

The Role of Spike Timing in the Coding of Stimulus Location in Rat Somatosensory Cortex

Stefano Panzeri,* Rasmus S. Petersen,†
Simon R. Schultz,‡ Michael Lebedev,§
and Mathew E. Diamond†||

*Neural Systems Group
Department of Psychology
Ridley Building
University of Newcastle upon Tyne
NE1 7RU
United Kingdom

†Cognitive Neuroscience Sector
International School for Advanced Studies
Via Beirut 2/4
34014 Trieste
Italy

‡Howard Hughes Medical Institute
Center for Neural Science
New York University
4 Washington Place
New York, New York 10003

§Laboratory of Systems Neuroscience
National Institutes of Health Animal Center
P.O. Box 608
Poolesville, Maryland 20837

Summary

Although the timing of single spikes is known to code for time-varying features of a sensory stimulus, it remains unclear whether time is also exploited in the neuronal coding of the spatial structure of the environment, where nontemporal stimulus features are fundamental. This report demonstrates that, in the whisker representation of rat cortex, precise spike timing of single neurons increases the information transmitted about stimulus location by 44%, compared to that transmitted only by the total number of spikes. Crucial to this code is the timing of the first spike after whisker movement. Complex, single neuron spike patterns play a smaller, synergistic role. Timing permits very few spikes to transmit high quantities of information about a behaviorally significant, spatial stimulus.

Introduction

The most common approach to analyzing how neurons represent peripheral stimuli has been to compare the number of spikes emitted during a selected epoch after the presentation of many different stimuli. Total spike counts typically vary across a stimulus set, such that spike count coding affords some degree of discriminability concerning which stimulus has occurred (Adrian, 1926; Werner and Mountcastle, 1965; Tolhurst et al., 1983; Tolhurst, 1989; Britten et al., 1992; Tovee et al., 1993; Petersen and Diamond, 2000). On the other hand, recent studies have revealed that the timing of individual

spikes—once assumed to be too variable to support robust computation—can represent the time structure of rapidly varying stimuli, such as movement within a visual scene, with remarkable accuracy (Bialek et al., 1991; Rieke et al., 1997; Ruyter van Steveninck et al., 1997; Borst and Theunissen, 1999; Ahissar et al., 2000).

Although it should not be surprising that the timing of individual spikes is important for coding a stimulus defined by its temporal structure, it would be more remarkable if spike timing were also to code spatial aspects of stimuli. Spatial features, stationary on a short time scale, constitute an essential part of the surrounding world. The present paper asks whether spike timing in the rat somatosensory cortex contributes to the coding of a nontemporal property, the location of a stimulus applied to a single whisker.

Neurons in the rat primary somatosensory cortex typically have multiwhisker receptive fields, one principal whisker evoking the strongest response and other whiskers evoking weaker responses (Welker, 1971; Simons, 1978; Armstrong-James and Fox, 1987). Principal whisker stimulation tends to evoke spikes a few milliseconds earlier than does surround whisker stimulation (Armstrong-James and Fox, 1987; Armstrong-James et al., 1992). From these well-documented facts, intuition dictates that spike count and spike timing both carry information about stimulus location. However, their relative importance is unknown and can be determined only by a quantification of the information they contribute. Moreover, intuition is of no value in speculating about the information contained in single spikes as compared to that in complex temporal spike patterns. In this paper, we show that many cells contain information in spike timing beyond that present in the total spike count. The response variable carrying most of this extra information is the time of the *first* spike following each whisker deflection; patterns within each spike train play a less important role. Thus, the timing of single spikes in the rat somatosensory cortex contributes to the coding of a spatial, behaviorally relevant feature of peripheral events.

Results

Distinguishing between “Spike Count” and “Spike Timing” Codes

We analyzed 106 single cells recorded from barrel-column D_2 in somatosensory cortex of urethane-anesthetized rats. We stimulated vibrissae C_{1-3} , D_{1-3} , and E_{1-3} one at a time to permit an analysis of how single cells encode stimulus location. The simplest hypothesis is that stimulus site is encoded purely by the total number of spikes fired in some poststimulus time window that is long compared to the typical interspike interval—spike count coding. This idea was tested by determining whether additional information was transmitted by spike time coding—either fast modulation in firing rate (Bialek et al., 1991) or repeated occurrence of precise sequences of spikes (Abeles et al., 1993).

|| To whom correspondence should be addressed (e-mail: diamond@sissa.it).

