_{i=1}^3 c_i)]}{1 + k_s \cdot \sum_{i \neq j} c_i}.$$

The variances of the three sensory responses were modeled as

$$\sigma_{s_j}^2 = k_v \cdot \bar{s}_j.$$

We described the outputs of the sensory pools as normal random processes to be able to treat the decision process as a standard diffusion process (based on Brownian motion), which is reasonable if the pools are not too small.

Model of the decision process. In principle, we would have to treat the race between the three integrators mathematically as a three-dimensional diffusion process. However, for the 2AFC case, the decision process has

often been described as a one-dimensional diffusion process with two boundaries instead of a two-dimensional diffusion process. This simplification can be done when one assumes that the two signals that are accumulated by the two integrators are only different in sign and identical in absolute value. Such a situation would result from all of the contributions that a particular pool of sensory neurons makes to the net evidence signals having the same origin. If we make the same assumption in our model, we can also reduce the dimensionality of the problem. We can write the three evidence signals as

$$e_1 = s_1 - \frac{1}{2}s_2 - \frac{1}{2}s_3$$

$$e_2 = s_2 - \frac{1}{2}s_1 - \frac{1}{2}s_3$$

$$e_3 = s_3 - \frac{1}{2}s_1 - \frac{1}{2}s_2.$$

e_3 can be rewritten as

$$e_3 = - \left(s_1 - \frac{1}{2}s_2 - \frac{1}{2}s_3 \right) - \left(s_2 - \frac{1}{2}s_1 - \frac{1}{2}s_3 \right) = -e_1 - e_2.$$

Thus, if e_1 and e_2 are known, e_3 is known. In our model, each of the three evidence signals is integrated over time (see Fig. 3A):

$$i_j(t) = \int_0^t e_j(\tau) d\tau.$$

Because integration is a linear operation, if i_1 and i_2 are known, we also know i_3 . We can therefore rewrite the decision criterion for choosing the third alternative:

$$i_3 > 1$$

$$-i_1 - i_2 > 1$$

$$i_2 < -i_1 - 1.$$

Thus, the third integrator exceeding a value of 1 is equivalent to crossing another linear boundary in the i_1 – i_2 plane. This is illustrated in Figure 3C. The diffusion process always starts at (0, 0) and stops when one of the three boundaries is crossed. The figure shows these boundaries: the solid line ($i_1 = 1$) is the decision boundary for the first alternative, the dashed line ($i_2 = 1$) is the boundary for the second alternative, and the dotted line ($i_2 = -i_1 - 1$) is the boundary for the third alternative.

The two-dimensional diffusion process is described by a drift vector and a covariance matrix. The drift vector is given by $[\bar{e}_1 \bar{e}_2]^T$, the means of the first two evidence signals. Because $[e_1 \ e_2]^T$ can be calculated as

$$\begin{bmatrix} e_1 \\ e_2 \end{bmatrix} = \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} \end{bmatrix} \cdot \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix},$$

$[\bar{e}_1 \ \bar{e}_2]^T$ is given by

$$\begin{bmatrix} \bar{e}_1 \\ \bar{e}_2 \end{bmatrix} = \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} \end{bmatrix} \cdot \begin{bmatrix} \bar{s}_1 \\ \bar{s}_2 \\ \bar{s}_3 \end{bmatrix}.$$

The covariance matrix Σ can be calculated as

$$\begin{aligned} \Sigma &= \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} \\ -\frac{1}{2} & -\frac{1}{2} & 1 \end{bmatrix} \cdot \begin{bmatrix} \sigma_{s_1}^2 & 0 & 0 \\ 0 & \sigma_{s_2}^2 & 0 \\ 0 & 0 & \sigma_{s_3}^2 \end{bmatrix} \cdot \begin{bmatrix} 1 & -\frac{1}{2} \\ -\frac{1}{2} & 1 \\ -\frac{1}{2} & -\frac{1}{2} \end{bmatrix} \\ &= \begin{bmatrix} \sigma_{s_1}^2 + \frac{1}{4}\sigma_{s_2}^2 + \frac{1}{4}\sigma_{s_3}^2 & -\frac{1}{2}\sigma_{s_1}^2 - \frac{1}{2}\sigma_{s_2}^2 + \frac{1}{4}\sigma_{s_3}^2 \\ -\frac{1}{2}\sigma_{s_1}^2 - \frac{1}{2}\sigma_{s_2}^2 + \frac{1}{4}\sigma_{s_3}^2 & \frac{1}{4}\sigma_{s_1}^2 + \sigma_{s_2}^2 + \frac{1}{4}\sigma_{s_3}^2 \end{bmatrix} \\ &= k_v \cdot \begin{bmatrix} \bar{s}_1 + \frac{1}{4}\bar{s}_2 + \frac{1}{4}\bar{s}_3 & -\frac{1}{2}\bar{s}_1 - \frac{1}{2}\bar{s}_2 + \frac{1}{4}\bar{s}_3 \\ -\frac{1}{2}\bar{s}_1 - \frac{1}{2}\bar{s}_2 + \frac{1}{4}\bar{s}_3 & \frac{1}{4}\bar{s}_1 + \bar{s}_2 + \frac{1}{4}\bar{s}_3 \end{bmatrix}. \end{aligned}$$

For calculating the predictions of the model (probabilities of the different choices and RT distributions), we discretized the two-dimensional diffusion process and projected it onto a Markov chain. Ditterich (2006a) has explained this in detail (his section B.5). The only difference between what has been explained there and what we have done here is that we had to use three instead of two decision boundaries. Thus, we had to add a third absorbing state to the Markov chain. The Matlab function `OU_2D_3B_MAR.M`, which has been used for performing the model calculations, is part of the Stochastic Integration Modeling Toolbox (written by J.D.), which can be downloaded from <http://master.peractionlab.org/software/>.

