J Neurophysiol 93: 2974–2986, 2005.
First published November 17, 2004; doi:10.1152/jn.00508.2004.

Using Neuronal Latency to Determine Sensory–Motor Processing Pathways in Reaction Time Tasks

James J. DiCarlo[1,2] and John H. R. Maunsell[1]

[1]*Howard Hughes Medical Institute and Division of Neuroscience, Baylor College of Medicine, Houston, Texas; and* [2]*McGovern Institute for Brain Research and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts*

Submitted 14 May 2004; accepted in final form 15 November 2004

DiCarlo, James J. and John H. R. Maunsell. Using neuronal latency to determine sensory–motor processing pathways in reaction time tasks. *J Neurophysiol* 93: 2974–2986, 2005. First published November 17, 2004; doi:10.1152/jn.00508.2004. We describe a new technique that uses the timing of neuronal and behavioral responses to explore the contributions of individual neurons to specific behaviors. The approach uses both the mean neuronal latency and the trial-by-trial covariance between neuronal latency and behavioral response. Reliable measurements of these values were obtained from single-unit recordings made from anterior inferotemporal (AIT) cortex and the frontal eye fields (FEF) in monkeys while they performed a choice reaction time task. These neurophysiological data show that the responses of AIT neurons and some FEF neurons have little covariance with behavioral response, consistent with a largely "sensory" response. The responses of another group of FEF neurons with longer mean latency covary tightly with behavioral response, consistent with a largely "motor" response. A very small fraction of FEF neurons had responses consistent with an intermediate position in the sensory-motor pathway. These results suggest that this technique is a valuable tool for exploring the functional organization of neuronal circuits that underlie specific behaviors.

INTRODUCTION

It is common in the study of higher brain regions to describe neuronal representations as being primarily "sensory" or "motor" based on whether modulation of activity is more closely linked to the delivery of a stimulus or the initiation of a response. However, most neurons may occupy positions that are intermediate between these extremes, and be poorly described by either label. Methods that assign neurons positions along a continuous sensory–motor transformation are likely to prove valuable for understanding the role of neurons and structures in particular behaviors.

In tasks where the subject must respond promptly, the timing of the onset of a neuron's response can provide an approach to evaluating its involvement in, and position along, a sensory–motor processing pathway supporting a particular behavior. The technique described here (termed the "RT–NL technique") is aimed at revealing sensory–motor processing pathways through a quantitative examination of mean neuronal latency (NL), mean reaction time (RT), and trial-by-trial covariance of NL and RT. We are not the first to examine neuronal timing in behavioral tasks (Commenges and Seal 1985, 1986; Lamarre and Chapman 1986; Lamarre et al. 1983; Seal et al. 1983; Thompson and Schall 2000), but we seek to further develop and apply these ideas.

An underlying assumption in this report is that, in RT tasks, neuronal activity generated by sensory transducers is transmitted through a potentially branching path of neuronal connections that ultimately leads to the motor neurons whose activity produces a behavioral response to the stimulus. We refer to the sensory transducers, motor neurons, and the neuronal elements that link their activity in a feed-forward way as the neuronal "processing chain" that mediates the behavior. We do not assume that all RT tasks are carried out by a fixed set of neuronal connections. However, we do assume that, when the RT task and context are held constant, the processing pathway underlying that particular RT task is also constant in that particular brain structures and neurons within those structures are responsible for the sensory–motor transformation from stimulus to behavioral response. Neurons whose activity is modulated after sensory stimulation and before the behavioral response and that contribute, however indirectly, to the initial activation of the motor neurons underlying the behavioral response are considered part of the processing chain.

Many neurophysiological studies have examined the mean NL of neurons in different brain areas (e.g., Robinson and Rugg 1988; Schmolesky et al. 1998), and it is obvious that such measures can be used to assign neurons along a sensory–motor continuum. However, the examination of another measure for each neuron, the *covariance* of NL and RT, has received little attention, even though it can provide additional information beyond that conveyed by the mean NL. For example, the mean NL will be expected to be large (i.e., long latency) for all neurons that are activated many synapses away from the sensory transducers—a potentially large number of neurons. However, the covariance of the NL and RT will be large only for neurons that are closely related to the transformation of the sensory signal to the motor response or the motor response itself—neurons that neurophysiologists are especially interested in. By examining both the mean NL and the covariance of NL and RT for individual neurons in different areas, the RT–NL technique seeks to provide new information about the circuits that underlie specific behaviors.

The goal of the current study was to use the RT–NL technique in neurophysiological experiments, both to test its feasibility and reliability and to potentially reveal new information about the role of specific brain areas in particular behaviors. The work described here suggests that the RT–NL technique may provide a valuable tool for assigning specific

Address for reprint requests and other correspondence: J. J. DiCarlo, Massachusetts Institute of Technology, E25-242, 45 Carleton Street, Cambridge, MA 02139 (E-mail: dicarlo@mit.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

 0022-3077/05 $8.00 Copyright © 2005 The American Physiological Society

functional relationships to different neurons and brain structures in different behaviors.

METHODS

In RT tasks, neurons located at anatomically early stages of processing ("sensory" neurons) respond with mean latencies that are short relative to behavioral responses, whereas neurons late in processing ("motor" neurons) become active only shortly before the RT. Hypothetical examples of the trial-by-trial relationship between NL and RT are shown in Fig. 1. The raster plot in Fig. 1*A* shows responses that might be recorded from a sensory neuron to repeated presentation of a given stimulus, whereas the plot in Fig. 1*B* shows activity that might be recorded from a motor neuron.

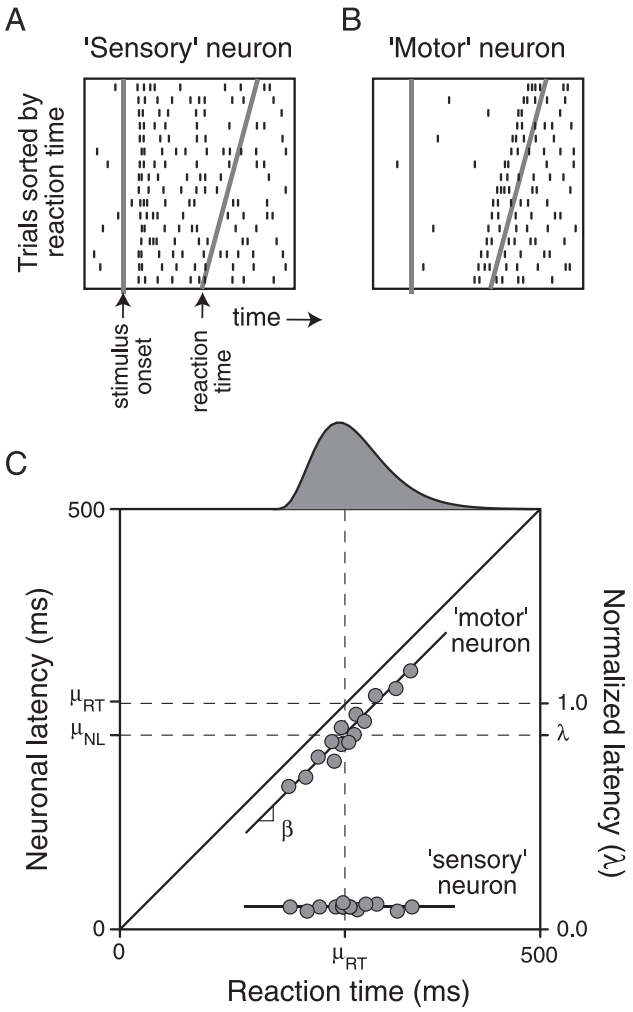

FIG. 1. Hypothetical observations from sensory and motor neurons during a reaction time task. *A* and *B*: hypothetical spike data from "sensory" and "motor" neurons during performance of a sensory–motor reaction time (RT) task. Abscissa shows time relative to stimulus onset. Trials have been sorted based on hypothetical RTs (with RT increasing uniformly from trial to trial for clarity). Each tick mark represents the occurrence of a spike. Sensory neuronal latency (NL) response reliably covaries with the time of stimulus onset (follows the stimulus onset by a fixed time interval), whereas the latency of the motor neuronal response covaries with RT (proceeds the RT by a fixed time interval). *C*: each point represents one trial for either the sensory or motor neuron. Abscissa shows the RT observed on that trial; the left ordinate shows the NL observed on that trial; the right ordinate shows the same NL normalized by the mean RT over all trials (μ_{RT}). A line has been fit through each set of points. Mean of the neuronal latencies of each group of data (μ_{NL}) divided by the mean RT is the normalized mean NL (λ). Slope of the line (β) is a normalized measure of the covariance of the NL and the RT across trials. Hypothetical distribution of RTs is shown at the *top* of the panel.

Each point in the scatter plot in Fig. 1*C* plots the latency for the change in activity for one of these neurons in one trial (NL) against the RT on that trial. Lines show the best least-squares linear fit to each set of data. These lines show that these neurons differ not only in the mean time of their responses (reflected by the vertical offsets of the lines), but also in the trial-by-trial covariance of those response times with RT (reflected by the slopes of the lines). Thus Fig. 1 suggests that neurons near the ends of neuronal processing chains can be distinguished based on either their mean NL or the trial-by-trial covariation of their NL and RT, and raises the possibility that these measures could be used to assign neurons to intermediate positions on processing chains.

We define two terms related to the mean NL and covariance between NL and RT. The mean NL for each neuron in Fig. 1*C* corresponds to the mean of the *y*-axis values of its points. We describe a neuron's normalized mean NL with a value termed λ, which is found by dividing its mean NL by the mean RT. Thus λ is a unitless value that progresses from near zero for neurons that, on average, become active immediately after stimulus onset, to near one for neurons that become active immediately before the response. This progression from zero to one suggests that normalized mean NL (λ) may be useful for determining a neuron's position along the processing chain. Indeed, measures of mean NL have been used in many studies in an effort to order neurons or brain areas along processing pathways (e.g., Bullier and Henry 1979; Gawne et al. 1996; Maunsell and Gibson 1992; Nowak and Bullier 1997; Robinson and Rugg 1988; Schroeder et al. 1998; Zeki 2001).