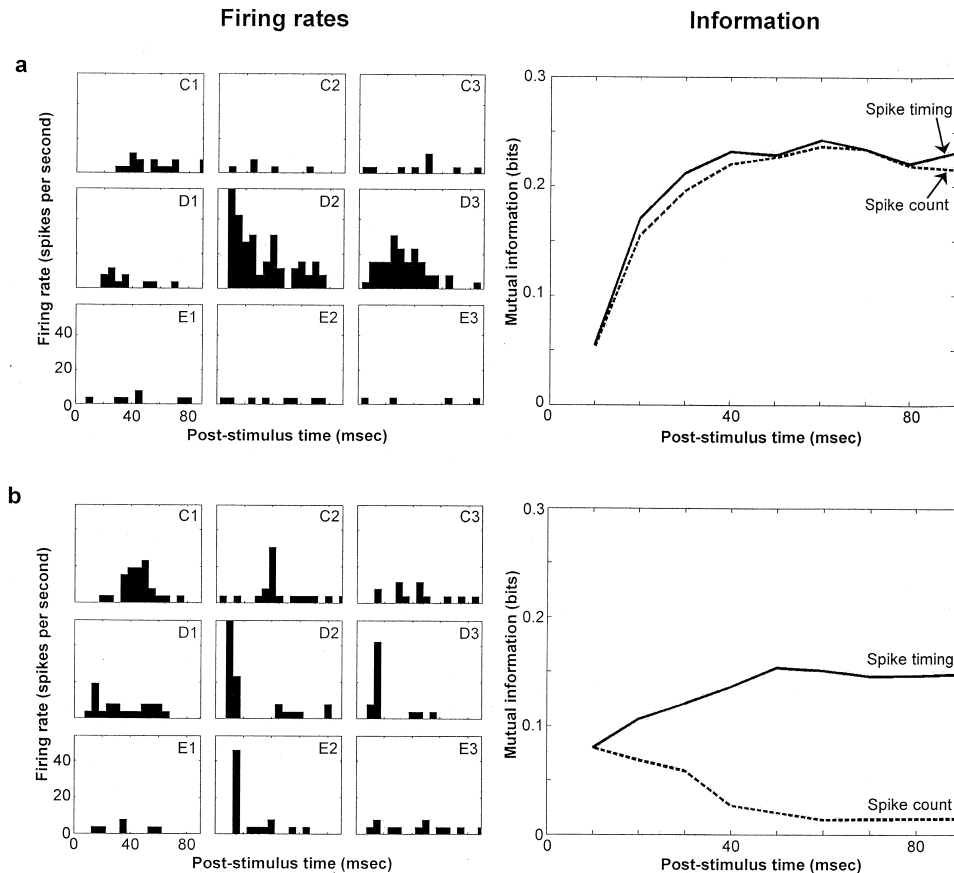


Figure 1. Coding by Spike Timing versus Spike Count

(A) Example of a neuron where all the information was in the spike count. PSTHs for each of the nine whisker stimuli are shown on the left. The mutual information between the stimulus set and neuronal response is shown on the right as a function of time after stimulus onset. Information in spike timing (solid line) and spike count (dotted line) were computed using the corresponding series expansions described in Experimental Procedures.

(B) A neuron for which there was significant information in spike timing. PSTHs and mutual information plotted as for (A).

We used an information theoretic method (Panzeri et al., 1999; Panzeri and Schultz, 2000) that separates the contributions from the different types of neural coding (see Experimental Procedures). The richness of possible spike patterns (and thus the complexity of possible neural codes) increases rapidly with the number of spikes emitted on each stimulus trial. Conversely, low firing rates limit the potential coding complexity. Typical spike counts for barrel cortical neurons are just 0–3 spikes/whisker deflection. This allows the information to be well approximated by a series expansion that depends only on the time-varying firing rate (conventionally illustrated as the poststimulus time histogram, PSTH) and the pairwise correlation between spike times (Panzeri et al., 1999; Panzeri and Schultz, 2000).

We computed both the information contained in the spike counts and that in the spike times with the series expansion approach. All analyses were based on a poststimulus time window 0– T . Results were computed as a function of the window size T and thus reveal the information accumulated from time 0 to time T . For spike count coding, the “response” on each trial was simply the number of spikes occurring in the time window. For spike time coding, the window was broken into bins of

size Δt and the response was a binary “word” of length $T/\Delta t$. The overall contribution of temporal coding was evaluated as the difference between the two information measures. Different values of Δt (from 20 to 2.5 ms) were considered when evaluating the spike time information. Furthermore, by considering individually the terms due to the PSTH and those due to spike time correlations, we assessed whether temporal encoding was due to firing rate modulation or complex spike patterns.

Comparison of Information Transmitted by Spike Timing versus Total Spike Count

Examples of two cells are shown in Figure 1. At the left, PSTHs to deflection of each of the nine whiskers (50 trials per whisker) are shown. One of the cells (Figure 1A) displayed a strong response to its principal whisker and very weak responses to its nonprincipal ones, apart from D_3 (left panel). Hence, essentially all the information about stimulus location was transmitted simply by the total spike count (right panel, dotted line), and there was no extra benefit from spike timing (solid line).

The second cell (Figure 1B) represented stimulus location in a different way. Here, the temporal structure of the PSTH differed among stimuli: whisker D_2 evoked

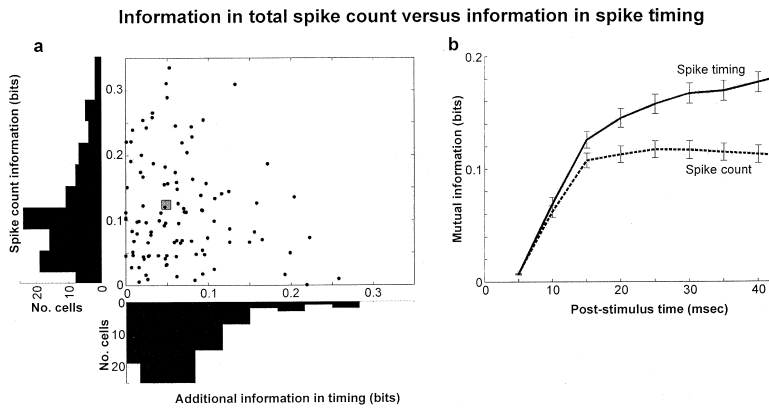


Figure 2. Variability in Coding Type across Cells

(A) Information in spike count is plotted versus the additional information in spike timing for each of the 106 neurons. The gray square indicates the mean values across cells. Histograms are shown on the respective axes. (B) Average of the mutual information functions shown in Figure 1B across cells: spike timing (solid line) and spike count (dotted line). The bars denote SEM.

spikes earlier than E_2/D_3 , which evoked spikes earlier than whiskers C_2/C_1 . At 10 ms poststimulus, there was equal information in spike count and spike timing (right panel), because only whisker D_2 had evoked a response by then (left panel). At 40 ms, however, other whiskers responded with a similar number of spikes, so the discriminability from the spike count code was lost. In contrast, discriminability was preserved in the spike timing code: from 20–30 ms onward there was much more information in spike timing (solid line) than in the total spike count (dotted line).

These results show that spike timing can be important for stimulus localization. How typical is the role of spike timing across the whole population of sampled neurons? Figure 2A is a plot for all cells of the spike count information versus the additional information in spike timing, both computed for the period 0–40 ms poststimulus onset at 5 ms resolution. The gray square indicates the mean values: 0.12 ± 0.01 (mean \pm SEM) bits for spike count information and 0.05 ± 0.01 bits for the additional information in timing. Note that there was little correlation at the population level between spike count information and the additional information in spike timing (Pearson correlation coefficient -0.1).

The time course of information (averaged over cells) is shown in Figure 2B. Early in the response (0–10 ms poststimulus onset), the spike count provided almost as much information (90% on average) as spike timing. Later, however, there was a significant advantage for the timing code. Whereas spike timing information continued to increase gradually, spike count information saturated. Indeed, at longer time windows (data not shown), spike count information actually decreased. This occurred for two reasons. First, for cells like that shown in Figure 1B, the essential temporal structure of the PSTH could not be reduced simply to counting spikes (see Panzeri and Schultz, 2000). Second, since the evoked response was transient (<50 ms), in long time windows, the spike count signal was degraded by “spontaneous” firing. At 40 ms, spike timing provided 44% more information than did total spike count.

