

Neural Coding of Spatial Phase in V1 of the Macaque Monkey

Dmitriy Aronov,^{1,2} Daniel S. Reich,^{1,3,4} Ferenc Mechler,¹ and Jonathan D. Victor^{1,4}

¹Department of Neurology and Neuroscience, Weill Medical College of Cornell University, New York City, New York 10021; ²Department of Biological Sciences, Columbia University, New York City, New York 10027; ³Johns Hopkins University, Departments of Radiology and Neurology, Johns Hopkins Hospital, Baltimore, Maryland 21287; and ⁴The Rockefeller University, New York City, New York 10021

Submitted 18 September 2002; accepted in final form 26 January 2003

Aronov, Dmitriy, Daniel S. Reich, Ferenc Mechler, and Jonathan D. Victor. Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 89: 3304–3327, 2003. First published January 29, 2003; 10.1152/jn.00826.2002. We examine the responses of single neurons and pairs of neurons, simultaneously recorded with a single tetrode in the primary visual cortex of the anesthetized macaque monkey, to transient presentations of stationary gratings of varying spatial phase. Such simultaneously recorded neurons tended to have similar tuning to the phase of the grating. To determine the response features that reliably discriminate these stimuli, we use the metric-space approach extended to pairs of neurons. We find that paying attention to the times of individual spikes, at a resolution of ~30 ms, and keeping track of which neuron fires which spike rather than just the summed local activity contribute substantially to phase coding. The contribution is both quantitative (increasing the fidelity of phase coding) and qualitative (enabling a 2-dimensional “response space” that corresponds to the spatial phase cycle). We use a novel approach, the extraction of “temporal profiles” from the metric space analysis, to interpret and compare temporal coding across neurons. Temporal profiles were remarkably consistent across a large subset of neurons. This consistency indicates that simple mechanisms (e.g., comparing the size of the transient and sustained components of the response) allow the temporal contribution to phase coding to be decoded.

INTRODUCTION

Understanding how brain neuronal activity represents and manipulates information is one of the central goals of neuroscience. In comparison to the responses of peripheral sensory neurons, the responses of cortical sensory neurons to multiple presentations of the same stimulus are remarkable for their variability. This variability increases the technical challenges of studying neuronal coding and, more significantly, calls for an explanation.

Neuronal hardware is not intrinsically noisy, however, as attested to by the machine-like precision of sensory neurons *in vivo* (Reich et al. 1997) or of a cortical neuron whose inputs are carefully controlled (Mainen and Sejnowski 1995). Some have argued that the variability of cortical neurons is not a limitation of neuronal hardware in general but rather, a clue that neuronal coding makes use of rich statistical structure (Bullock 1997). Specific possibilities include the precise times of occurrence of impulses (Berry et al. 1997; Gawne 2000; Softky 1994; Théunissen et al. 1996), the pattern of intervals (Sen et al. 1996), correlations (Dan et al. 1998; Meister et al.

1995), and oscillations (Gray et al. 1989). The alternative view (Shadlen and Newsome 1998) is that variability is not a fundamental part of neural coding but rather represents a limitation that the nervous system has to deal with. In this view, reliable signals are extracted from unreliable neurons by averaging their activity across a population of similar neurons, and across appropriate intervals of time, and that detailed spatio-temporal structure simply is not relevant. Both the temporal and spatial aspects of neural coding, although often expressed as dichotomies, are best considered as continua. The timing of individual impulses must matter; the question is, over what time scale. The notion of population averages is only intended to apply to a local population; the question is, how local.

We address these issues by examining neural coding of spatial phase by neighboring pairs of neurons in primary visual cortex of the anesthetized, paralyzed macaque. Spatial phase is a fundamental aspect of spatial vision, crucial both for the extraction of local features (Burr et al. 1989) and overall scene perception (Oppenheim and Lim 1981). Certain properties of receptive fields are well known to be clustered (topographic location, size, and orientation tuning, for example). The functional organization of spatial phase tuning, however, is unknown. DeAngelis et al. (1999) study spatial phase as defined *relative to the receptive field envelope* and find that there is essentially no correlation among neighboring neurons. Here, we study a distinct quantity, spatial phase defined relative to a particular position in space (see DISCUSSION). We find that nearby neurons often have similar spatial phase tuning, although there is also some degree of variability. The presence of such variation does not suffice to imply that neural populations represent spatial phase on a neuron-by-neuron basis rather than via a summed population code. For coding purposes, it is useful to keep individual neurons' responses separate if their spatial phase preferences differ, but only if these differences are not overshadowed by trial-to-trial variability. Moreover, although individual neurons' responses may demonstrate a systematic dependence of their temporal characteristics on spatial phase (Victor and Purpura 1998), it is unclear how this temporal information would be useful if the temporal patterns were idiosyncratic to the individual neurons.