Model fit and predictions. The model parameters were identified by an optimization procedure based on the mean RTs. A multidimensional simplex algorithm (provided by Matlab's Optimization Toolbox) was used to minimize the sum of the squared differences between the mean RTs in the data and the mean RTs predicted by the model, taking the standard errors of the estimated means into account. We used the mean RTs for each combination of coherences, regardless of choice (15 data points). For the model, these were obtained by calculating a weighted sum of the predicted mean RTs for the different choices based on the predicted probabilities of these choices.

For the pooled dataset and for S1 and S3, the decision time distributions were calculated to a maximum decision time of 5 s with a temporal resolution of 25 ms during the optimization process and 10 ms for the optimized model. Because of wider distributions and reaching computer memory limits, we had to use different parameters for S2: during the optimization process, the decision time distributions were calculated to a maximum decision time of 7 s with a temporal resolution of 35 ms; for the optimized model, the decision time distributions were calculated to a maximum decision time of 8 s with a temporal resolution of 15 ms.

Results

Data

To study the neural mechanisms of perceptual decision making between multiple alternatives, we asked human observers to judge the direction of the strongest of three motion components embedded into a new (nontransparent) multicomponent random-dot motion stimulus. The observers were free to watch the stimulus for as long as they wanted to and responded with a goal-directed eye movement (saccade) to one of three choice targets whenever they were ready. We measured how often they picked each target (relative frequencies) and their RTs as a function of the motion strengths (coherences) of the three motion

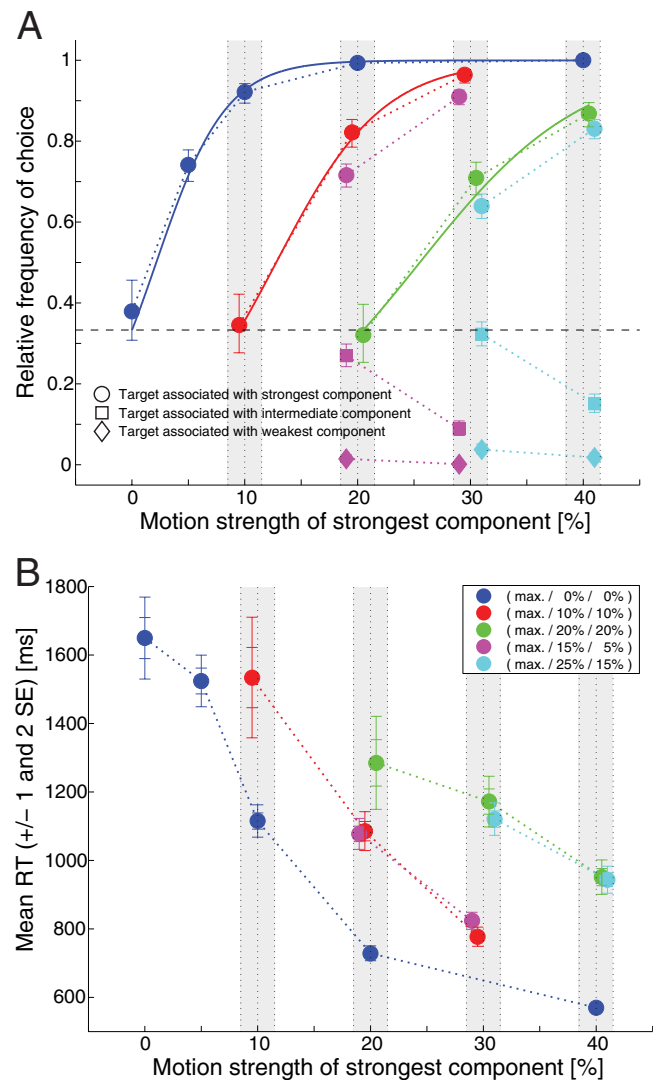


Figure 2. Experimental results. **A**, Relative frequency of choice as a function of the motion strength of the strongest component. The symbols represent the data points, and the error bars represent 95% confidence intervals for the estimated probabilities (see text for details). The dotted lines connect neighboring data points for the same trial type (different combinations of motion strengths of the two weaker components). The trial type is indicated by the color (see **B** for the legend). The solid lines show the function fits used for quantifying the slope of the psychometric functions (see text). The shape of the symbol indicates the choice [circle, correct choice (strongest component has been chosen); square, target associated with the intermediate component has been chosen; diamond, target associated with the weakest component has been chosen]. The dashed line indicates chance level. The symbols (and lines) have been shifted horizontally to avoid overlap. This is indicated by the light gray areas. All symbols in such an area would normally be located on the central vertical line. **B**, Mean RT as a function of the motion strength of the strongest component. The symbols represent the data points, and the error bars represent the mean \pm 1 and 2 SEs. The dotted lines connect neighboring data points for the same trial type. The color code is identical to the one used in **A** (see legend). The light gray areas again indicate a horizontal shift (see **A** for details).

components. A schematic of the experimental task is shown in Figure 1. Please see Materials and Methods for further details. We will first focus on the pooled data from all three subjects (12,247 valid trials), but we will show individual results later.

Choice behavior

Figure 2A shows how often a particular target was chosen as a function of the coherence (fraction of dots that move coherently in a particular direction) of the strongest motion component. The color of the symbols codes for the coherences of the two

weaker motion components (Fig. 2*B*, legend). The shape indicates the chosen target [round, correct target (target associated with strongest motion component); square, target associated with intermediate motion component; diamond, target associated with weakest motion component]. The dotted lines connect symbols of the same color and shape. The error bars represent 95% confidence intervals for the estimated probabilities, calculated according to the method proposed by Goodman (1965). The dashed line indicates chance level (1/3 for a choice between three alternatives). Some of the symbols have been shifted horizontally to avoid overlapping error bars. This is indicated by the light gray areas: any symbol located within such an area is supposed to be located exactly at 10, 20, 30, or 40%.