Information Transmitted about Each Stimulus

Whisker D_2 is located in the center of a 3×3 array made up of whiskers $C_1-C_2-C_3$, $D_1-D_2-D_3$, and $E_1-E_2-E_3$. While the preceding section refers to the information averaged

over all nine stimulus sites, a second issue of interest is how well each particular whisker is represented in the spikes of a neuron in barrel-column D_2 . To answer this, we measured the “stimulus-specific information,” which reports how much the response to a given whisker differs from the average response across all whiskers (Equation 2 of Experimental Procedures). Figure 3A shows that, for neurons in barrel-column D_2 , the best encoded stimulus was the principal whisker D_2 , both when spike count (black bars) and spike timing (gray bars) were considered. In other words, neurons in barrel-column D_2 best reported whether or not the stimulus site matched the principal whisker. If the stimulus site was *not* D_2 , individual neurons in barrel-column D_2 provided a much less effective encoding of which other whisker might have been deflected. This condition can

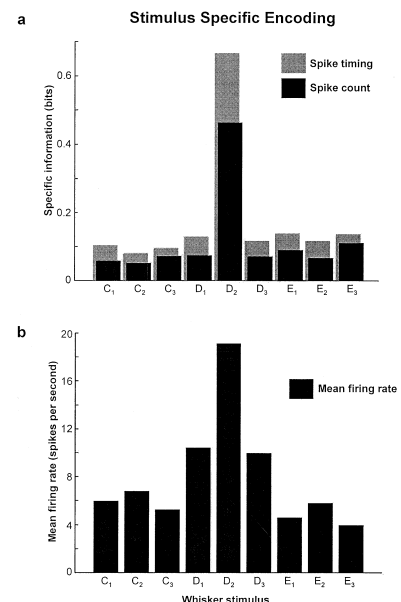


Figure 3. Stimulus-Specific Encoding

(A) Specific information about each stimulus location, averaged over all 106 cells, was computed for both spike timing (gray) and spike count (black) at 0–40 ms post stimulus onset, as described in Experimental Procedures. (B) Mean firing rate for the same time interval, averaged over the same cells.

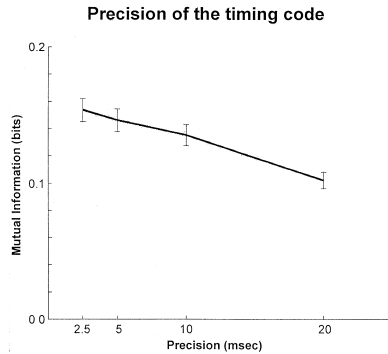


Figure 4. Precision of the Timing Code

Mutual information in spike timing, using the series expansion method, was computed for the interval 0–20 ms relative to stimulus onset for various time bin sizes. The values plotted here are bits in the total 0–20 ms poststimulus interval, averaged over all 106 cells. Bars denote SEM.

be understood by considering the number of spikes evoked by each stimulus (Figure 3B). Whisker D₂ evoked a large response, well differentiated from that of the other whiskers. The latter evoked smaller responses which were less discriminable from one another (for the receptive field structure of barrel D₂ neurons, see Armstrong-James et al., 1994). The stimulus-specific information added by spike timing was greatest about principal whisker stimulation, contributing 0.19 stimulus-specific bits when the stimulus site was D₂ and an average of 0.04 bits when the stimulus site was a surrounding whisker. Thirty-eight percent of the extra mutual information contributed by timing was specific to the principal whisker.

Precision of Spike Timing

What degree of precision underlies the spike timing code? To answer this question, we varied the resolution at which the spike times were binned and computed the average information across stimuli as a function of bin size (Figure 4). If information increases as bin size

is decreased, timing must be precise on the scale of the smaller bin size. Using a single, 20 ms bin, the average information across all 106 sampled neurons was 0.10 ± 0.006 bits. Reducing bin size to 10 ms, the information present in the 0–20 ms interval increased to 0.14 ± 0.008 bits. Further reductions of bin size to 5 ms and finally to 2.5 ms yielded additional increases in information to 0.146 ± 0.008 and 0.154 ± 0.008 bits over the 0–20 ms poststimulus interval, respectively. The shortest bin that could be robustly estimated was 2.5 ms. Hence, the precision of the code was at least 2.5 ms.

How was such high precision generated? The critical factor was that the latency difference between the response to whisker D₂ and that to surrounding whiskers “survived” the trial-to-trial variability in first spike time. The average latency to D₂ deflection was 10 ms, whereas that to surround deflection was 14 ms, consistent with Armstrong-James et al. (1992). Following D₂ deflection, the standard deviation of the first spike time across trials (i.e., the jitter) was just 2.6 ± 0.3 ms. Therefore, the reliability of first spike timing was sufficient to support reliable temporal coding of stimulus location based on latency differences.

Nature of the Temporal Code

Thus, spike timing, and not merely total spike count, conveys information about whisker location. Next, we consider the nature of the temporal coding. The simplest temporal code is one where stimuli are distinguished by different temporal profiles of firing rate modulation. Rate modulation coding is present whenever the profile of the PSTH differs across stimuli (Figure 1B) and can in fact carry information even in the absence of total spike count differences: the areas of the PSTHs can be equal but their profiles different. If the spike times are independent given the time-dependent firing rate, then the PSTH is a complete description of the response. In contrast, if spike times are not independent, more complex types of temporal coding are possible. To uncover complex temporal coding, it is necessary to measure the trial-to-trial correlation between the responses in different time