Additional information can be gained from a statistical measure of the trial-by-trial association of RT and NL. We measure the association of RT and NL as the covariance of those variables, cov(*RT*, *NL*). Because we are not interested in the absolute covariance per se, but the fraction of the RT variance that is associated with the NL, we define a normalized measure of association, β, which is the covariance of NL and RT divided by the RT variance [$\beta = \text{cov}(RT, NL)/\sigma^2_{RT}$]. Thus β is a unitless value that progresses from near zero for neurons that have little correlation with RT (Fig. 1*A*) to near one for neurons that have activity closely correlated with the timing of the behavioral response (Fig. 1*B*). Graphical intuition about this measure can be gained by realizing that this definition of normalized covariance (β) is also the definition of the slope of the best-fitting line resulting from the linear regression of unbiased trial-by-trial estimates of NL on RT (i.e., the linear regression for hypothetical data plotted like those in Fig. 1*C*). However, it is important to emphasize that our analyses do not estimate the normalized covariance (β) by performing this linear regression (see following text), nor do they rely on a linear relationship between RT and NL. Moreover, although regression is often used to predict one variable given the value of another, this makes little physical sense in this situation because the value of RT on a particular trial cannot cause the value of an earlier NL (see DISCUSSION). We simply use this definition of β because it captures the association between NL and RT, and thus can inform about neuronal processing chains.

Although λ and β are both expected to progress from zero to one along a neuronal processing chain, they need not take the same value for a given neuron. For example, if most of the variance in RT is generated in later ("motor") stages, then β will remain small until those stages, and only neurons with large λ values would have β values much greater than zero. Only if RT variance accumulates uniformly along the processing pathway will λ and β increase in tandem. Thus by examining both of these values for individual neurons and the distribution of these values across all activated neurons we can gain insights into the position of neurons and structures along the processing pathway and the way that NL and RT variance accumulate along it (see DISCUSSION).

Measuring mean NL (λ) and covariance of NL and RT (β)

The obvious approach to determining mean NL and the covariance between NL and RT would be to measure NL and RT for many individual trials, but measuring the latency of spiking neurons on individual trials is problematic. Spike times provide only a sparsely sampled estimate of the assumed underlying rate function. As a result, even if the underlying rate function of a neuron changes rapidly (at the NL we seek to determine), the spikes from a single trial cannot generally determine NL with precision. Thus although methods exist for assigning latencies to spike trains from single trials (e.g., Commenges et al. 1986a), our attempts with this approach showed that the variance of the trial latencies overwhelmed measurements of λ and β. For this reason we focused on alternative approaches that estimate λ and β from the combined data from all trials. Although these approaches do not provide neuronal latencies from individual trials, they provide estimates of mean NL and covariance between NL and RT that are more reliable and less biased that those based on neuronal latencies from individual trials.

We used two methods to measure λ and β from the spike trains of neurons. The basic idea behind these methods is illustrated in Fig. 2, which shows simulated data from one neuron from several trials. The rasters lines have been sorted by RT, which is marked by a heavy sigmoidal line, aligned to stimulus onset (heavy vertical line) and truncated at fixed offsets before stimulus onset and after the RT for each trial. To determine λ and β, we assumed that at the start of every trial the underlying rate function is constant, and, at some time in the trial, changes to another constant (i.e., the true latency is a step change to either a higher or a lower underlying rate). The algorithms we used to determine this step change in rate are reasonably robust to departures from these assumptions.

Finding the best estimates of λ and β can be visualized as searching for the horizontal offset (λ) and deformation (β) of the rate-change

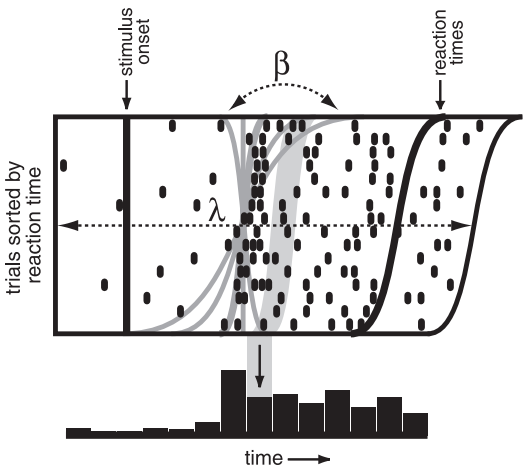

FIG. 2. Methods used to determine λ and β for neuronal data. The method applied to each data set can be visualized by plotting the spike times from each trial relative to the time of stimulus onset (abscissa). Trials in the raster plot are sorted by RT. Two methods used to determine λ and β [least-squares estimation (LSE) and maximum likelihood estimation (MLE)] can both be visualized as seeking to position a rate change line so that the mean firing rates on each side of the line (over all trials) is most different. Optimal rate change line (heavy gray line) and several others (light gray lines) are drawn in the center of the raster. Horizontal position of the midpoint of the line corresponds to λ. Optimal deformation of the line between vertical ($\beta = 0$) and the sigmoidal shape of the RT line ($\beta = 1$) corresponds to β. That is, β is the constant of proportionality used to shift each trial in time according to the RT on that trial so that the spikes from all trials are brought into optimal alignment. The LSE method finds the rate change line that yields the smallest error for the rates over the surfaces on either side of the line. The MLE method uses interspike intervals from combined trials. Trials are combined with different time offsets corresponding to different rate change lines. Gray band on the raster plot indicates one set of time offsets. See text for details.

border on the raster plot of Fig. 2 that best divide the neuronal activity into two regions of constant firing rate (i.e., pre-NL firing rate to the left of the border, post-NL firing rate to the right). Deformations of the rate-change border are restricted to correspond to fixed proportions of RT: A rate-change border corresponding to a β of zero is a vertical line, and rate-change borders corresponding to nonzero values of β are produced by adding a horizontal offset on each raster line that is a given fraction of the RT for that trial. This is equivalent to borders that are straight lines in plots like the one in Fig. 1*C*.

The first of the two methods was a least-squares estimate (LSE). This approach assigned a squared error function to the data under each assumption of λ and β (i.e., each assumption about the offset and deformation of the rate-change border). An average spike rate was determined as the mean firing rate on each side of the line, and the error was computed as the sum of the squared differences over all 1-ms bins for all trials between the rate for each bin and the average rate for its side of the border. The final estimates of λ and β were taken as the horizontal offset ($\hat{\lambda}$) and deformation ($\hat{\beta}$) that produced the minimum error.

The second method was a maximum likelihood estimate (MLE). This was an extension of the maximum likelihood method described by Seal and Commenges (Commenges et al. 1986a; Seal et al. 1983), which finds the time in a single trial where it is maximally likely that the interspike intervals observed before and after that time were drawn from different distributions. We collapsed all the trials into a single spike sequence before applying their method. For Poisson spike processes, collapsing across all trials does not change the underlying statistics of the spike process (normalized for spike rate) because superimposed Poisson processes give a Poisson process (Cinlar 1975). By shifting each trial in time by an amount proportional to the RT for that trial before collapsing, we could test different β values (i.e., different proportional shifts). The gray band in the raster plot of Fig. 2 illustrates this shifting for a particular β value (although no binning occurred in the actual analysis). For each β value tested, spikes were compiled into a single interspike interval distribution, after which the method of Commenges and Seal (Commenges et al. 1986a) was used to estimate λ. The final estimates of λ and β were taken as those that yielded the overall maximum likelihood.

With either method, we based measurements on spikes collected from 100 ms before the stimulus onset to 100 ms after RT on each trial. For the MLE method, it was necessary to compensate for fewer trials contributing to the extremes of the collapsed histogram following time shifts to avoid artificially large interspike intervals in these regions. For both methods, we used a brute-force search over a range of λ values from about −0.4 (100 ms before stimulus onset) to about 1.4 (100 ms after the mean RT) and β values from −0.5 to 1.5 in steps of 0.02 and 0.04, respectively. We then applied successively smaller searches of reduced ranges of λ and β values until the best pair of values was obtained with a resolution of 0.005 (λ) and 0.01 (β).

We used simulated data to confirm that these methods of estimating λ and β were not biased and to assess the reliability of those estimates (see Fig. 4). For each simulation, 200 trials of simulated spike data were created by a Poisson spike generator driven by a constant rate function that stepped from 5 to 25 spikes/s at a particular point (NL) in each trial (the median pre-NL and post-NL rates observed in the neurophysiologic recordings were 6 and 27 spikes/s, and a median of 187 trials were obtained). Reaction times for each trial were drawn from a normal distribution (mean 270 ms, SD 40 ms; comparable to that seen in the behavioral task). The bias of the LSE and MLE estimates of λ and β were examined with simulated data sets in which λ was varied from 0.0 to 1.2 and β was varied from −0.2 to 1.2. The maximum bias (mean of the absolute difference from the true value) was 0.005 ($\hat{\lambda}$) and 0.017 ($\hat{\beta}$) (medians 0.001 and 0.005) for the LSE method and 0.018 ($\hat{\lambda}$) and 0.025 ($\hat{\beta}$) (medians 0.005 and 0.004) for the MLE method. In sum, both methods produced estimates of λ and β that had no appreciable bias over the range of physiologically plausible values.

Bootstraps (Efron and Tibshirani 1998) were used to assign confidence intervals to values of $\hat{\lambda}$ and $\hat{\beta}$ determined for each recorded neuron. To do this, the entire analysis was rerun 100 times using exactly the same methods except that the trials included in each of the 100 runs were resampled with replacement from the original set of trials. The SDs of the distributions of $\hat{\lambda}$ and $\hat{\beta}$ obtained from these 100 runs were taken as the SEs of $\hat{\lambda}$ and $\hat{\beta}$ for each neuron (Efron and Tibshirani 1998).