The dark blue symbols represent the proportion of correct choices for the traditional dot stimulus with only a single coherent motion component. As in the 2AFC case, a gradual improvement of the observed performance is seen with an increase of the motion strength. The choice performance was at chance for a pure noise stimulus (0% coherence) and approached perfect accuracy for strong motion (> 20% coherence). For stimuli with the two weaker motion components having coherences of 10% each (shown in red), chance performance is again observed for all motion strengths being identical, corresponding to a shift of the psychometric function to the right. The psychometric function was also shallower than in the single motion component case. We used function fits, represented by the solid lines in Figure 2*A*, for estimating the slopes of the psychometric functions (for details, see Materials and Methods). The slope (α) of the blue psychometric function was estimated to be 32.3 ± 1.0 (SE), and the slope of the red function was estimated to be 21.2 ± 0.7 . For stimuli with the two weaker motion components having coherences of 20% each (shown in green), the psychometric function was again shifted to the right (chance performance for all coherences being 20%) and was even shallower than in the red case (estimated slope, 13.9 ± 0.5).

For the stimuli with three different motion coherences, we analyzed not only the frequency of correct responses, but also how often the targets associated with the two weaker components were chosen. This is shown in purple for stimuli with the two weaker coherences being 15 and 5% and in cyan for the two weaker coherences being 25 and 15%. The relative frequencies of choices followed the order of the motion strengths: the correct target was chosen most frequently, the target associated with the intermediate component was chosen with an intermediate frequency, and the target associated with the weakest motion component was chosen least frequently.

Response times

Figure 2*B* shows the associated mean RTs. The error bars represent ± 1 and ± 2 SEs (for approximating 95% confidence intervals). As in Figure 2*A*, the symbols were again shifted horizontally to prevent overlapping error bars (again indicated by the light gray areas). Similarly to what has been reported in the 2AFC case (Palmer et al., 2005), our subjects waited longest for the most difficult decisions (on average ~ 1650 ms for pure noise stimuli) and responded much faster when the decision was easy (~ 550 ms for stimuli with a single 40% coherence motion component). For a given coherence of the strongest component, choices were both less accurate (Fig. 2*A*; comparison of blue, red, and green data points in a single gray column) and slower (Fig. 2*B*; same comparison) with increasing motion strengths of the two weaker components. Furthermore, the green chronometric function (stronger distracting motion components) seems shallower than

both the red and the blue ones (weaker distracting motion components). It is also striking that the purple (max/15/5%) and the red (max/10/10%) data points and the cyan (max/25/15%) and the green (max/20/20%) data points tend to overlap. These different stimulus categories are characterized by identical mean values of the coherences of the two weaker components. Thus, the RT seems to be determined by only two degrees of freedom (the coherence of the strongest component and the average coherence of the two weaker components), rather than three (all three motion coherences).

Interestingly, identical means of the coherences of the two weaker motion components did not induce identical choice performance, as can be seen in Figure 2*A*: the purple circles are clearly located below the red circles, and the cyan circles are clearly located below the green circles. This suggests that in our task, accuracy and RT are not controlled by a unique single variable (for a discussion of this observation, see supplemental material, available at www.jneurosci.org).

Between-stimulus-category speed–accuracy effects

This notion is further supported by another observation in our dataset. Within each stimulus category (unique color in Fig. 2), we have observed a characteristic relationship between accuracy and RT, which has been reported in a multitude of previous studies: more accurate choices in response to a more informative sensory stimulus are also faster. However, the introduction of our multicomponent stimulus also allows us to look at between-category speed–accuracy effects. For example, choices in response to (0/0/0%), (10/10/10%), and (20/20/20%) stimuli were all at chance level (Fig. 2*A*) (because of the identical coherence levels for all three directions), but the mean RTs were clearly different (Fig. 2*B*): responses were faster for higher coherence levels. This might appear counterintuitive at first, because one might think of the higher-coherence stimulus as a higher-conflict situation, but we will provide an explanation for this phenomenon based on our computational model in the Discussion section.

Model

The critical question for understanding the neural mechanisms underlying the decision process is whether we can find a quantitative explanation for the observed data pattern. We therefore developed a computational model. We based its architecture on ideas that had proven successful in explaining both the behavior and the neural activity in the parietal cortex of monkeys performing the 2AFC version of the task (Ditterich, 2006a,b). The basic idea is that two integrators (one for each possible choice) compete with each other for reaching a critical activity level or threshold. One of the integrators reaching the threshold terminates the decision process and therefore determines the choice and the decision time. The assumption is that these integrators accumulate net sensory evidence for a particular choice. In the case of discriminating between two (opposing) motion directions, these net evidence signals would be based on reading out the activity of two task-relevant pools of motion-sensitive neurons, each of the pools being tuned to one of the two possible directions of coherent motion. The net evidence signals would result from subtracting the activity of one of the sensory pools from the activity of the other sensory pool.

Model of the decision mechanism

Our task required a choice between three alternatives. We therefore assumed a race between three integrators (one for each

choice). Whichever integrator would reach a critical activity level (threshold) first would determine the choice and the decision time. The outputs of these integrators are labeled as i_1 , i_2 , and i_3 in Figure 3A. Because we cannot measure this decision time directly, but only the RT, we have to make an assumption how these two measures are related. Similar to the assumptions in previous decision models (Luce, 1986), we assumed that the RT is composed of two independent and additive components: the decision time and a residual time combining the durations of all non-decision-related processes, such as, e.g., providing the decision process with the necessary sensory information and executing the eye movement. Expecting the trial-by-trial variability of the decision component of the RT to be much larger than the variability of the residual component, we assumed a constant residual time (first free parameter of the model). The validity of this assumption will be demonstrated later when examining the RT distributions.