Coding by stimulus-independent correlations

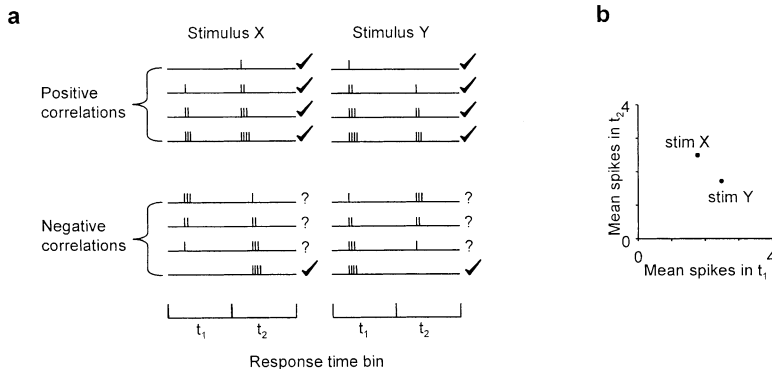


Figure 5. Influence of Stimulus-Independent Correlations on Temporal Coding

(A) shows raster plots for the response of a hypothetical neuron to two stimuli (X and Y) across two time bins (t_1 and t_2). The four rows of raster plots correspond to four different presentations of the stimuli. In the upper set, there are strong positive stimulus-independent correlations between time bins (e.g., on trials where there are more spikes in t_1 , there tend to be more spikes in t_2 , so called “correlated noise”). In the lower part, there are strong negative correlations (e.g., on trials where there are more spikes in t_1 , there tend to be fewer spikes in t_2 , “anti-correlated noise”). Although the mean response across trials for a given stimulus is equal for both types of noise correlation (B), there is a large difference in stimulus discriminability. Trials in which the stimuli can be discriminated (the

spike pattern provides an unambiguous stimulus identification) are indicated by check marks in (A). For positive stimulus-independent correlations, all eight trials are discriminable. For negative correlations, only two are discriminable and six trials are confusable (indicated by question marks).

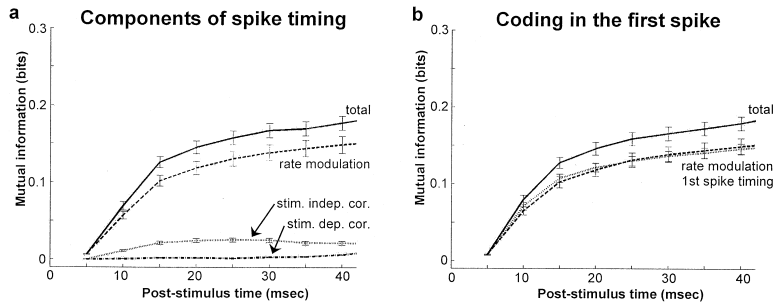


Figure 6. Nature of the Timing Code

(A) Each of the terms in the series expansion of the mutual information in spike timing (averaged over all 106 cells) is shown separately. The total (solid line) can largely be accounted for by the time-varying firing rate (or PSTH, dashed line). There is a smaller contribution from stimulus-independent spike patterns (dotted line) and a very small one from stimulus-dependent spike patterns (dash-dotted line). Bars denote SEM.
 (B) Information in the first spike. Mutual information in spike timing was computed using only the first spike on each trial. Results were

averaged over all 106 cells and plotted as the dotted line. For convenience of comparison, the total information in spike timing, using all spikes on each trial, is replotted (solid line), as is the time-varying firing component using all spikes (dashed line). Bars denote SEM.

bins. Any correlations that cannot be explained by independent firing can reflect the presence of characteristic information-bearing spike patterns (such as bursts).

Two distinct ways in which spike patterns can play a role in stimulus coding have been suggested: (1) stimulus-dependent correlations and (2) stimulus-independent correlations. In the first case, a neuron “tags” each stimulus by firing a stimulus-specific spike pattern (Abeles et al., 1993). Such patterns induce stimulus-dependent correlations that differ across stimuli and thus always increase the information content. In the second case, the coding is affected by interaction between the relative mean responses in different time bins. This interaction takes the form of correlations across stimuli (“signal correlation”) and correlations in the variability around the mean response (“noise correlation”); see Snippe (1996), Oram et al. (1998), Abbott and Dayan (1999), and Panzeri and Schultz (2000). These concepts are illustrated in Figure 5. The general principle is that when signal and noise correlations have the same sign, information is degraded; when they have the opposite sign, it is enhanced.

By measuring not only the PSTH, but also the stimulus-dependent and stimulus-independent correlations between pairs of spike times (see Experimental Procedures), our information estimation algorithm permits these various types of temporal code to be teased apart. The results of the analysis (averaged across the sample of 106 neurons) are plotted in Figure 6A. We found that at 40 ms poststimulus onset, PSTH modulations conveyed 83% of the total information. However, there was also a small but significant contribution (13.5%) from stimulus-independent spike patterns. Interestingly, there was very little contribution (3.5%) from stimulus-dependent spike patterns. Individual rat somatosensory cortical neurons appear to code stimulus location largely by differences in time-varying firing rate. Spike patterns play a smaller, synergistic role in a stimulus-independent manner.

Information Transmitted by Single Spikes

How important is each individual spike? How many spikes are necessary for accurate coding? To address these questions, we compared the spike timing information transmitted using all spikes after whisker deflection on each trial to that using only the first, second, third, or fourth spike, separately. The result was that 83% of

the total information in the spike train for 0–40 ms was available in the time of the first spike alone. If only the second spike of the train was considered, 26% of the information in the entire spike train was present. For the third and fourth spikes, the values were 10% and 3%, respectively. Clearly, the bulk of the information was conveyed by the first spike. The fact that the total information present when the spikes were considered individually exceeded 100% of that present in the full spike train indicates that spikes subsequent to the first one were partially redundant. This observation emphasizes the significance of the first spike. Moreover, for each time step in the 0–40 ms interval, the first spike accounted for essentially all of the information in firing rate modulation (Figure 6B). Hence, the *only* information-bearing part of the PSTH was the response onset. Note that it was not only the occurrence of the first spike (a spike count code) that mattered, but also its time. Over the period 0–40 ms poststimulus, the timing of the first spike carried 210% as much information as the simple presence of a spike.

The large amount of information contained in individual spikes is further highlighted in Figure 7, showing the distribution, across all cells, of the average information per spike within the first 10 ms after stimulus onset. The neurons transmitted on average 1.59 ± 0.1 bits/spike about which of the nine whiskers was stimulated. In the

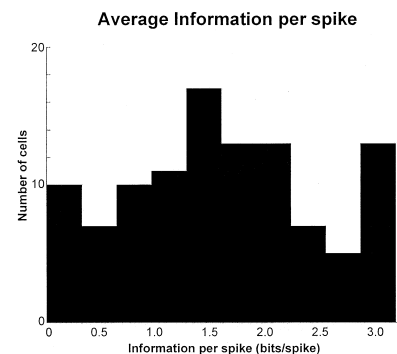


Figure 7. Average Information per Spike

For a given time interval, the information per spike is the mutual information in that interval divided by the trial-average spike count in that interval. This was computed for all 106 cells for the interval 0–10 ms relative to stimulus onset and shown as a histogram.

previous figures, results were given in units of bits per bin; here they are in bits per spike. (The values in Figure 7 are higher because the mean number of spikes across stimuli in this time interval was typically low, 0.1 spikes/trial). Although the most informative spikes were those in the interval 0–10 ms poststimulus, high values were also obtained for other response intervals (e.g., 1.28 bits/spike for 0–15 ms, 1.16 bits/spike for 0–20 ms).