Both of the methods assume that the RT value measured in each behavioral trial in the experiments (i.e., the time that the saccade was detected to start) is a good estimate of the true RT. Positive bias in our measure of RT would bias $\hat{\lambda}$ toward zero, and variability in our measure of RT will bias $\hat{\beta}$ toward zero. Both of these effects are small in practice. In particular, given the distribution of RTs observed (mean ~ 270 ms, SD ~ 40 ms), both theory and simulations show that a 5-ms bias in the measured RT would produce about 2% bias in $\hat{\lambda}$ and 5 ms of uncorrelated variability (i.e., SD of random error) in the measured RT would produce about 2% bias in $\hat{\beta}$. Our measurement of the response saccade start time (i.e., RT) is at least this accurate and precise (DiCarlo and Maunsell 2000, 2003).

There is reason to question whether these methods for estimating λ and β could confound correlation between NL and RT with correlation between response magnitude and RT. For example, if the magnitude of a response with a finite rise time varied with RT, a measure of latency based on a particular rate of firing would similarly vary, even if the onset of the response did not. Although the potential for this confound exists, most of the neuronal responses we recorded did not have enough variance in response magnitude to explain the range of correlation between NL and RT that we observed.

Animals and surgery

Two male monkeys (*Macaca mulatta*) were used in this study (weighing 4.5 and 4.7 kg). Before behavioral training, aseptic surgery was performed to attach a head post to the skull and to implant a scleral search coil in the right eye. After 2–3 mo of behavioral training (below), a second surgery was performed to place a recording chamber to reach the anterior portion of the left inferotemporal cortex (AIT; Horsley–Clark chamber center = 15 mm A, 12 mm L). After several weeks of recording from AIT, a second chamber was placed over the right frontal eye field (FEF; Horsley–Clark chamber center = 22–23 mm A, 17–19 mm L).

Horizontal and vertical eye positions were monitored using the scleral search coil (Robinson 1963). Saccades greater than about 0.2° were reliably detected in real time using speed criteria (details described elsewhere: DiCarlo and Maunsell 2000). All animal procedures complied with the standards of the Baylor College of Medicine Animal Research Committee and the American Physiological Society.

Behavioral task

The animals performed a visual-shape identification task in which two visual shapes required different motor responses (saccades). For each animal, each target shape was assigned a different response location, and this mapping never changed. When a shape appeared, the animal was required to signal its identity by making a saccade directly to one of two fixed locations (Fig. 3*A*). These two response locations were continuously indicated by identical white squares ($0.6 \times 0.6°$, 46 cd/m^2). Saccades that ended within a window ($\pm 3°$ h and $\pm 3°$ v) centered on each response location were scored as a response to that location. Correct responses produced a juice reward and a brief tone. Reaction time was defined as the period between visual shape onset and the start of the response saccade.

Each trial began with the presentation of a small, white fixation point ($0.1 \times 0.1°$) near the display center. The animal was required to bring and hold its gaze within 0.4° of the point. The fixation point was extinguished 300 ms after acquisition, and one of the 2 shapes was randomly chosen and immediately presented directly at the current center of gaze (Fig. 3*B*). Because position variability on the retina can introduce neuronal response variability (DiCarlo and Maunsell 2003; Gur and Snodderly 1987), the stimulus position was always specified relative to the animal's center of gaze at the end of the fixation period.

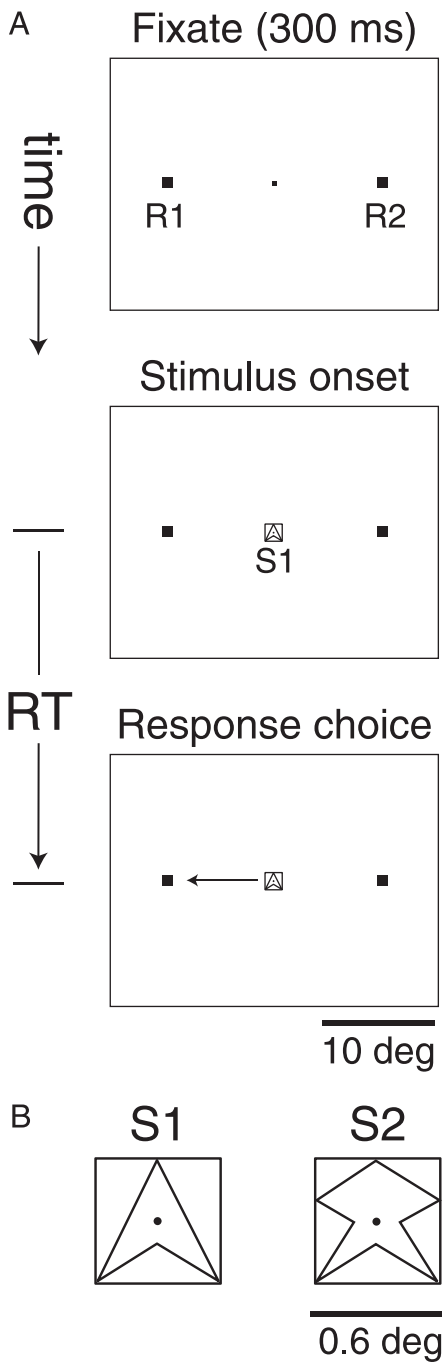

FIG. 3. Behavioral task performed by the monkeys. *A*: temporal sequence of events on each behavioral trial. Each rectangle is a schematic illustration of the visual display. Two response targets (*left* = R1; *right* = R2) were always illuminated at the same locations (10° eccentricity). After a fixation period, one of 2 possible stimuli (S1 or S2, *B*) was presented at the center of gaze. Animal was required to identify the stimulus by making an eye movement (saccade) to the appropriate (previously learned) response target (stimulus S1 requires response R1; stimulus S2 requires response R2). Animal was free to make its response saccade as rapidly as it liked, and the time between stimulus onset and the start of the response saccade was defined as the RT on that trial. *B*: schematics of the 2 stimuli used with one of the monkeys.

After a target shape appeared, the animal was allowed to respond as rapidly as it liked. Trials in which the animal failed to make a direct saccade to one of the 2 response locations were excluded from the analyses (<1% of trials). The visual shape was extinguished immediately when the animal's gaze left the fixation window.

Visual stimuli were presented on a video monitor (37.5 × 28.1 cm, 75 Hz, 1,600 × 1,200 pixels) positioned 62 cm from the monkey so that the display subtended ±17° (h) and ±13° (v) of visual angle. Both animals worked with a fixed set of 2 achromatic shapes (Fig. 3*B*). Each shape was constructed by connecting line segments (0.02° width) to form the stimulus outline (about 0.6° width). This outline shape was then convolved with a difference-of-Gaussians spatial filter (0.01° SD positive, 0.02° SD negative) so that the average luminance over each form was the same as the monitor background. The peak luminance of each stimulus was set to the monitor maximal white (46 cd/m^2).

Neuronal recording

Recordings were made from the left anterior inferotemporal cortex (AIT) and the right frontal eye field (FEF) in both animals. For AIT, a 23G guide tube was used to reach AIT from a dorsal approach. The superior temporal sulcus (STS) and the ventral surface were identified by comparing gray and white matter transitions and the depth of the skull base with atlas sections. Penetrations were made over an approximately 10 × 10-mm area of the ventral STS and ventral surface (Horsley–Clark AP: 10–20 mm, ML: 14–24 mm) of each animal. In both animals, the penetrations were concentrated near the center of this region, where form-selective neurons were more reliably found. Using electrolytic lesions and fluorescent dye (DiI, Molecular Probes) to coat the electrode (DiCarlo et al. 1996), we confirmed that the bulk of the AIT recordings from the first animal were on the ventral surface, centered about 10.5 mm posterior to the pole of the temporal lobe, about 3 mm lateral of the anterior middle temporal sulcus (AMTS). Based on the anterior–posterior coordinates, and the sulci, this region is approximately the anterior third of IT (AIT), and is contained in area TE (Felleman and Van Essen 1991; Logothetis and Pauls 1995; Logothetis and Sheinberg 1996).

The FEF chamber was targeted for the genu of the arcuate sulcus. A 23G guide tube was used to just penetrate the dura. The FEF was mapped in each animal using low-amplitude microstimulation to evoke saccades. Brief bursts of current were delivered through the recording electrode (50 μA, biphasic 200-μs pulses, cathode leading, 200 Hz, 200-ms duration, beginning 25 ms after a saccade) using an isolated stimulator (Bak Electronics). Consistent with previous reports (e.g., Bruce and Goldberg 1985; Schall 1997), such microstimulation could reliably produce saccades (50- to 100-ms latency from the first current pulse) at many locations along the anterior bank of the arcuate sulcus, and the saccade amplitude was largest for medial positions and smallest for lateral positions. We concentrated our recordings near the cortical region where low-amplitude microstimulation produced saccades of about 10° (the saccade amplitude required for the behavioral task). However, all neurons that were recorded along penetrations where low-amplitude microstimulation could reliably evoke a saccade were considered part of the FEF. Using electrolytic lesions, we confirmed that the bulk of the penetrations from the first animal were indeed through the anterior bank of the arcuate sulcus, near the genu.

In both AIT and FEF, single-unit recordings were made using glass-coated Pt/Ir electrodes (0.5–1.5 MΩ at 1 kHz) and spikes from individual neurons were amplified, filtered, and isolated using conventional equipment. The animal performed the task as the electrode was advanced into either AIT or FEF. Responses from every isolated neuron were assessed with an audio monitor and on-line histograms, and data were collected from even marginally responsive cells under the assumption that longer periods of observation might reveal statistically detectable effects. Because we sought to collect many trials of data from neurons that were modulated by the task (our goal was about 200 trials in each task condition), recordings were halted for neurons that did not show clear task modulation after 20–50 trials. Nevertheless, data from each recorded neuron were considered for further analysis if isolation was maintained for at least 10 presentations of each target form. In total, the responses of 63 AIT neurons (Monkey 1 = 25, Monkey 2 = 38) and 133 FEF neurons (Monkey 1 = 58, Monkey 2 = 75) were recorded. Most of the AIT neurons were located on the ventral surface (87%); the rest were in the ventral bank of the STS. Only neuronal responses collected during correctly completed behavioral trials were included in the analyses (about 90% of trials).