We arbitrarily defined the threshold that had to be crossed by any integrator to terminate the decision process as 1. (One of the parameters of a bounded diffusion model can always be chosen arbitrarily without restricting the generality of the model.) As in the 2AFC case, we assumed that the integrators accumulate net sensory evidence. These evidence signals are labeled e_1 , e_2 , and e_3 in Figure 3A. Because our multicomponent stimulus could contain coherent motion in three different directions, we assumed that three task-relevant pools of sensory neurons would have to be read out to accomplish the task, each of them being tuned to one of these three possible directions of motion. These sensory signals are labeled s_1 , s_2 , and s_3 in Figure 3A. But how would these signals have to be combined to obtain the net sensory evidence for a particular choice? The sensory pool that is tuned to the direction associated with a particular choice target should provide evidence for making this choice, whereas the other two pools should provide evidence against it. We used a linear combination of the three sensory signals with a weight of +1 for the signal providing evidence for a particular choice and weights of -0.5 for each of the signals providing evidence against it. This selection of weights makes sure that both pools of sensory neurons providing evidence against a particular choice have the same amount of influence on the decision. Because the three directions of motion are equally spaced, there should be no advantage to either one of them. Furthermore, the selection also makes sure that all weights sum to zero. Thus, when all three pools are equally active, there is, on average, no net evidence for either choice.

Model of the neural representation of the sensory stimulus

How would we describe the response of the three sensory pools to our multicomponent random-dot stimulus? To our knowledge,

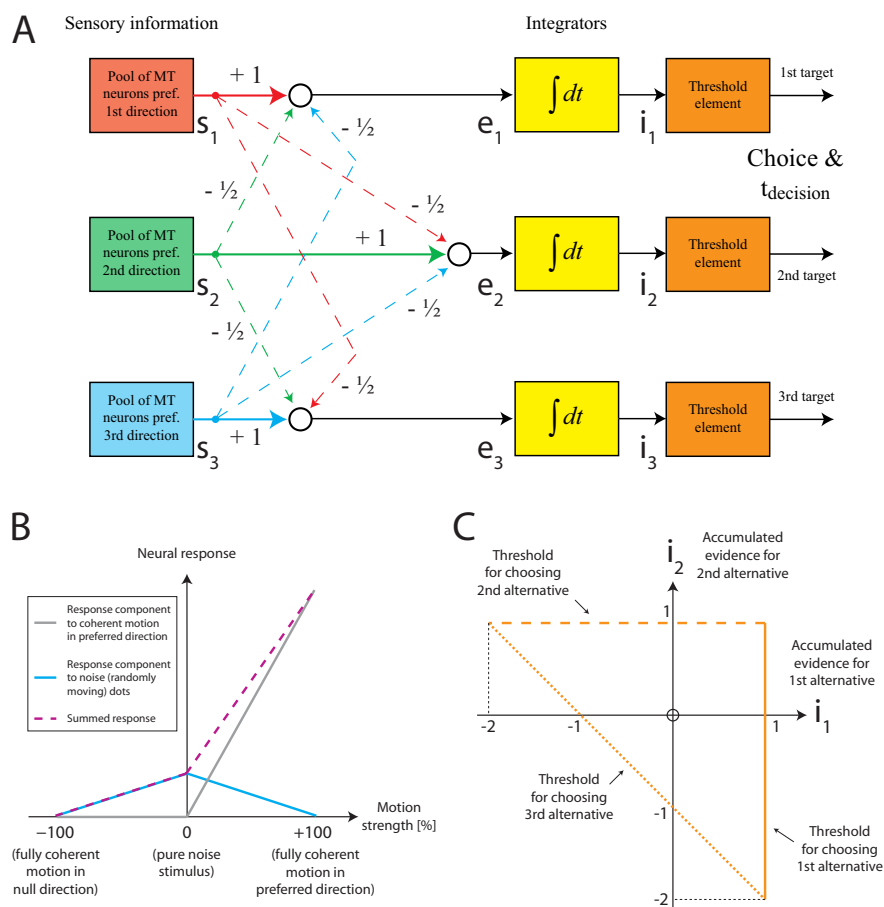


Figure 3. Computational model. **A**, Structure of the model. Three integrators (each associated with one of the three alternatives) race against each other. The integrator output signal (i_1 , i_2 , or i_3) reaching a decision threshold first determines the choice and terminates the decision process. The integrator input signals (e_1 , e_2 , and e_3) are net evidence signals, which are linear combinations of the three relevant sensory signals (s_1 , s_2 , and s_3). Solid arrows indicate positive weights (excitatory connections), and dashed arrows indicate negative weights (inhibitory connections). **B**, Linear response model for the sensory pools. The piecewise linear response function (purple dashed line) is modeled as the sum of two linear response components (gray and blue solid lines; see text for details). **C**, Model implementation as a two-dimensional diffusion process with three boundaries (see text for details). Each trial starts at (0, 0). Which of the three thresholds (solid, dashed, and dotted lines) is crossed first determines the choice. The time of the threshold crossing determines the decision time.