Information measures always depend on the stimulus set; an expanded stimulus set may inflate the information per spike values. To evaluate this factor, we calculated the information per spike for stimulus sets containing different numbers of whiskers. Reducing the stimulus set to two whiskers (D_2 and a nonprincipal one) yielded values of 0.7 bits/spike in the 0–10 ms poststimulus window, where 1 bit/spike would be the physical maximum.

The overall analysis thus demonstrates that the first spike emitted after whisker deflection was not only the most informative spike in the train but also that it coded stimulus position very efficiently and reliably. The finding was true even when the considered stimulus set was small.

Discussion

In recent years, there have been a number of convincing demonstrations of sensory encoding by the timing of individual spikes (Rieke et al., 1997; Borst and Theunissen, 1999). These studies have considered dynamic stimuli with a complex variation across time and have shown that highly precise spike timing can serve to code when certain features within the stimulus occur (Ruyter van Steveninck et al., 1997; Buracas et al., 1998; Reinagel and Reid, 2000). However, it has remained unclear whether spatial stimulus features, not defined by temporal properties, might also be encoded by the timing of individual spikes. Here, we have provided quantitative evidence that the timing of single spikes—in particular the first stimulus-evoked spike—can also be informative about where a stimulus has occurred. Previous information theoretic studies reported evidence of temporal encoding of static visual stimuli by single units in the mammalian cortex, but only in long response windows typically containing many spikes, and with coarse precision (in the range of 25–50 ms) (Heller et al., 1995; Victor and Purpura, 1998).

A recent study (Ghazanfar et al., 2000) considered coding of whisker location in rat S1 and thalamus. Instead of using an information theoretic approach, the authors trained an artificial neural network to reconstruct stimulus location, using spike trains simultaneously recorded across a population of ~ 30 neurons as input. They found that the fraction of trials in which the reconstruction was correct increased as the size of the time bins decreased to 6 ms, implying that there was a contribution of temporal coding beyond the total spike count. Ghazanfar et al. also reported that disrupting correlations between neurons in the ensemble decreased coding accuracy. The advantage of their method is that it permits statements to be made about coding across large neuronal populations: this is very difficult with information theory due to the issue of lim-

ited sampling (see Experimental Procedures). On the other hand, the advantage of the present series expansion method (Panzeri et al., 1999; Panzeri and Schultz, 2000) is that it permits coding mechanisms in single neurons to be investigated more exhaustively than can be done with a neural network. Information analysis revealed (1) the cardinal role of the first spike, which contained 83% of total spike train information; (2) the relatively small role played by multispike patterns in single neurons; (3) the very high average information content per spike; and (4) the high precision of the temporal code—at least 2.5 ms.

We investigated the role of spike patterns in coding stimulus positions and found that complex, single neurons spike pattern play a much less critical (though synergistic) role than the first spike. This result is consistent with recent findings obtained with dynamic stimuli in the fly visual system (Brenner et al., 2000).

How many neurons in barrel-column D_2 must be sampled for reliable decoding of stimulus location? On the admittedly coarse assumption that neurons code stimulus location independently, we can make a simple estimate. Consider the problem of discriminating whisker D_2 from any other whisker: this is a stimulus set with two elements, D_2 and “not D_2 ”. The average neuron supplied 0.11 bits of mutual information about this stimulus set in the period 0–40 ms poststimulus onset. Since perfect discrimination requires $-1/9\log_2 1/9 - 8/9\log_2 8/9 = 0.5$ bits, 5 “average” neurons are necessary. While population recordings should be used to check for effects of redundancy or synergy, it is nevertheless clear that a very small number of neurons per column can be sufficient to support stimulus localization. Moreover, since the average number of spikes per trial fired by one neuron in the period 0–40 ms was 0.33, the implication is that the number of spikes necessary for reliable performance is 1.5. Although this is only a first approximation, the inevitable conclusion is that small numbers of spikes are capable of contributing to robust stimulus representation. This idea is consistent with recent work on a variety of sensory systems from fly (Bialek et al., 1991) to primate (Thorpe et al., 1996; Buracas et al., 1998; Reinagel and Reid, 2000).

A limitation of studies such as the present one is that they only reveal what information is available in the neuronal responses. How the nervous system might make use of the available information is a separate issue. Two specific questions are relevant. First, how rapidly is the encoded information actually integrated? At 10 ms poststimulus, the information in spike timing is only 10% greater than that in spike count. At 40 ms the total information is 225% higher than at 10 ms, and 44% of this is due to spike timing. It would be interesting to know which information is truly utilized by the cortex—the earliest available signal or the more reliable, later signal. The second question is whether the latency signal, which we have shown to be extremely informative, can be decoded by the cortex. Unlike the experimenter, the rat has no independent information about when each stimulus occurred. Hence, the animal itself cannot know that the latency on a particular trial was 5 ms as opposed to 10 ms, for example. In reply to this point, we note that the timing information can be decoded, in principle, by neurons with input from multiple columns. Consider

stimulation of whiskers D_1 and D_2 . Whisker D_1 evokes a short latency response in barrel D_1 and a longer latency response in barrel D_2 , while whisker D_2 evokes the opposite sequence. Thus, without knowledge of when the stimulus occurred, whisker identity can be decoded by comparing the order in which the two barrels become active. Information in spike timing is hence decodable from relative latencies, a very simple temporal population code.

The importance of such encoding would be paramount in situations where the specific site of peripheral receptor activation guides the animal's behavior. Rats navigate and localize objects using their whiskers (Hutson and Masterton, 1986; Brecht et al., 1997), and such behaviors rely upon somatosensory cortical processing (Hutson and Masterton, 1986; Hurwitz et al., 1990; Barneoud et al., 1991; Pazos et al., 1995). Indeed, in some conditions the position of a stimulated whisker can govern behavioral outcome (Harris et al., 1999). Spike timing could contribute to the speed and reliability of the cortical representation of peripheral events.

In general, many important features of the natural environment—the spatial structure of a visual scene, for example—are characterized by static rather than temporal properties. Use of the temporal dimension to encode nontemporal stimulus features, as appears to occur in barrel cortex, could contribute to the information carrying capacities of sensory cortex.