RESULTS

For the RT–NL technique to be useful, it must produce precise estimates of λ and β from data obtained in the limited period during which neurons are typically held for recording. To examine this, we plotted the reliability of the λ and β estimates from recorded neurons as a function of the number of trials obtained and the task-modulated firing rate of each neuron (Fig. 4). The results were similar to those predicted from simulations of Poisson spiking neurons: $\hat{\lambda}$ was usually more reliable than $\hat{\beta}$, and the reliability of both estimates depended strongly on the number of trials and the firing rate (see Fig. 4). However, Fig. 4 shows that reasonably reliable λ and β estimates could be obtained with moderate firing rate modulations (about 20 spikes/s) and about 150 behavioral trials (about 10 min for trials of the type used here; about 20 min if 2 trial types are interleaved, as done here).

To examine λ and β values across the population of recorded neurons, we focused on neurons where we had obtained reliable estimates of those parameters in the stimulus–response condition requiring a leftward response saccade (contralateral to recorded FEF; see Fig. 3). Neurons were considered to have reliable λ and β estimates if their SEs were <0.1 and 0.2, respectively. Sixty-eight of the 196 neurons recorded met both these criteria (AIT: 25 of 63; FEF: 43 of 133). This low yield is largely because many of the isolated neurons did not reveal significant responsivity after 20–50 trials of observation and recording of these was halted (see METHODS).

For neurons with reliable estimates, the values of λ and β produced by both the LSE and the MLE methods were highly correlated (correlation coefficients were 0.98 for λ and 0.97 for β) and the median absolute difference between the methods was small (0.01 for λ and 0.06 for β). For brevity, in the following sections we report the LSE-determined values. Where appropriate, the MLE-determined values are given in the figure captions for comparison.

Examples of mean NL and RT–NL covariance

An example of how the RT–NL technique was applied to recordings from a single neuron is shown in Fig. 5. These data were obtained from an AIT neuron while the animal performed the 2-choice visual-discrimination task. In this plot and others like it, only correctly completed trials from one stimulus–response condition are analyzed (leftward response saccade; see Fig. 3). The trials in the raster plot are sorted by reaction times, which are marked in heavy black. The mean RT across these trials was 334 ms (SD 63 ms). The plot shows that the neuron responded with a latency of about 125–150 ms, consistent with previous reports of IT neuronal latencies (e.g.,

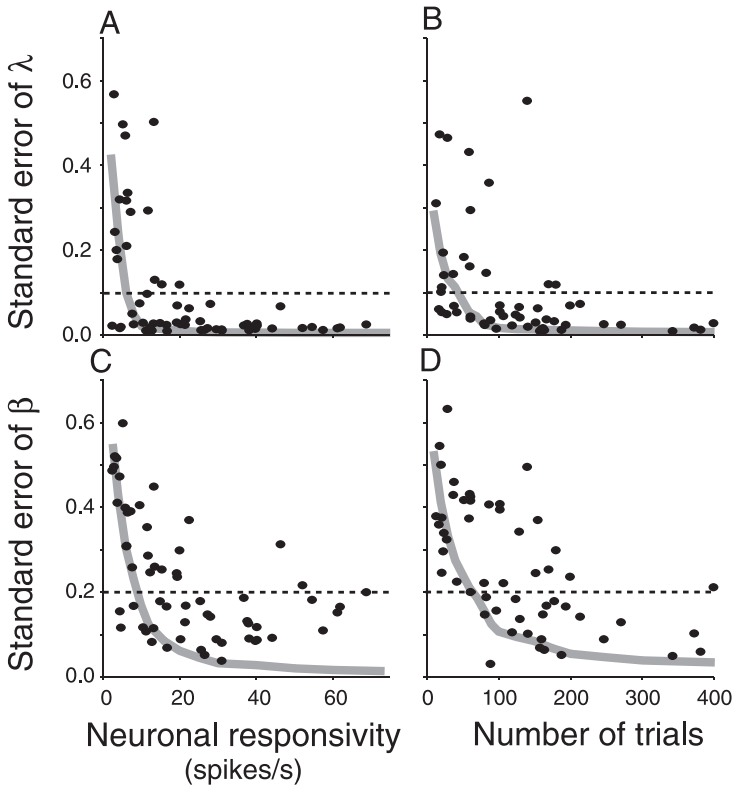

FIG. 4. Reliability of λ and β estimates. *A–D*: effect of neuronal responsivity (post-NL rate minus pre-NL rate; *A* and *C*) and number of trials (*B* and *D*) on the SEs of the λ and β estimates (determined by bootstrap; see METHODS). Each point shows the result from an individual recorded neuron. Only neurons with 100–300 trials are included in *A* and *C*. Only neurons with responsivities of 15–30 spikes/s are included in *B* and *D*. Solid lines show the reliability predicted from simulations of step changes in firing rate and Poisson spiking statistics. Dashed lines indicate the reliability levels used to include results in the final data set (e.g., Fig. 7; see text). For all simulations, λ and β were set to 0.5 and pre-NL rate was 5 spikes/s. In *A* and *C*, the number of simulated trials was 200. In *B* and *D*, simulated responsivity was 20 spikes/s. Results shown were obtained with the LSE method (see METHODS); the MLE method produced similar results (not shown).

Baylis et al. 1987; DiCarlo and Maunsell 2000; Vogels and Orban 1994).

To quantify the mean NL and the RT–NL covariation, we applied both the LSE and MSE techniques. The best fit for the rate change determined by the LSE is indicated by the gray curve in Fig. 5. This curve corresponds to a mean NL of 147 ms or a normalized mean NL (λ) of 0.44 (SE 0.01). The gray curve also corresponds to a normalized RT–NL covariation (β) of 0.19 (SE 0.06), a value somewhat greater than zero.

To confirm these estimates of λ and β, it is helpful to examine these data in another way. We divided the trials in Fig. 5*A* into 3 equal groups based on RT (slow, medium, and fast RT). These groups are indicated at the *right side* of the rasters. We then compiled response histograms for each group. The histograms in Fig. 5*B* were constructed with each trial aligned on the stimulus onset (vertical bar), yielding conventional poststimulus time histograms (PSTHs). The histograms in Fig. 5*C* were constructed with each trial aligned on a point halfway from stimulus onset to RT on that trial (vertical bar). The histograms in Fig. 5*D* were constructed with each trial aligned on the RT (vertical bar). The alignment on stimulus onset (Fig. 5*B*) provides the best overlap of the 3 histograms, consistent with a value of β that is near zero. Close inspection of this panel shows that the histograms from the intermediate (dashed) and slow (dotted) RT trials are offset slightly to the right of the histogram from the fast RT trials (solid). This suggests that the actual value of β is slightly greater than zero, in agreement with the small positive β value returned by the LSE and MLE methods.

Data collected from a FEF neuron are shown in Fig. 6, which has the same format as Fig. 5. The mean NL for this neuron is much longer than that for the AIT neuron, and the onset of the neuron's activity varies closely with RT. The LSE method yielded a best-fit rate change at the gray line, which corresponds to a λ value of 0.93 (SE 0.01) and a β value of 1.02 (SE 0.06). These values correspond to a NL that leads the RT on each trial by about 25 ms and are consistent with previously published reports describing saccade-linked activity in FEF (e.g., Bruce and Goldberg 1985; Hanes and Schall 1996). The β value near one is expected for pure motor neurons, and the λ value is smaller (longer saccade lead time) than that seen in primate abducens motor neurons (0.96; about 10-ms lead time; Sylvestre and Cullen 1999). The histograms at the bottom of Fig. 6 support the conclusion that the optimal alignment (i.e., the best-fit β) is close to one. The close overlap of the plots in Fig. 6*D* shows that the value of β returned by the LSE and MLE analyses arose from a change in the onset of the response, not a change in the magnitude of the response (see also Figs. 5*B*, 8*D*, and 9*C*).

Distributions of λ and β in different brain regions

Figure 7 shows the relationship between λ and β values for neurons in AIT (filled symbols) and FEF (open symbols). The distributions of λ and β values for AIT and FEF were not significantly different in the 2 monkeys (four 2-sample Kolmogorov–Smirnov tests, $P > 0.05$), with one exception: the distribution of λ values in FEF ($P = 0.02$). This difference was attributed to the presence of several more intermediate λ values (in the range 0.6–0.8) in Monkey 2 than in Monkey 1. For brevity, the data from both animals have been combined in Fig. 7 (Monkey 1: 11 AIT neurons, 14 FEF neurons; Monkey 2: 14 AIT neurons, 29 FEF neurons).

Figure 7 shows that most AIT neurons had λ values between 0.3 and 0.7 (median 0.50), corresponding to a NL of 85–195 ms (at the mean RT). This is consistent with previous reports

of AIT latency (e.g., Baylis et al. 1987; DiCarlo and Maunsell 2000; Vogels and Orban 1994). Most AIT neurons had β values near zero, with a median β of 0.06. Only 4 of the 25 AIT neurons (16%) had 95% confidence intervals for β that did not include zero (all >0). This indicates that a large fraction of task-modulated AIT neurons have neuronal latencies that covary much more closely to the time of stimulus onset than they do to the RT, even in the presence of substantial RT variability.

Figure 7 reveals that most FEF neurons fall into 2 clusters: *1*) neurons with λ and β values similar to those for AIT and *2*) neurons with λ and β values that are both close to one. The former group of neurons largely overlaps the population of AIT neurons. Within this group are FEF neurons with stimulus-related response latencies that are shorter than any of the AIT neurons in the sample (as short as about 50 ms). Short latency responses have previously been described in the FEF (e.g., Schall 1991; Schmolesky et al. 1998).