nobody has recorded the physiological responses of motion-sensitive neurons to this type of stimulus so far. However, Britten et al. (1993) have recorded the response of neurons in macaque area MT to the single-component version of the stimulus, having a net motion component either in the preferred direction of the recorded cell or in the opposite (null) direction. They found that, after an initial transient response, the instantaneous activity of a single neuron fluctuated over time, but that the mean firing rate was largely stationary. When characterizing the response to stimuli with different coherences in one of the two possible directions, they also found that the firing rate of an MT neuron could be approximated by a piecewise linear function with two different slopes (one for each direction of coherent motion in the stimulus). Ignoring spontaneous activity, such a response profile could be thought of as resulting from two additive and approximately linear response components: a strong response to coherent motion in the preferred direction of the neuron and a weaker response to the noise dots. Because the fraction of noise dots in the stimulus is determined by the coherence (it is 1 minus the coherence of the stimulus), the latter response component would be strongest for a 0% coherence (pure noise) stimulus and would fall




Figure 4. Comparison between the pooled dataset and the model predictions. **A**, The model was fitted to the mean RTs. The symbols represent the data points as in Figure 2, but the solid lines now connect the model results. Color conventions and horizontal shifts (light gray areas) are as in Figure 2. **B**, Comparison between the probabilities of the particular choices predicted by the model (connected by solid lines) and the relative frequencies of the choices in the data (symbols; not used for the model fit). Shape conventions are as in Figure 2. **C**, Comparison between the RT distributions predicted by the model (solid blue lines) and the RT distributions observed in the experiment (gray histograms). The distributions are shown for the four trial types with the largest numbers of observations.

off linearly for stronger coherent motion in either direction. This is plotted in blue in Figure 3B. The former response component would be strongest for 100% coherent motion in the preferred direction of the neuron, would fall off linearly with weaker motion coherence, would reach zero for a pure noise stimulus, and would stay there for coherent motion in the null direction. This is shown in gray in Figure 3B. Adding these two components together results in a piecewise linear function with two different slopes (Fig. 3B, dashed purple line).

Based on these observations for the single-component stimulus, we assumed that pools of similarly tuned MT neurons would also be driven by two linear and additive components when presented with our multicomponent stimulus: one response component would be driven by the dots moving coherently in the preferred direction of the neurons, and the other component would be driven by the noise dots. Because the other two directions of coherent motion had a separation of 120° from the preferred direction of a particular pool, it is safe to assume that these motion components would not elicit a considerable neural response. The typical direction tuning width of neurons in MT has been reported to be on the order of 40–50° (half-width at half-height) (Albright, 1984; Snowden et al., 1992; Treue et al., 2000). However, it has been suggested that the MT circuitry contains a normalization mechanism, and this normalization mechanism might be triggered by the fact that other subpopulations of neurons are strongly driven by the other coherent motion components. Simoncelli and Heeger (1998) have suggested that this normalization mechanism is divisive. Adding such a normalization mechanism to our MT response model turned out to be essential for being able to explain the behavior (for further information, see supplemental material, available at www.jneurosci.org).

We have therefore modeled the mean response of a population of motion-sensitive neurons to a three-component random-dot stimulus to be driven by two additive, linear response components and to be affected by a divisive normalization mechanism (for details, see Materials and Methods). The description of the mean response adds three free parameters to our model: the first one is the overall gain g of the sensory response (relationship between neural activity and motion strength). We need this flexibility in the model because we have arbitrarily fixed the decision threshold. The second one, k_n , describes the relative gain of the response to the noise dots compared with the response to an identical fraction of dots moving coherently in the preferred

direction. The third one, k_s , describes the strength of the divisive normalization.

As we have already mentioned above, Britten et al. (1993) also demonstrated that the MT response to a random-dot pattern fluctuates over time. Thus, in addition to specifying the mean response, we also have to define its variability. Britten et al. (1993) showed that the variance of the number of spikes that are emitted by a single neuron within a particular time interval is approximately proportional to the mean spike count. Our model requires a description of the variability of the response of a pool of similarly tuned motion-sensitive neurons. How variable this response is depends on a number of factors: the Fano factor (variance-to-mean ratio) of the response of a single neuron, the number of neurons in the pool, the correlation between neurons in the pool (Zohary et al., 1994), and an arbitrary scaling of the input signals to the integrators (resulting from the arbitrary selection of a decision threshold). However, none of the mentioned factors should interfere with the variance remaining proportional to the mean of the response. We therefore introduced the variance-to-mean ratio k_v as the fifth and last free parameter of the model. Additional details regarding the model can be found in the Materials and Methods section. There we also derive how the model can be implemented as a two-dimensional diffusion process with three decision boundaries, which is illustrated in Figure 3C.

Model fit

We determined the model parameters by fitting the mean RTs (for details, see Materials and Methods). The resulting model parameters are shown in Table 1. k_n turned out to be on the order of 10%, which means that we would expect the sensory neural response to a certain number of noise dots to be $\sim 1/10$ of the response of the same number of dots moving coherently in the preferred direction of the motion-sensitive neurons. k_s was ~ 2 , which means that the divisive normalization would be expected to reduce the neural response by $\sim 50\%$ when the sum of the coherences away from the preferred direction of the examined pool is 50%. The residual time was on the order of 350 ms. Thus, given a range of ~ 550 – 1650 ms of mean RTs depending on the difficulty of a trial, the expected range of mean decision times would be between 200 and 1300 ms. For the most difficult trials, the decision process is therefore expected to account for $\sim 80\%$ of the overall RT. Figure 4A shows the results of the fitting process. The symbols, as in Figure 2B, represent the data, but the lines now represent the model. The expected RTs have been calculated