Experimental Procedures

Electrophysiology

Methodology is described in detail by Lebedev et al. (2000). All procedures conformed to NIH and international standards concerning the use of experimental animals. Twenty-two adult male Wistar rats weighing ~ 350 g were used. Anesthesia was induced by urethane (1.5 g/kg body weight, i.p.). The subject was placed in a stereotaxic apparatus (Narishige, Tokyo) and left somatosensory cortex exposed by a 4 mm diameter craniotomy. Body temperature was maintained near 37.5°C and, during the recording session, anesthetic depth was held at a consistent depth by monitoring hindpaw withdrawal, corneal reflex, and respiration rate.

At the end of the experiment, subjects were perfused with saline followed by 4% paraformaldehyde. After postfixation in 20% sucrose, a flattened slab of neocortex was frozen, cut into $40\ \mu\text{m}$ tangential sections, and processed to label nitric oxide synthase activity (Valtschanoff et al., 1993) in order to visualize barrel-columns in layer IV. To determine the columnar location of sampled neurons, electrode penetration sites were identified relative to the histological sections.

An array of six tungsten electrodes, arranged either as a single row or as a 2×3 matrix, with $300 \pm 50\ \mu\text{m}$ horizontal separation between adjacent electrode tips, was advanced into the cortical barrel field, focused on barrel-column D_2 . Typically, 1–2 electrodes penetrated barrel D_2 . The whole array was advanced in $100\ \mu\text{m}$ steps, and an effort was made to sample recording sites throughout the depth of cortex. Nonetheless, the great majority of neurons were likely to have been located between $300\text{--}950\ \mu\text{m}$, the depths at which thalamic axons from the ventral posterior medial nucleus terminate (Lu and Lin, 1993). Since the methodology did not permit us to register electrode depth with accuracy better than about $100\ \mu\text{m}$, we did not attempt to classify single neurons according to laminar position.

Neuronal activity was amplified and band-pass filtered in the range $300\text{--}7500$ Hz. Action potentials were digitized at 25 KHz, 32 points per waveform, time-stamped with 0.1 ms precision (Data-wave, Boulder, CO), and stored on a Pentium PC for offline analysis. Single-unit action potentials were discriminated by differences in

shape and amplitude using a custom template-matching algorithm. Typically, 1–2 single units were discriminated per electrode.

Individual whiskers were stimulated 3 mm from their base by a piezoelectric wafer (Morgan Matroc, Bedford, OH) that was controlled by a voltage pulse generator (A.M.P.I., Jerusalem). The stimulus was an up-down step function of $80\ \mu\text{m}$ amplitude and 100 ms duration, delivered once per second, 50 times for each vibrissa. The stimulated vibrissae were C_{1-3} , D_{1-3} , and E_{1-3} .

Data Analysis

In order to quantify how units code stimulus location, we applied information theoretic methods. Mutual information (Shannon, 1948) quantifies how well an ideal observer of neuronal responses can discriminate between all the different stimulus locations (S), based on a single response trial:

$$I(S;R) = \left\langle \sum_r P(r | s) \log_2 \frac{P(r | s)}{P(r)} \right\rangle_s, \quad (1)$$

where $P(r | s)$ is the conditional probability of observing a neuronal response r given a deflection of whisker s , $P(r)$ is the unconditional probability of response r (the average of $P(r | s)$ across stimuli), and $\langle \dots \rangle_s$ denotes an average across stimuli weighted by the probability $P(s)$ of deflecting whisker s .

The mutual information can be considered as a weighted average of the following stimulus-specific contributions:

$$I(s;R) = \sum_r P(r | s) \log_2 \frac{P(r | s)}{P(r)}. \quad (2)$$

The quantity $I(s;R)$ above measures how “far apart” the conditional response probability $P(r | s)$ is from the unconditional response probability $P(r)$. Hence, it quantifies how well an ideal observer can discriminate whether a specific stimulus s occurred, based on the neuronal response on a single trial. Using conventional terminology (Borst and Theunissen, 1999), we refer to $I(s;R)$ in Equation 2 as the stimulus-specific information. Note that this quantity has also been referred to as the stimulus-specific surprise (DeWeese and Meister, 1999).

To calculate neuronal responses, we considered a time window $0\text{--}T$ ms after the onset of whisker deflection. In calculating the information conveyed only by the total spike count, the response r was computed as the number of spikes emitted in this time window on one trial. To study information conveyed by spike timing, we divided the spike train into bins of size Δt , so that each possible spike train is a binary “word” of length $T/\Delta t$.

In practice, the probabilities in the above formulae must be estimated from a limited number of experimental trials N . This can lead to a significant upward bias in the mutual information given by Equation 1. Suppose that a neuron fires in a purely noisy way, irrespective of which stimulus occurs (zero mutual information). This means that the true underlying probabilities $P(r | s)$ will be the same for all s . Conversely, a neuron which fires selectively to certain stimuli (and does carry information) will have nonuniform $P(r | s)$: it will be high for some stimuli and low for others. When we estimate the set of $P(r | s)$ for the noisy neuron from a limited sample of experimental trials, our estimates of $P(r | s)$ will necessarily fluctuate around their true values. Thus, the estimated $P(r | s)$ will be nonuniform, as if the neuron was genuinely information bearing, and the estimated mutual information will be above zero. The magnitude of the bias depends on the number of trials. As N increases, the estimated probabilities become more accurate, and the bias in the mutual information decreases. A precise expression for the bias has been calculated mathematically, which is accurate provided that N is at least the number of different possible responses (Panzeri and Treves 1996; Golomb et al., 1997). The formula can be used to improve the accuracy of empirical information measurements: given the number of trials available in a particular experiment, the expected bias can be calculated and subtracted from the raw estimate of the mutual information. In the present case, the maximum number of spikes observed in $0\text{--}100$ ms was 9. Thus, the number of different spike counts was 10 (spike counts $0\text{--}9$), and the information in the total spike count was well sampled with 50 trials. For spike timing, response words of only length 4–5 could be considered by the brute force use of Equations 1 and 2.