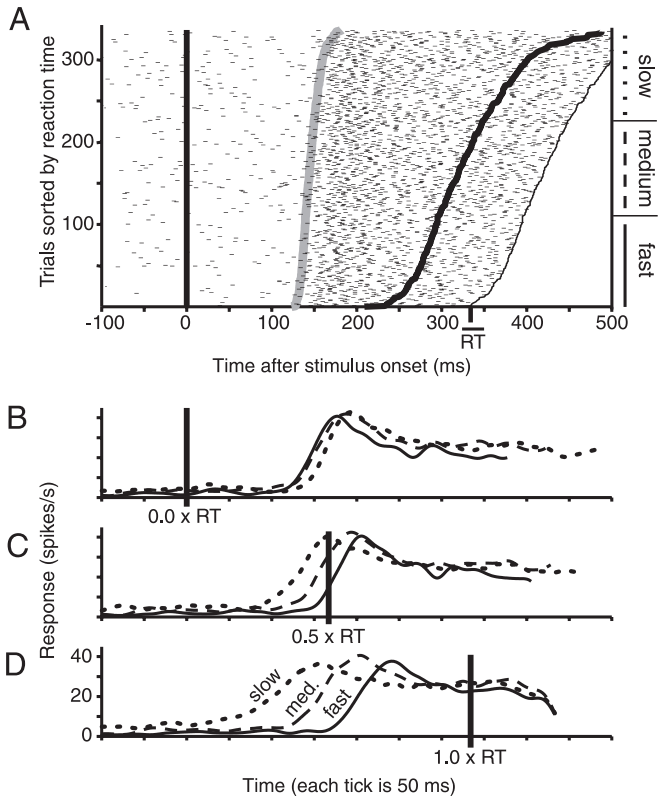

FIG. 5. Example of data and LSE fit for an anterior inferotemporal (AIT) neuron with a β value near 0. *A*: raster plot of the data collected during all trials in which visual stimulus S1 was presented and (correct) response R1 was given by the monkey. Each tick mark indicates the occurrence of an action potential. Abscissa is the time from stimulus onset (heavy black line). Trials are sorted by RT (heavy black curve). Because of the nonuniform distribution of RTs, the plot of RTs is a curve. Raster starts at 100 ms before stimulus onset, which was the earliest time considered in the LSE and MLE analyses. Thin black curve at the *right* of the plot marks 100 ms after RT on each trial, which is the latest time considered in the LSE and MLE analyses. Gray curve indicates the fit of the LSE algorithm to these data. MLE method found similar values for λ and β (0.42 $\pm$ 0.01 and 0.11 $\pm$ 0.05, respectively). *B–D*: each shows average histograms created from 3, equal-sized groups of trials divided based on RT (termed fast, medium, and slow RT trials and indicated with solid, dashed, and dotted histograms; see *A*). In *B* all trials were aligned at the stimulus onset before creating the 3 histograms. In *D* all trials were aligned at the RT. In *C* all trials were aligned at a time point halfway between the stimulus onset time and the RT. Each histogram was smoothed with a Gaussian filter (σ = 10 ms) and truncated 100 ms after the median RT of the trials on which it is based.

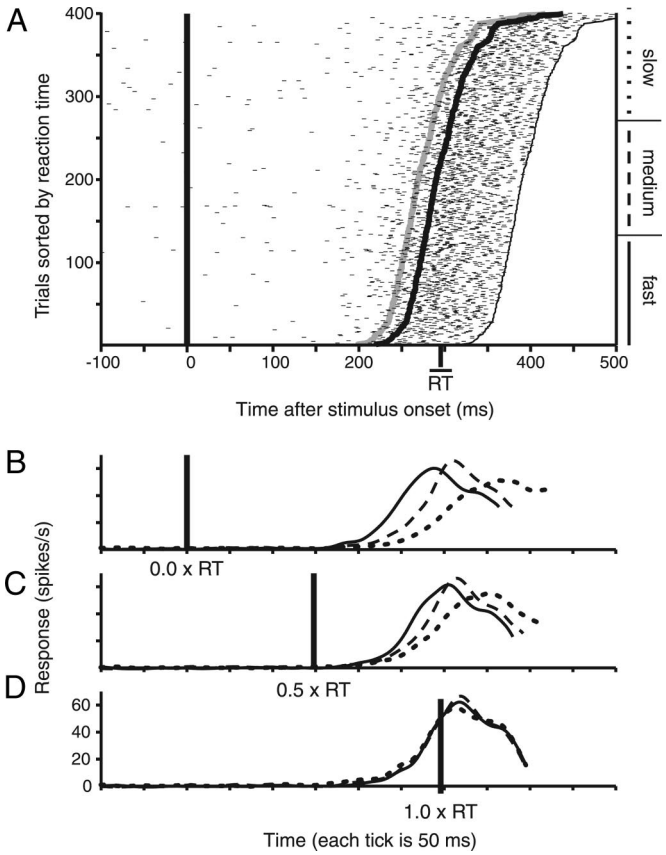

FIG. 6. Example of data and LSE fit for a frontal eye field (FEF) neuron with a β value near 1. Same format as in Fig. 5. MLE method produced values of λ and β for this neuron that were 0.82 $\pm$ 0.02 and 0.92 $\pm$ 0.14, respectively.

Among the FEF neurons with λ and β values near one, some had λ values slightly less than one (e.g., Fig. 6). These are the values expected for FEF "motor" neurons (e.g., Bruce and Goldberg 1985). However, most of the FEF neurons with β values near one had λ values that were *greater* than one. This indicates that these neurons became active *after* the start of the response saccade, usually before the saccade had ended. A typical FEF neuron with a λ value greater than one is shown in Fig. 8, which has the same format as Figs. 5 and 6. It became active (gray line) about 30 ms after the start of each saccade (heavy curve), regardless of the time of that saccade relative to stimulus onset. According to the LSE method, the λ value for this neuron was 1.11 (SE 0.01), and the β value was 1.03 (SE 0.05). Neurons like this one are not causal in producing the saccade because they are not modulated until after the saccade begins, but may carry signals related to the execution of the saccade, or possibly proprioceptive or visual signals generated by the movement. Although not quantified as we have done here, "postsaccadic" responses have been described previously in the FEF (Bruce and Goldberg 1985; Schall 1997).

Figure 7 also shows that a few FEF neurons have β values that are reliably intermediate between zero and one (e.g., 2 indicated near the arrow, Fig. 9). These also have λ values that are intermediate between the centers of the 2 main clusters. These intermediate values are consistent with neurons that lie at intermediate levels of a processing chain. Because neurons with intermediate β values may provide important clues about the functional organization of the processing chain, we show

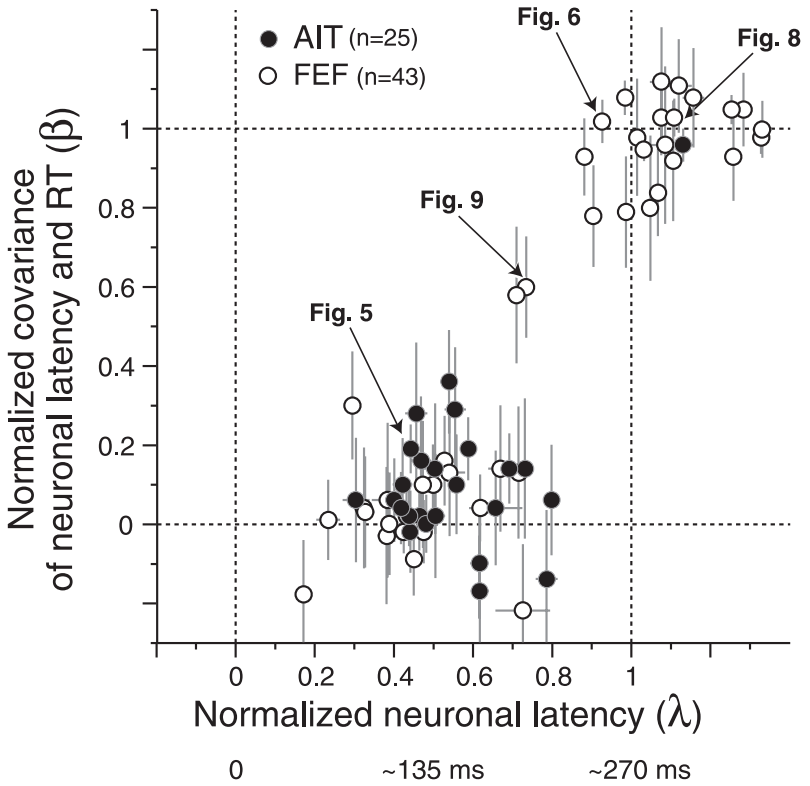

FIG. 7. Distribution of λ and β values observed in AIT and FEF. Normalized NL (λ) is plotted on the abscissa; normalized covariance of NL and RT (β) is plotted on the ordinate. Each point shows the λ and β estimates obtained from a single neuron where those estimates were determined to be reliable (see text). Error bars show 1 SE of the estimates for each neuron, determined by bootstrap (see METHODS). All data are from one sensory–motor condition (stimulus S1, response R1; see Fig. 3). By construction, the dashed vertical line at $\lambda = 1$ is the time of the start of the response saccade. End of the response saccade occurred at a time corresponding to a λ of about 1.2. Labels indicate example neurons shown in Figs. 5, 6, 8, and 9.

the activity of one of these neurons in Fig. 9 in the same format as other figures. The LSE provided a β estimate of 0.58 (SE 0.17) for this neuron. Consistent with this, the rasters show that the onset of the activity is slightly later on trials with long RT. By dividing the data into 3 RT groups, the *bottom panels* of Fig. 9 show that the histograms overlap best when the trials are aligned on a time point on each trial about halfway between stimulus onset and the RT (Fig. 9*C*). Alignment on either the time of stimulus onset (Fig. 9*B*) or on the RT (Fig. 9*D*) causes the histograms to misalign in the expected directions.