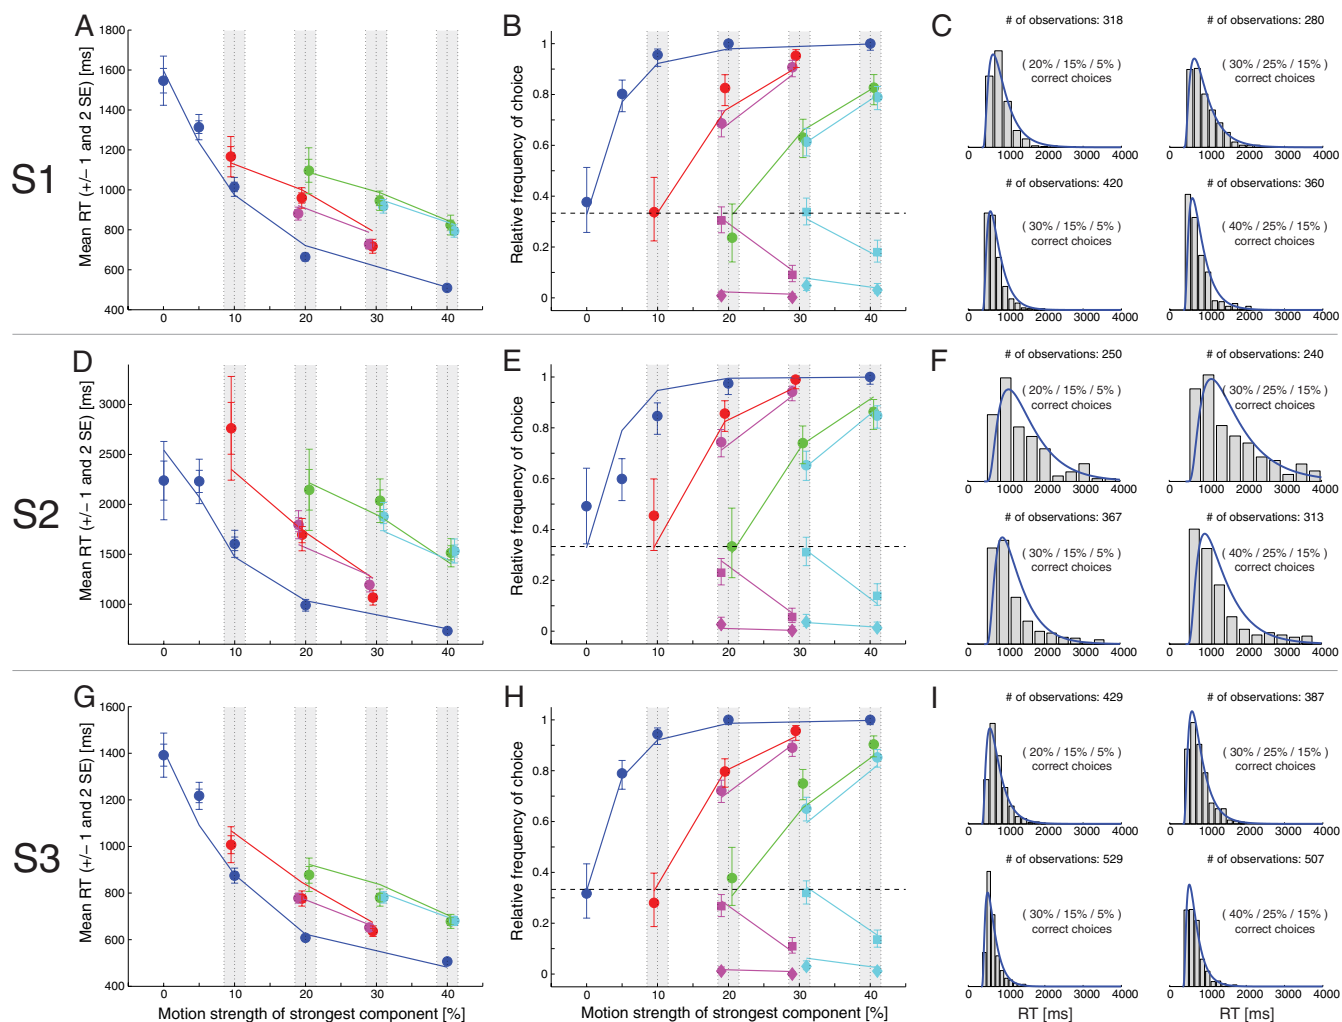


Figure 5. Comparison between the individual datasets and the model predictions. Each row represents one experimental subject. Otherwise, this figure follows the conventions of Figure 4.

for the model for the same combinations of coherences that have been used in the experiment, and these points have been connected by line segments.

Model predictions

Now that we have determined the optimal model parameters, we can look at model predictions and compare them to aspects of the dataset, which have not been used during the fitting process, for testing the model. First, we examined the relative frequency/probability of making a particular choice. Figure 4B shows a comparison between the data (symbols; identical to Fig. 2A) and the model predictions (lines). As can be seen from the figure, the model predictions matched the data quite well. Second, we examined the shape of the RT distributions. This comparison is shown in Figure 4C. The gray histograms represent the data, the blue solid lines the model predictions. We show a good match for the four experimental conditions with three different coherences, because we have the largest numbers of observations for these conditions (resulting from the largest numbers of possible permutations), but a similar match was observed for the other conditions (data not shown).