We were able to improve on this temporal precision by using a time series expansion method (Panzeri et al., 1999; Panzeri and Schultz, 2000) that expresses the mutual information as an infinite series in the time window T . The order zero term depends only on the PSTH. The first order term additionally depends on pairwise correlations between spike times. Higher order terms introduce dependencies on successively higher order correlations. The key is that low firing rates, as found in rat somatosensory cortex, restrict the potential complexity of the neuronal code, making the series accurate even when curtailed at second order. Since only the PSTH and pairwise correlations must be estimated from the experimental data, limited sampling is less of a problem than it is with brute force use of Equations 1 and 2. In our case, response words of length 8–9 could be considered.

The second order approximation for the mutual information (Equation 1) can be expressed as a sum of three terms. One basic constituent of these equations is $r_s(t)$. This is the mean firing rate at time bin t following application of stimulus s —the PSTH. The first term in the approximation depends only on this quantity and hence measures the information conveyed by modulation of firing rates (the PSTH):

$$I_{PSTH}(S;R) = \Delta t \sum_t \left[\langle r_s(t) \log_2(r_s(t)) \rangle_s - \langle r_s(t) \rangle_s \log_2 \langle r_s(t) \rangle_s \right] + \frac{\Delta t^2}{2 \ln(2)} \sum_{t_1, t_2} \left[\langle r_s(t_1) r_s(t_2) \rangle_s (1 + \log_2 \langle r_s(t_1) \rangle_s \langle r_s(t_2) \rangle_s) - \log_2 \langle r_s(t_1) r_s(t_2) \rangle_s - \langle r_s(t_1) \rangle_s \langle r_s(t_2) \rangle_s \right]. \quad (3)$$

The other terms quantify the contribution of temporal spike patterns to the neural code. This is reflected by the appearance of the constituent $C_s(t_1, t_2)$, which is the probability density of observing spikes at times t_1 and t_2 given stimulus s , and $Ce_s(t_1, t_2) = r_s(t_1) r_s(t_2)$, which is the expected value of $C_s(t_1, t_2)$ in the case that the spike times are statistically independent.

The second term in the series expansion approximation measures the effect of stimulus-independent correlations:

$$I_{cor-ind}(S;R) = -\frac{\Delta t^2}{2 \ln(2)} \sum_{t_1, t_2} \left[\left\langle C_s(t_1, t_2) - Ce_s(t_1, t_2) \right\rangle \log_2 \frac{\langle Ce_s(t_1, t_2) \rangle_s}{\langle r_s(t_1) \rangle_s \langle r_s(t_2) \rangle_s} \right]. \quad (4)$$

The first multiplicative factor here is the spike time covariance (“noise correlation”), averaged over stimuli. The second multiplicative factor measures similarity of PSTHs across stimuli (“signal correlation”): it is positive when they are similar and negative when they are dissimilar. Figure 5 explains how stimulus-independent correlations can affect temporal coding.

The third term in the series expansion approximation quantifies the effect of stimulus-dependent correlations:

$$I_{cor-dep}(S;R) = \frac{\Delta t^2}{2 \ln(2)} \sum_{t_1, t_2} \left\langle C_s(t_1, t_2) \log_2 \left[\frac{C_s(t_1, t_2)}{Ce_s(t_1, t_2)} \div \frac{\langle C_s(t_1, t_2) \rangle_s}{\langle Ce_s(t_1, t_2) \rangle_s} \right] \right\rangle. \quad (5)$$

This term measures how well stimulus location is tagged by differences in trial-to-trial spike time correlations across the stimuli. The logarithmic factor in Equation 5 quantifies how much the normalized correlation for each stimulus $C_s(t_1, t_2)/Ce_s(t_1, t_2)$ differs from the average across stimuli: this factor is zero for any spike time pairs with no such normalized stimulus-dependent correlations.

To further verify that the spike timing information we estimated was accurate and robust, we made several checks. First, we verified that the mathematical assumptions of the series method (Panzeri et al., 1999) held for all cells. Firing rates must be low, in the sense that the number of spikes per trial (averaged over stimuli) is less than one in all time intervals considered. Also, correlations between spike times must have finite precision: in particular, the normalized correlation $C_s(t_1, t_2)/Ce_s(t_1, t_2)$ must not diverge for any time resolution Δt . Second, we computed a lower bound on the full spike timing information Equation 1 (introduced in the context of cortical spike train analysis by Reich et al. [2000]) that is very robust to sampling problems. We compared this to the series expansion for spike timing and found agreement to within 2% up to time window 0–50 ms (binsize 5 ms). Third, we compared the spike timing expansion with

brute force use of Equation 1, for time windows in which the latter was well sampled (40 ms time window for binsize 10 ms, and 20 ms for binsize 5 ms) and found agreement to within 1.5%. Last, to check the accuracy of the series expansion for spike counts, we compared the information computed in this way to that for the brute force use of Equation 1 for spike counts. We found agreement to within 0.5% for each time step up to 100 ms poststimulus. The above tests show that our findings are consistent and robust across different information methods and that the series expansion methods are accurate in the present case. The results reported here all use the series expansion methods.

Acknowledgments

We thank I. Erchova and G. Mirabella for assisting in the collection of some data and O. Lebedeva for histological processing of cortical sections. Supported by NIH grant NS32647, Telethon Foundation grant 984, and MURST. S. P. is supported by the Wellcome Trust, S. R. S. by the Howard Hughes Medical Institute.

Received September 14, 2000; revised December 27, 2000.