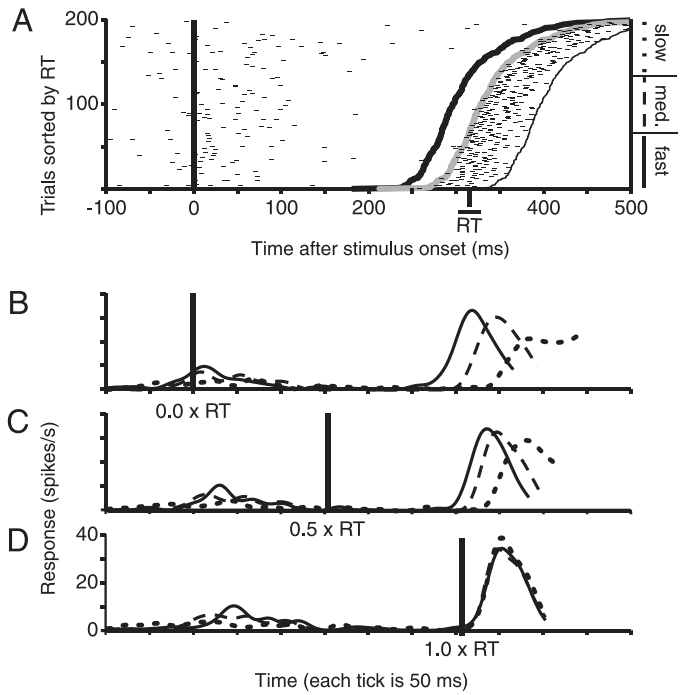

FIG. 8. Example of data and LSE fit for FEF neuron with a λ value >1. Same format as in Fig. 5. MLE method produced values of λ and β for this neuron that were 1.11 ± 0.01 and 1.03 ± 0.05, respectively.

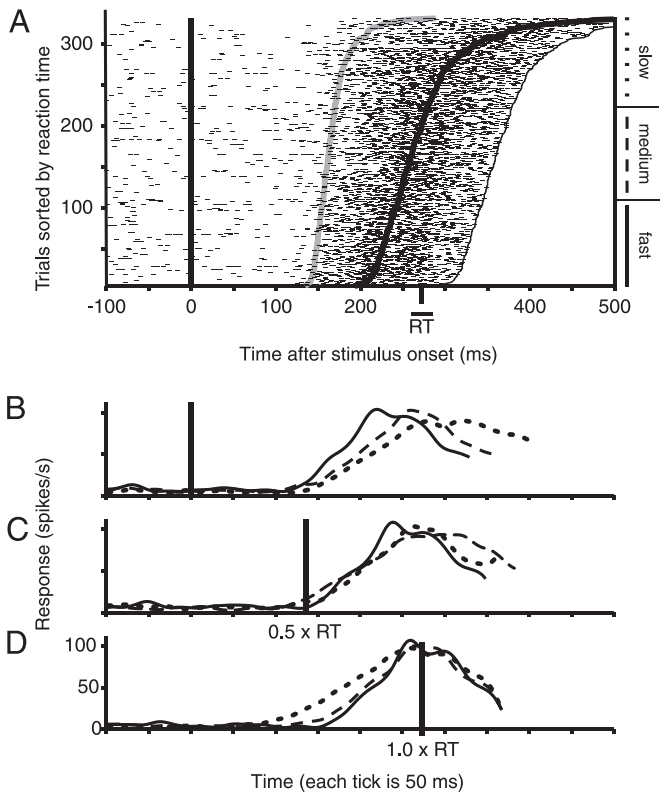

FIG. 9. Example of data and LSE fit for an FEF neuron with a β value intermediate between 0 and 1. Same format as in Fig. 5.

These intermediate β values are not an artifact of the LSE method because the MLE method also returned reliable, intermediate β values for both neurons (0.60, SE 0.18; 0.33, SE 0.07). Although hints of such neuronal responses patterns have been reported (Thompson and Schall 2000), we believe that this is the first quantitative demonstration of statistically reliable patterns of this sort. Both of these neurons were highly selective in that they were 3–4 times more modulated during the task requiring the leftward response saccade than in the task requiring the rightward response saccade (driven rates 68 spikes/s for leftward vs. 15 spikes/s for rightward saccades, and 21 spikes/s for leftward vs. 4 spikes/s for rightward saccades). However, there was no overall tendency in the population for selective neurons to have particular values of β (data not shown).

DISCUSSION

We have examined how the timing of neuronal responses can be used to infer relative positions of individual neurons on processing pathways involved in the initiation of specific behaviors. The neurophysiological results show that single-unit recordings yield reliable estimates of mean NL and RT–NL covariance (appropriately normalized and defined as λ and β). Collectively, this work suggests that the RT–NL technique is a useful tool for assessing neuronal pathways underlying specific behaviors.

λ and β of individual AIT and FEF neurons

To our knowledge, no previous report has examined the degree to which AIT response latencies covary on a trial-by-trial basis with motor responses. Although the values of λ we obtained for AIT neurons show that they become active near the middle of the stimulus–response interval, the small values of β we obtained for AIT neurons suggest that there is little motor component in the latency of AIT activity. Thus variance in RT in this task cannot be attributed to variance that accumulates in processing up to the level of AIT. If the activity of AIT neurons we recorded contributed to the initiation of saccades in this task, this result implies that most RT variance is generated in stages beyond AIT.

Although there is some evidence of saccade-locked activity in AIT when saccades are made in total darkness (Ringo et al. 1994), this activity is weak (about 0.5 spike/saccade) compared with stimulus-driven activity. A recent report (Eifuku et al. 2004) found that a few, very long latency AIT neurons had mean latencies that covaried with the mean RT of different tasks (i.e., longer mean NL for stimuli that resulted in longer mean RT). This suggests that λ (not β) for some AIT neurons varies from task to task. Because that study did not examine the trial-by-trial covariation of NL and RT, it is not possible to relate those results to the results on RT–NL covariance presented here.

The λ and β values of FEF neurons place most into 2 distinct groups. Previous studies of FEF have described distinct "visual" and "saccade-related" neurons in the FEF (Thompson and Schall 2000) and it is likely that the 2 groups in this study reflect that distinction. We were surprised to see so many FEF neurons with λ greater than one (Figs. 7 and 8), indicating that they became active after the response saccade had begun. However, many reports of saccade-related FEF activity define neurons as having perisaccadic or postsaccadic activity (see Schall and Thompson 1999) and show examples of activity that begins after the start of a saccade. Thus we do not believe there was anything unusual about the saccade-related activity we saw in the FEF.

Distribution of λ and β across AIT and FEF

In addition to confirming and quantifying existing observations, the distribution of values in a plot of λ versus β (Fig. 7) provides new clues about the role of AIT and FEF neurons in visuomotor tasks. Although the limited data presented here cannot provide definitive answers about the pathways supporting this task, they provide constraints and point to several nonexclusive explanations of how RT variance arises in this task. The accumulation of variance along the processing chain can be analytically determined under certain conditions. We computed the expected λ and β values under conditions where *1*) statistically independent random delays are added at each processing level, *2*) the NL for each level is the sum of all the preceding delays, and *3*) the RT is the sum of all delays along the entire processing chain.

Figure 10 shows 3 different plots of λ and β relationships that are consistent with the distribution of λ and β values obtained (Fig. 7). In each case, the filled points are λ and β values corresponding to neurons on the processing chain. Figure 10*A* illustrates a case in which β rises linearly with levels on the processing chain, with $\beta \cong \lambda$ throughout the chain (filled circles). This pattern would be expected if independent random delays with the same mean and variance were introduced at each processing level. The neurophysiological data from AIT and FEF (Fig. 7) almost all lie to the right of the line marked by the processing chain in Fig. 10*A*. Values lying to the right of the processing chain can occur if neurons are driven, directly or indirectly, by neurons on the processing chain, but do not themselves contribute to the behavioral response. The open circles in Fig. 10*A* represent chains of neurons that branch off the processing chain at different levels. Successive levels on these branches can be expected to have increasing values of λ because of the delay added at each processing stage. Because neurons on a branch do not contribute to the behavioral response, the NL variability that accumulates along the branch will not covary with RT, and β values for neurons on a branch will never rise above β at the level of the processing chain where the branch occurred. It is possible that only a tiny fraction of neurons in any area lie on the processing chain for a given task, and the odds of sampling these neurons are extremely small. The neurons we recorded might have been driven by neurons on the processing chain, but not contributed to the initiation of the response. In that case, each of the points with β values near 0 in Fig. 7 sits below and to the right of points corresponding to neurons on the processing chain, which were never sampled and thus do not appear in the plot.

Figure 10*B* illustrates an alternative case in which β values rise rapidly from near zero to near one at an intermediate value of λ. This would imply that most of the subject's RT variance arises in a few central stages, such that the AIT and FEF neurons with low β values precede a highly variable central stage, and the FEF neurons with high β values follow it. In this

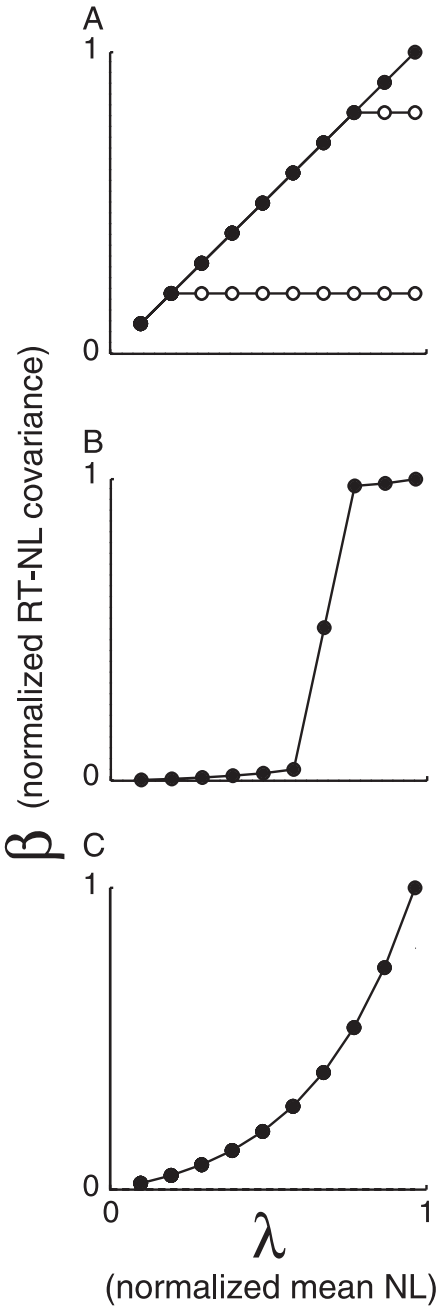

FIG. 10. Hypothetic distributions of λ and β. In each plot the filled circles represent values corresponding to neurons lying on the processing chain that generate a behavioral response. Connecting lines between points indicate anatomical pathways between groups of neurons, with activity propagating from *left* to *right* (increasing values of λ). See text for details.

case, it would be appropriate to consider most of these neurons as either predominantly sensory or predominantly motor. If most variance accumulated in only a few stages, it would also explain why so few intermediate β values were encountered.