Individual subjects

So far, we have only looked at the pooled dataset, but we were curious whether the model could also capture the individual behavior of our three subjects. We therefore fitted our model to the

individual datasets of the three subjects (again based on the mean RTs). The resulting optimal model parameters are listed in Table 1, and a comparison between the individual datasets and the best-fitting models is shown in Figure 5 (each row represents one subject). As can be seen from Table 1, the parameter sets obtained for the three subjects were not radically different. S2 showed the largest deviations from the parameter set for the pooled data in terms of a reduced overall gain and an increased residual time, reflecting overall increased RTs in our inexperienced subject. Figure 5D shows that S2 had mean RTs of >2 s for some of the more difficult experimental conditions. Overall, Figure 5 demonstrates that our model was able to capture the individual datasets quite successfully. The first column depicts the fitting results to the individual mean RT data. The second column shows a comparison between the relative frequencies of particular choices and the probabilities that were predicted by the models. The third column compares the observed RT distributions with the ones predicted by the models. Having a closer look at Figure 5, C and I, which shows the RT distributions with a better resolution than Figure 4C, reveals a surprisingly good match for the shape of the RT distributions. The model had more difficulty reproducing the data of our inexperienced subject S2. When looking at Figure 5F, the top two histograms show some evidence of bimodal distributions, which suggests that S2 has been hesitating on a subset of mainly the more difficult trials. Given that our model cannot

capture this aspect of the behavior, it still matches the data surprisingly well.

Discussion

We have presented human behavioral data from a three-choice version of a random-dot motion-discrimination task. We have introduced a new (nontransparent) multicomponent version of the random-dot motion stimulus, which provided us with simultaneous experimental control over the sensory evidence for all three choice alternatives. We have further demonstrated that this task provides a rich quantitative dataset, spanning a wide range of accuracy levels and mean RTs, for testing computational models of multialternative decision making. We were able to demonstrate that a relatively simple model with only five free parameters could capture multiple aspects of the behavior, including mean RTs, relative frequencies/probabilities of particular choices, and RT distributions. The model is based on the idea of a race between multiple integrators to a decision threshold. Each of the integrators accumulates a net evidence signal for a particular choice. These net evidence signals are computed as linear combinations of the activities of task-relevant pools of sensory neurons.

Model-based explanation of the between-stimulus-category speed–accuracy effects

In the Results section, we pointed out that the use of our multicomponent stimulus allowed us to create situations with identical choice performance but different mean RTs: although all stimuli with three identical motion strengths were associated with chance performance, responses were faster for higher coherence. This can be explained in the context of our computational model. The net evidence signals e_j feeding into the integrators are stochastic (fluctuating) signals. They are characterized by an expected value and by a measure of their variability. Whereas the structure of our model predicts that the expected values of the net evidence signals should be zero for stimuli with three identical coherence levels, the variability of the net evidence signals is expected to increase with increasing coherence. This is because of the scaling of the variance of the sensory signals s_j with their mean. The integrators therefore accumulate more noise, which makes an earlier threshold crossing more likely and therefore leads to, on average, faster responses.

If a feedback inhibition model (see below, Other computational models of decision making) were also able to account for our dataset, the faster RTs would not necessarily require an increase in the noise level in addition to an increase in the mean of the sensory signals. For binary choices, Bogacz et al. (2006) pointed out that the mutual inhibition model is approximately equivalent to a drift diffusion model with a decision criterion that moves closer to the starting point of the random process as the sum of the means of the sensory signals increases, which would also lead to a reduction in the expected RTs.

Residual time

The residual times reported here ranged from 326 to 413 ms, which is in accordance with a previous human study of the 2AFC version of the experiment (Palmer et al., 2005). We observed very good matches between the real and predicted shapes of the RT distributions, which justifies our simplifying assumption of no trial-by-trial variability in the residual time. Thus, most of the

trial-by-trial variability in the RT was contributed by variations in the decision time.

A previous study of discriminating multiple directions of motion

A previous microstimulation study of discriminating between multiple possible directions of visual motion in a random-dot pattern by Salzman and Newsome (1994) suggested that different subpopulations of motion-sensitive neurons can be read out individually for making a decision. These authors also found that the relative frequencies of the observed choices were well described by a polychotomous logistic regression, but they did not discuss how the decision-making mechanism might be implemented. In contrast, our model suggests how the different subpopulations of motion-sensitive neurons are read out and how these signals are combined and processed for making a decision.

Other computational models of decision making

The structure of our current model of three-alternative decision making provides generalizations for two key elements of our previous model for making choices between two alternatives (Ditterich, 2006a,b): the race to threshold between two competing integrators is replaced by a race between three competing integrators; and the difference between the activities of two relevant sensory pools is replaced by a linear combination of the activities of three relevant sensory pools. An extension of this mechanism to more than three alternatives is straightforward, as discussed below.

Compared with multialternative decision field theory (MDFT) (Roe et al., 2001), which is based on psychological concepts, our model is more physiologically motivated and operates with neural responses. However, there are structural similarities between both models: both models make use of linear combinations, using the same set of weights (MDFT does so when calculating the valences from the weighted evaluations; our model uses them for deriving the net evidence signals from the sensory responses), and both models are based on temporal integration. Furthermore, our decision rule is equivalent to the rule for internally controlled decisions in MDFT. In contrast to MDFT, which also allows the competing integrators to influence each other, our model is able to explain the observed behavior without relying on lateral interactions between the integrators.

How does our model compare with the leaky, competing accumulator model (LCAM) (Usher and McClelland, 2001)? Both models are based on a race between a number of accumulators (or integrators), one for each possible choice. The integrator that reaches a decision criterion first determines the choice and terminates the decision process. However, the integrators in our model are perfect, whereas the accumulators in the Usher and McClelland model are leaky. As we have pointed out in the discussion of our previous model for the 2AFC version of the task (Ditterich, 2006a), based on the available data we were only able to provide a lower bound for the integration time constant, but it was not possible to constrain the time constant to a narrow interval. We therefore also do not expect our current dataset to be able to distinguish between leaky and perfect integration. Whereas the time constant of integration is not too critical in our model as long as it is not too small, it is a critical parameter in the LCAM. We will get back to this issue below. Another difference between the two models concerns the type of inhibition. Whereas our

model is based on feedforward inhibition, the Usher and McClelland model relies on feedback inhibition and the sensory inputs to the accumulators are only excitatory. We did not test explicitly whether the Usher and McClelland model would be able to explain our dataset, because this would have required a completely different approach from the one we have taken (simulation based rather than numerical evaluation of the predicted model behavior). However, for choices between two alternatives, Bogacz et al. (2006) have been able to demonstrate that for a particular parameter range (when the model is balanced, which means that decay and inhibition are equal, and when these parameters are not too small), the dynamics of the mutual inhibition model closely approximates the dynamics of the drift diffusion model (and therefore a perfect integrator model with feedforward inhibition). We would therefore expect that a mutual inhibition model could not be ruled out on the basis of our behavioral dataset.