References

- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* **70**, 1629–1638.
- Abbott, L.F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101.
- Adrian, E.D. (1926). The impulses produced by sensory nerve endings: part I. *J. Physiol. Lond.* **61**, 49–72.
- Ahissar, E., Sosnik, R., and Haidarlu, S. (2000). Transformation from temporal to rate coding in a somatosensory thalamocortical pathway. *Nature* **406**, 302–306.
- Armstrong-James, M., and Fox, K. (1987). Spatiotemporal convergence and divergence in the rat S1 “barrel” cortex. *J. Comp. Neurol.* **263**, 265–281.
- Armstrong-James, M., Fox, K., and Das Gupta, A. (1992). Flow of excitation within rat barrel cortex on striking a single vibrissa. *J. Neurophysiol.* **68**, 1345–1358.
- Armstrong-James, M.A., Diamond, M.E., and Ebner, F.F. (1994). An innocuous bias in whisker use in adult rats modifies receptive fields of barrel cortex neurons. *J. Neurosci.* **14**, 6978–6991.
- Barneoud, P., Gyger, M., Andres, F., and van der Loos, H. (1991). Vibrissa-related behavior in mice: transient effect of ablation of the barrel cortex. *Behav. Brain Res.* **44**, 87–99.
- Bialek, W., Rieke, F., Ruyter van Steveninck, R.R., and Warland, D. (1991). Reading a neural code. *Science* **252**, 1854–1857.
- Borst, A., and Theunissen, F.E. (1999). Information theory and neural coding. *Nat. Neurosci.* **2**, 947–957.
- Brecht, M., Preilowski, B., and Merzenich, M.M. (1997). Functional architecture of the mystacial vibrissae. *Behav Brain Res.* **84**, 81–97.
- Brenner, N., Strong, S.P., Koberle, R., Bialek, W., and de Ruyter van Steveninck, R.R. (2000). Synergy in a neural code. *Neural Comput.* **12**, 1531–1552.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**, 4745–4765.
- Buracas, G.T., Zador, A.M., DeWeese, M.R., and Albright, T.D. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959–969.
- DeWeese, M.R., and Meister, M. (1999). How to measure the information gained from one symbol. *Network* **10**, 325–340.
- Ghazanfar, A.A., Stambaugh, C.R., and Nicolelis, M.A. (2000). Encoding of tactile stimulus location by somatosensory thalamocortical ensembles. *J. Neurosci.* **20**, 3761–3775.
- Golomb, D., Hertz, J., Panzeri, S., Treves, A., and Richmond, B. (1997). How well can we estimate the information carried in neuronal responses from limited samples? *Neural Comput.* **9**, 649–665.

- Harris, J.A., Petersen, R.S., and Diamond, M.E. (1999). Distribution of tactile learning and its neural basis. *Proc. Natl. Acad. Sci. USA* 96, 7587–7591.
- Heller, J., Hertz, J.A., Kjaer, T.W., and Richmond, B.J. (1995). Information flow and temporal coding in primate pattern vision. *J. Comput. Neurosci.* 2, 175–193.
- Hurwitz, B.E., Dietrich, W.D., McCabe, P.M., Watson, B.D., Ginsberg, M.D., and Schneiderman, N. (1990). Sensory-motor deficit and recovery from thrombotic infarction of the vibrissal barrel-field cortex. *Brain Res.* 512, 210–220.
- Hutson, K.A., and Masterton, R.B. (1986). The sensory contribution of a single vibrissa's cortical barrel. *J. Neurophysiol.* 56, 1196–1223.
- Lebedev, M.A., Mirabella, G., Erchova, I., and Diamond, M.E. (2000). Experience-dependent plasticity of rat barrel cortex: Redistribution of activity across barrel-columns. *Cereb. Cortex* 10, 23–31.
- Lu, S.M., and Lin, R.C. (1993). Thalamic afferents of the rat barrel cortex: a light- and electron-microscopic study using Phaseolus vulgaris leucoagglutinin as an anterograde tracer. *Somatosens. Mot. Res.* 10, 1–16.
- Oram, M.W., Foldiak, P., Perrett, D.I., and Sengpiel, F. (1998). The 'Ideal Homunculus': decoding neural population signals. *Trends Neurosci.* 21, 259–265.
- Panzeri, S., and Treves, A. (1996). Analytical estimates of limited sampling biases in different information measures. *Network* 7, 87–107.
- Panzeri, S., and Schultz, S. (2001). A unified approach to the study of temporal, correlational and rate coding. *Neural Comput.*, in press.
- Panzeri, S., Schultz, S.R., Treves, A., and Rolls, E.T. (1999). Correlations and the encoding of information in the nervous system. *Proc. R. Soc. Lond. B Biol. Sci.* 266, 1001–1012.
- Pazos, A.J., Orezza, S.L., McCabe, P.M., Dietrich, W.D., and Green, E.J. (1995). Recovery of vibrissae-dependent behavioral responses following barrelfield damage is not dependent upon the remaining somatosensory cortical tissue. *Brain Res.* 689, 224–232.
- Petersen, R.S., and Diamond, M.E. (2000). Spatial and temporal distribution of neural activity in rat barrel cortex and the coding of stimulus location. *J. Neurosci.* 20, 6135–6143.
- Reich, D.S., Mechler, F., Purpura, K.P., and Victor, J.D. (2000). Interspike intervals, receptive fields, and information encoding in primary visual cortex. *J. Neurosci.* 20, 1964–1974.
- Reinagel, P., and Reid, R.C. (2000). Temporal coding of visual information in the thalamus. *J. Neurosci.* 20, 5392–5400.
- Rieke, F., Warland, D., Ruyter van Steveninck, R.R., and Bialek, W. (1997). *Spikes: Exploring the Neural Code* (Cambridge, MA: MIT Press).
- Ruyter van Steveninck, R.R., Lewen, G.D., Strong, S.P., Koberle, R., and Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science* 275, 1805–1808.
- Shannon, C. (1948). A mathematical theory of communication. *Bell. Sys. Tech. J.* 27, 379–423.
- Simons, D.J. (1978). Response properties of vibrissa units in rat SI somatosensory neocortex. *J. Neurophysiol.* 41, 798–820.
- Snippe, H.P. (1996). Parameter extraction from population codes: a critical assessment. *Neural Comput.* 8, 511–529.
- Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *Nature* 381, 520–522.
- Tolhurst, D.J. (1989). The amount of information transmitted about contrast by neurones in the cat's visual cortex. *Vis. Neurosci.* 2, 409–413.
- Tolhurst, D.J., Movshon, J.A., and Dean, A.F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23, 775–785.
- Tovee, M.J., Rolls, E.T., Treves, A., and Bellis, R.P. (1993). Information encoding and the response of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* 70, 640–654.
- Valtschanoff, J.G., Weinberg, R.J., Kharazia, V.N., Schmidt, H.H., Nakane, M., and Rustioni, A. (1993). Neurons in rat cerebral cortex that synthesize nitric oxide: NADPH diaphorase histochemistry, NOS immunocytochemistry, and colocalization with GABA. *Neurosci. Lett.* 157, 157–161.
- Victor, J.D., and Purpura, K.P. (1998). Spatial phase and the temporal structure of the response to gratings in V1. *J. Neurophysiol.* 80, 554–571.
- Welker, C. (1971). Microelectrode delineation of fine grain somatotopic organization of (Sml) cerebral neocortex in albino rat. *Brain Res.* 26, 259–275.
- Werner, G., and Mountcastle, V.B. (1965). Neural activity in mechanoreceptive cutaneous afferents: stimulus-response relations, Weber functions, and information transmission. *J. Neurophysiol.* 28, 359–397.