Finally, Fig. 10*C* illustrates a pattern in which β remains small over most of the range of λ values, and then rises in an accelerating way as λ approaches one. This pattern could come about in 2 ways. First, it would occur if the random delays added in later stages were consistently more variable than those in earlier stages. Second, it would occur if random delays with the same variance were added at each level, but converging inputs at all levels acted to reduce response variance. If each level in the processing chain consists of many neurons that project in a convergent way to neurons in the next level, neurons in the later level can "average out" the variability that exists across the neurons in the earlier level, providing the variability in different neurons is statistically independent (Marsalek et al. 1997). The effect would be to compress all β values toward zero, with β for neurons in earlier levels more affected because the variance they contribute would be repeatedly reduced in successive stages, yielding little covariance with RT. β would rise exponentially in later stages because the variance those stages contribute to RT would be attenuated by fewer stages intervening before the behavioral response. The steepness of the rise would depend on what fraction of the variance was eliminated by convergence. If there is an accelerating pattern of β values, few intermediate values of β might have been encountered in this study because β goes from low to high values over only a few levels of processing. The presence of a few AIT neurons with β values reliably greater than zero (Fig. 7) is consistent with this scenario.

The distribution in Fig. 7 might represent any of the possibilities shown in Fig. 10, or a combination of these and other effects. Distinguishing among possible explanations for the distribution of λ and β values observed will require recording from many neurons in many different structures, but it is a step forward to have an approach that may make it possible to resolve them. In this regard, the observation of FEF neurons with intermediate RT–NL covariance values (i.e., β values significantly greater than zero and significantly less than one) is particularly relevant. Although they were a small fraction of the neurons we recorded, they demonstrate that it is possible to find neurons with physiological signatures consistent with intermediate positions on a processing chain.

Related work

Correlating neuronal activity and behavioral responses is a well-established tool. Over the last few decades dozens of studies have shown that average neuronal thresholds correlate with perceptual thresholds (e.g., Connor and Johnson 1992; Johnson 1974; Mountcastle et al. 1969; Parker and Newsome 1998; Shadlen et al. 1996; Tolhurst et al. 1983). More recently, neurophysiological recordings from behaving animal subjects have revealed correlations between the *magnitude* of the responses of sensory neurons and perceptual reports on a trial-by-trial basis (e.g., Britten et al. 1996; Dodd et al. 2001; Thiele et al. 1999; Zhang et al. 1997a,b). The goal of this study was to extend this approach by developing methods of examining correlations between the *latency* of neuronal responses (NL) and RT on a trial-by-trial basis.

Many previous studies have used mean latency as a tool for exploring the functional relationships between neurons or neuronal structures. Some have compared NL and RT in the responses combined across many trials or many neurons (Cook and Maunsell 2002; Eifuku et al. 2004; Schmied et al. 1979; Thompson and Schall 2000), but the covariance between the latency of individual neurons and behavioral RT has received little attention. Lamarre and colleagues (Lamarre and Chapman 1986; Lamarre et al. 1983; Spidalieri et al. 1983) examined the trial-by-trial relationship between NL and RT in a manner equivalent to that described here. However, they interpreted β

in a binary way, with values near zero assigned as sensory activity and values near one assigned as motor. Lamarre and Chapman (1986) reported that a small number of neurons with movement-related activity in the dentate nucleus had properties consistent with β values intermediate to zero and one. However, they said that their data could not be used to decide whether this finding had true functional significance, and excluded those neurons from their analysis. Other studies (e.g., Ageranioti-Belanger and Chapman 1992; Berthier and Moore 1990; Berthier et al. 1991; Chapman and Ageranioti-Belanger 1991; Ghez et al. 1983) examined trial-by-trial covariance between NL and RT but limited their interpretation to assigning the activity as more closely related to either sensory stimulation or to motor response.

Overall, trial-by-trial covariance between NL and RT has received remarkably little attention. A contributing factor may have been a report by Commenges and Seal (1986) that suggested the approach had little merit. They provided a mathematical analysis that led them to conclude that the slope of the trial-by-trial relationship between NL and RT was a poor approach to evaluating neurophysiological processing levels. Although we agree with their analysis, we believe that their conclusion was too sweeping because it was based on one particular analysis. Specifically, they showed that a linear regression of RT on NL will typically yield a slope close to one, regardless of whether a neuron is early or late in a processing chain. The regression they considered is equivalent to fitting regression lines to the data points in Fig. 1*C* after flipping the axes. In that configuration, the slope of the best fitting line is equal to the covariance between NL and RT normalized by the variance of NL. Because the variance of NL is small for neurons with little covariance between NL and RT ("sensory" neurons) and large for neurons with a lot of covariance between NL and RT ("motor" neurons), this normalization results in no appreciable difference in the slope of the best-fitting line for neurons at different levels in the processing chain. In our analysis, the covariance between NL and RT is normalized by a constant (the variance of RT), so this problem is avoided.

The RT–NL technique may be complementary or perhaps synergistic with other measures of neuronal and behavioral response. Several previous studies have examined the trial-by-trial covariance between the *magnitude* of response and behavioral *choice* (called "choice probability" or "detect probability"; e.g., Britten et al. 1996; Dodd et al. 2001; Thiele et al. 1999). Response magnitude and response latency are different measures of neuronal response, and they are not always correlated (e.g., Richmond and Optican 1987; Richmond et al. 1990). Similarly, behavioral choice and behavioral latency (RT) are different measures of behavior that are not perfectly correlated. Examination of the association of each behavioral measure with each neuronal measure might provide useful information about the contribution of neurons to behavior. For example, it has been shown that measures of the trial-by-trial association of behavioral choice and the magnitude of neuronal activity can be used to place neurons along a sensory-motor continuum (Zhang et al. 1997a,b). However, the examination of the association of behavioral latency (RT) and neuronal latency (i.e., the NL–RT technique described here) has some potentially important advantages. First, although the covariation of NL and RT can be expected to monotonically increase along most processing chains, the covariation of neuronal response magnitude and behavioral choice (e.g., choice–probability) is less certain to exhibit such a monotonic increase (although it can; Williams et al. 2003). In particular, if signals in the processing pathway undergo substantial convergence or divergence between processing levels, the choice–probability may either rise or fall along the pathway. Tasks could be contrived so that no individual motor neuron provided a reliable indication of the behavioral response. Although convergence and other factors can affect the rate at which covariance between NL and RT increases along the processing chain, it can be expected to rise largely monotonically. Second, clean application of choice-related measures requires the subject to be working near sensory discrimination thresholds (so that different choices are made on identical stimuli), whereas the covariance between NL and RT can be measured in any RT task, even if behavior is well above threshold (as in the task reported here). Third, unlike choice-related measures, the NL–RT technique can be applied even if the animal is not required to make a choice, but simply respond to a stimulus.

In conclusion, the RT–NL technique described here has several attractive features. First, it does not entail assignments of responses into "sensory," "motor," or other categories. Instead, it defines a quantitative continuum of processing, which is a less constrained, possibly more accurate, approach to the organization of the CNS. Second, although accurate measurements of second-order statistics like covariance often require a large amount of data, data collection for measurement of NL–RT covariance is feasible because trials in highly practiced choice reaction time tasks are short, and the analysis technique presented here produces unbiased, reliable estimates in a reasonable number of trials. Finally, it has the potential to identify the sources of variability that underlie RT variability for any RT task.

Although the approach has attractive features, it also has limitations. The values of λ and β for each neuron are defined for a specific stimulus–response pairing. Because neurons will likely contribute differently to different behaviors, they can be expected to have different values of λ and β during different tasks. Indeed, the activity of a given neuron will not be modulated during many tasks (for which λ and β would be undefined). This is a characteristic of brain organization rather than a flaw in the method, but it does mean that, with standard single-unit methods, useful comparisons between neurons will typically require behavioral performance that is stable over days or weeks. Fortunately, relatively stable behavior is not difficult to achieve for choice reaction time tasks.

It should also be recognized that the NL–RT technique cannot address some types of neuronal control of behavior. The concept of a processing chain on which the current work is based may be useful only for well-defined behaviors with clear onsets, such as choice reaction time tasks (see Smith 1968), in which a small, fixed set of sensory stimuli is associated in a fixed way with a small set of behavioral responses. It is not clear how representative such reaction time tasks are of everyday behaviors. Furthermore, the NL–RT technique reveals only some aspects of the circuits underlying behavior in choice reaction time tasks. Because it measures only the onset of neuronal and behavioral responses, it does not address circuits that modulate a behavioral response as it unfolds.

Although we believe that the current data demonstrate the potential of the NL–RT technique, further refinement could make it more powerful. For example, it would be valuable to develop methods to measure trial-by-trial variance in NL that does not covary with RT (i.e., NL variance that is not simply explained by RT variance). This variance appears in Fig. 1*C* as the scatter of the points around the best-fitting line. This variance may change in predictable ways across levels and in ways that will further constrain the position of neurons within the processing chains.

The methods describe here assess covariance using an approach that assumes a single, instantaneous change in neuronal activity. Measurements that use a response profile match to each neuron's response, rather than a step, would probably yield results that are more precise. It might also be possible to examine whether a neuron's response profile covaries with RT. Superior performance might obtain from more general methods that capture the trial-by-trial covariance of all aspects of a neuron's response, incorporating the latency, magnitude, and response profile. Although we did not encounter neurons with activity that combines distinct sensory and motor components, such neurons have been described (e.g., Sparks et al. 2000). The NL–RT technique might be extended to allow for multiple changes in activity during trials, as has been done for detecting multiple changes in activity within individual trials (Commenges et al. 1986b). An approach that captures all these aspects of neuronal responses is likely to be more powerful than any in isolation.

The work presented represents a first step in developing a new approach for evaluating the functional relationships of individual neurons in the context of particular behavioral tasks. We believe that the current data show that it can be used to dissect functional relationships among brain structures that contribute to specific behaviors and provides a tool that complements other approaches (e.g., Britten et al. 1996; Zhang et al. 1997a).