McMillen and Holmes (2006) performed a more detailed analysis of the LCAM and its optimality. For the 2AFC problem, it has been shown theoretically that the optimal statistical procedure (in the sense of minimizing the sample size or sampling time, when assuming a constant rate of information arrival, for a given error rate) is the sequential probability ratio test (SPRT) (Wald and Wolfowitz, 1948), which is implemented by the drift diffusion model. For choices between multiple alternatives, it has been shown that the multihypothesis sequential probability ratio test (MSPRT) is asymptotically optimal (Dragalin et al., 1999). Dragalin et al. considered two different versions of the test: MSPRTa, which stops when the largest posterior probability exceeds a threshold, and MSPRTb, which stops when the ratio between the largest and the second largest posterior probabilities exceeds a threshold. Bogacz and Gurney (2007) discussed a potential neural implementation of MSPRTa in the basal ganglia. McMillen and Holmes (2006) pointed out that the equivalent of MSPRTb in the context of leaky, competing accumulators would be a max-versus-next test, using a decision rule that terminates the decision process when the difference between the two largest accumulator values exceeds a threshold. However, they further demonstrate that for a balanced LCAM with decay and inhibition not being too small, the absolute test (stopping the decision process when any accumulator value exceeds the decision threshold) is nearly indistinguishable from the max-versus-ave test (stopping the decision process when the difference between the largest integrator value and the average of the other integrator values exceeds a threshold), which, in turn, is a good approximation of the asymptotically optimal max-versus-next test. Our model makes use of an absolute decision rule, but the net evidence signals feeding into the integrators already have max-versus-ave structure as a result of the feedforward inhibition. Similar to the way in which the drift diffusion model (and, thus, the stationary version of our previous model) implements the SPRT in the 2AFC case, we expect the mechanism presented here to provide a good approximation of the MSPRT.

Extension to more than three alternatives

Our proposed decision mechanism can easily be extended to an arbitrary number of alternatives. The number of integrators would then have to be increased to one per alternative. The decision rule would remain the same: the first integrator to cross the decision threshold determines the choice and terminates the decision process. The net evidence signals sent

into the integrators would all result from linear combinations of the available sensory signals, representing the difference between the neural activity providing direct evidence for a particular choice and the average of the neural activities providing evidence for the other alternatives and therefore against this particular choice. However, we would assume that as the available options would become more similar, or move closer together spatially when thinking of goal-directed movements (as the saccades in our case), it might very well be the case that lateral inhibition could no longer be neglected. In this case, it might be necessary to replace the discrete integrators by a continuous dynamic medium (a “map” or “dynamic field”), as has been suggested in dynamic field theory (Wilimzig et al., 2006).

In summary, we have seen that our behavioral dataset on multialternative decision making is well explained by a straightforward decision mechanism based on linear combinations of sensory evidence signals and integration to threshold. We expect that physiological experiments will help us distinguish between different proposed models of multialternative decision making and provide us with further insight into how the calculations proposed by the computational models are implemented biophysically.

References

- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130.
- Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput* 19:442–477.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765.
- Bogacz R, Usher M, Zhang J, McClelland JL (2007) Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos Trans R Soc Lond B Biol Sci* 362:1655–1670.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Ditterich J (2006a) Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw* 19:981–1012.
- Ditterich J (2006b) Evidence for time-variant decision making. *Eur J Neurosci* 24:3628–3641.
- Dragalin VP, Tartakovsky AG, Veeravalli VV (1999) Multihypothesis sequential probability ratio tests. I. Asymptotic optimality. *IEEE Trans Inf Theory* 45:2448–2461.
- Goodman LA (1965) On simultaneous confidence intervals for multinomial proportions. *Technometrics* 7:247–254.
- 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. *Proc Natl Acad Sci USA* 103:10023–10028.
- Luce RD (1986) Response times: their role in inferring elementary mental organization. New York: Oxford UP.
- McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. *J Math Psychol* 50:30–57.
- Meeker WQ, Escobar LA (1998) Statistical methods for reliability data. New York: Wiley.
- Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5:376–404.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–367.
- Roe RM, Busemeyer JR, Townsend JT (2001) Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol Rev* 108:370–392.

- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
- Salzman CD, Newsome WT (1994) Neural mechanisms for forming a perceptual decision. *Science* 264:231–237.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Simoncelli EP, Heeger DJ (1998) A model of neuronal responses in visual area MT. *Vision Res* 38:743–761.
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–168.
- Snowden RJ, Treue S, Andersen RA (1992) The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp Brain Res* 88:389–400.
- Treue S, Hol K, Rauber HJ (2000) Seeing multiple directions of motion-physiology and psychophysics. *Nat Neurosci* 3:270–276.
- Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *Ann Math Stat* 19:326–339.
- Wilimzig C, Schneider S, Schoner G (2006) The time course of saccadic decision making: dynamic field theory. *Neural Netw* 19:1059–1074.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* [Erratum (1994) 371:358] 370:140–143.