ACKNOWLEDGMENTS

We thank E. Cook for helpful suggestions and assistance with the simulations prepared for this report and W. Bosking, D. Cox, G. Ghose, C. Hung, J. Schall, and S. Shorter-Jacobi for helpful comments on earlier version of this manuscript. We also thank C. Hocker, D. Murray, and T. Williford for excellent technical assistance.

GRANTS

This work was supported by National Eye Institute Grant R01EY-05911. J.H.R. Maunsell is an Investigator with the Howard Hughes Medical Institute.

REFERENCES

Ageranioti-Belanger SA, and Chapman CE. Discharge properties of neurones in the hand area of primary somatosensory cortex in monkeys in relation to the performance of an active tactile discrimination task. II. Area 2 as compared to areas 3b and 1. *Exp Brain Res* 91: 207–228, 1992.

Baylis GC, Rolls ET, and Leonard CM. Functional subdivisions of the temporal lobe neocortex. *J Neurosci* 7: 330–342, 1987.

Berthier NE, Barto AG, and Moore JW. Linear systems analysis of the relationship between firing of deep cerebellar neurons and the classically conditioned nictitating membrane response in rabbits. *Biol Cybern* 65: 99–105, 1991.

Berthier NE and Moore JW. Activity of deep cerebellar nuclear cells during classical conditioning of nictitating membrane extension in rabbits. *Exp Brain Res* 83: 44–54, 1990.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, and Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13: 87–100, 1996.

Bruce CJ and Goldberg ME. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53: 603–635, 1985.

Bullier J and Henry GH. Neural path taken by afferent streams in striate cortex of the cat. *J Neurophysiol* 42: 1264–1270, 1979.

Chapman CE and Ageranioti-Belanger SA. Discharge properties of neurones in the hand area of primary somatosensory cortex in monkeys in relation to the performance of an active tactile discrimination task. I. Areas 3b and 1. *Exp Brain Res* 87: 319–339, 1991.

Cinlar E. *Introduction to Stochastic Processes.* Englewood Cliffs, NJ: Prentice-Hall, 1975.

Commenges D, Pinatel F, and Seal J. A program for analysing single neuron activity by methods based on estimation of a change-point. *Comput Methods Programs Biomed* 23: 123–132, 1986a.

Commenges D and Seal J. The analysis of neuronal discharge sequences: change-point estimation and comparison of variances. *Stat Med* 4: 91–104, 1985.

Commenges D and Seal J. The formulae-relating slopes, correlation coefficients and variance ratios used to determine stimulus-or movement-related neuronal activity. *Brain Res* 383: 350–352, 1986.

Commenges D, Seal J, and Pinatel F. Inference about a change in point experimental neurophysiology. *Math Biosciences* 80: 81–108, 1986b.

Connor CE and Johnson KO. Neural coding of tactile texture: comparisons of spatial and temporal mechanisms for roughness perception. *J Neurosci* 12: 3414–3426, 1992.

Cook EP and Maunsell JHR. Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci* 5: 985–994, 2002.

DiCarlo JJ, Lane JW, Hsiao SS, and Johnson KO. Marking microelectrode penetrations with fluorescent dyes. 64: 75–81, 1996.

DiCarlo JJ and Maunsell JHR. Form representation in monkey inferotemporal cortex is virtually unaltered by free viewing. *Nat Neurosci* 3: 814–821, 2000.

DiCarlo JJ and Maunsell JHR. Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *J Neurophysiol* 89: 3264–3278, 2003.

Dodd JV, Krug K, Cumming BG, and Parker AJ. Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J Neurosci* 21: 4809–4821, 2001.

Efron B and Tibshirani R. *An Introduction to the Bootstrap.* Boca Raton, FL: CRC Press, 1998.

Eifuku S, De Souza WC, Tamura R, Nishijo H, and Ono T. Neuronal correlates of face identification in the monkey anterior temporal cortical areas. *J Neurophysiol* 91: 358–371, 2004.

Felleman DJ and Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1: 1–47, 1991.

Gawne TJ, Kjaer TW, Hertz JA, and Richmond BJ. Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cereb Cortex* 6: 482–489, 1996.

Ghez C, Vicario DS, Martin JH, and Yumiya H. Sensory motor processing of target movements in motor cortex. In: *Motor Control Mechanisms in Health and Disease*, edited by Desmedt JE. New York: Raven Press, 1983, p. 61–92.

Gur M and Snodderly DM. Studying striate cortex neurons in behaving monkeys: benefits of image stabilization. *Vision Res* 27: 2081–2087, 1987.

Hanes DP and Schall JD. Neural control of voluntary movement initiation. *Science* 274: 427–430, 1996.

Johnson KO. Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. *J Neurophysiol* 37: 48–72, 1974.

Lamarre Y, Busby L, and Spidalieri G. Fast ballistic arm movements triggered by visual, auditory, and somesthetic stimuli in the monkey. I. Activity of precentral cortical neurons. *J Neurophysiol* 50: 1343–1358, 1983.

Lamarre Y and Chapman C. Comparative timing of neuronal discharge in cortical and cerebellar structures during a simple arm movement in the monkey. *Exp Brain Res Ser* 15: 14–27, 1986.

Logothetis NK and Pauls JP. Psychophysical and physiological evidence for viewer-centered object representation in the primate. *Cereb Cortex* 5: 270–288, 1995.

Logothetis NK and Sheinberg DL. Visual object recognition. *Ann Rev Neurosci* 19: 577–621, 1996.

Marsalek P, Koch C, and Maunsell J. On the relationship between synaptic input and spike output jitter in individual neurons. *Proc Natl Acad Sci USA* 94: 735–740, 1997.

Maunsell JHR and Gibson JR. Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol* 68: 1332–1344, 1992.

Mountcastle VB, Talbot WH, Sakata H, and Hyvarinen J. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *J Neurophysiol* 32: 452–484, 1969.

Nowak LG and Bullier J. The timing of information transfer in the visual system. In: *Cerebral Cortex: Extrastriate Cortex in Primate*, edited by Rockland K, Kaas JH, and Peters A. New York: Plenum Press, 1997, p. 870.

Parker AJ and Newsome WT. Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21: 227–277, 1998.

Richmond BJ and Optican LM. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. II. Quantification of response waveform. *J Neurophysiol* 57: 147–161, 1987.

Richmond BJ, Optican LM, and Spitzer H. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus–response relations. *J Neurophysiol* 64: 351–369, 1990.

Ringo JL, Sobotka S, Diltz MD, and Bunce CM. Eye movements modulate activity in hippocampal, parahippocampal, and inferotemporal neurons. *J Neurophysiol* 71: 1285–1288, 1994.

Robinson DA. A method of measuring eye movements using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* 101: 131–145, 1963.

Robinson DL and Rugg MD. Latencies of visually responsive neurons in various regions of the rhesus monkey brain and their relation to human visual responses. *Biol Psychol* 26: 111–116, 1988.

Schall JD. Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *J Neurophysiol* 66: 559–579, 1991.

Schall JD. Visuomotor areas of the frontal lobe. In: *Cerebral Cortex: Extrastriate Cortex in Primate*, edited by Rockland K, Kaas JH, and Peters A. New York: Plenum Press, 1997.

Schall JD and Thompson KG. Neural selection and control of visually guided eye movements. *Annu Rev Neurosci* 22: 241–259, 1999.

Schmied A, Benita M, Conde H, and Dormont JF. Activity of ventrolateral thalamic neurons in relation to a simple reaction time task in the cat. *Exp Brain Res* 36: 285–300, 1979.

Schmolesky MT, Wang Y, Hanes DP, Thompson KG, Leutgeb S, Schall JD, and Leventhal AG. Signal timing across the macaque visual system. *J Neurophysiol* 79: 3272–3278, 1998.

Schroeder CE, Mehta AD, and Givre SJ. A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cereb Cortex* 8: 575–592, 1998.

Seal J, Commenges D, Salamon R, and Bioulac B. A statistical method for the estimation of neuronal response latency and its functional interpretation. *Brain Res* 278: 382–386, 1983.

Shadlen MN, Britten KH, Newsome WT, and Movshon JA. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16: 1486–1510, 1996.

Smith EE. Choice reaction time: an analysis of the major theoretical positions. *Psychol Bull* 69: 77–110, 1968.

Sparks D, Rohrer WH, and Zhang Y. The role of the superior colliculus in saccade initiation: a study of express saccades and the gap effect. *Vision Res* 40: 2763–2777, 2000.

Spidalieri G, Busby L, and Lamarre Y. Fast ballistic arm movements triggered by visual, auditory and somesthetic stimuli in the monkey. II. Effects of unilateral dentate lesion on discharge of precentral cortical neurons and reaction time. *J Neurophysiol* 50: 1359–1379, 1983.

Sylvestre PA and Cullen KE. Quantitative analysis of abducens neuron discharge dynamics during saccadic and slow eye movements. *J Neurophysiol* 82: 2612–2632, 1999.

Thiele A, Distler C, and Hoffmann KP. Decision-related activity in the macaque dorsal visual pathway. *Eur J Neurosci* 11: 2044–2058, 1999.

Thompson KG and Schall JD. Antecedents and correlates of visual detection and awareness in macaque prefrontal cortex. *Vision Res* 40: 1523–1538, 2000.

Tolhurst DJ, Movshon JA, and Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775–785, 1983.

Vogels R and Orban GA. Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. *J Neurophysiol* 71: 1428–1451, 1994.

Williams ZM, Elfar JC, Eskandar EN, Toth LJ, and Assad JA. Parietal activity and the perceived direction of ambiguous apparent motion. *Nat Neurosci* 6: 616–623, 2003.

Zeki SM. Localization and globalization in conscious vision. *Ann Rev Neurosci* 24: 57–86, 2001.

Zhang J, Riehle A, and Requin J. Analyzing neuronal processing locus in stimulus–response association tasks. *J Math Psychol* 41: 219–236, 1997a.

Zhang J, Riehle A, Requin J, and Kornblum S. Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. *J Neurosci* 17: 2227–2246, 1997b.