Our analysis addresses both of these questions. By an extension of the metric-space approach (Victor and Purpura 1997) to single-trial responses of pairs of neurons, we show

Address for reprint requests: J. D. Victor, Dept. of Neurology and Neuroscience, Weill Medical College of Cornell Univ., 1300 York Ave., New York, NY 10021 (E-mail: jdvicto@med.cornell.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.

that spatial phase coding is enhanced by keeping track of which neuron fires which spike. However, we also find that the extreme of a "labeled line" code is not necessary to realize these advantages. A previous shortcoming of the metric-space method is that it provided no way to compare detailed temporal characteristics across neurons. We overcome this problem here via the notion of "temporal profiles." Temporal profiles provide a quasilinear framework that accounts for the bulk of the temporal features identified by the metric-space method. Comparison of temporal profiles reveals a remarkable similarity across the population of neurons. Moreover, these temporal profiles are readily described in terms of "transient" and "sustained" components and thus provide a substrate to extract temporal information in a universal and straightforward fashion.

METHODS

Physiological preparation

Standard acute preparation techniques were used for electrophysiological recordings from single units in the primary visual cortex (V1) of the primate (cynomolgus monkeys, *Macaca fascicularis*). All procedures used were in accordance with institutional and National Institutes of Health guidelines for the care and experimental use of animals. Some details of the techniques used were given earlier (Mechler et al. 1998).

Experiments were performed on 14 adult animals, weighing 3–4.5 kg. Prior to surgery, animals were given atropine (0.1 mg/kg im) and then anesthetized with ketamine (10 mg/kg im). Anesthesia was maintained with sufentanil ($3\text{--}6\ \mu\text{g} \cdot \text{kg}^{-1} \cdot \text{h}^{-1}$ iv). Paralysis was induced following all surgical procedures and maintained with pancuronium ($0.4\ \text{mg} \cdot \text{kg}^{-1} \cdot \text{h}^{-1}$ iv). Dexamethasone ($1\ \text{mg} \cdot \text{kg}^{-1} \cdot \text{day}^{-1}$ im) and gentamicin ($5\ \text{mg} \cdot \text{kg}^{-1} \cdot \text{day}^{-1}$ im) were given to prevent the development of cerebral edema and infection, respectively. The animal was ventilated through an endotracheal tube. Heart rate, echocardiography (EKG), arterial blood pressure, and end-tidal CO_2 were continuously monitored with a Model 78354A Hewlett-Packard Patient Monitor and kept in the normal physiological range. Core body temperature was maintained between 37 and 38°C using a thermostatically controlled heating pad. The EEG was obtained from frontal leads and monitored on an oscilloscope.

A limited unilateral craniotomy to expose the primary visual cortex was made overlying and posterior to the lunare sulcus (the Horsley-Clarke stereotaxic coordinates were typically 14–16 mm posterior and 14–16 mm lateral). A 1–2 mm durotomy was made for the recording electrode, which was stabilized after insertion by agarose gel.

Extracellular recording

Spike responses of single units were recorded extracellularly. We used single-tip (Ainsworth et al. 1977; Merrill and Ainsworth 1972) electrodes, typical resistance 2 M Ω , or quartz-coated platinum-tungsten fibers tetrodes (Thomas Recording, Giessen, Germany). Tetrodes had a conical tip, with four contacts of $\sim 1\ \text{M}\Omega$ each, $\sim 25\ \mu\text{m}$ apart: one at the apex and three arranged in radial symmetry on the walls of the conical surface. A stepper motor advanced either type of electrode in 1- μm steps.

The signals from the electrode or tetrode channels were passed through a multi-channel differential head-stage amplifier (NB Labs, Denison TX, or Neuralynx, Tucson AZ) and then further amplified and filtered (0.3- to 6-kHz pass-band, Lynx-8 multichannel differential amplifier, Neuralynx). Candidate spike waveforms, detected by a threshold criterion, were digitized at 25 kHz within a short (~ 1.2 ms) temporal window containing the peak amplitude and digitally recorded (Discovery software, DataWave Technologies, Longmont,

CO). Multiple single units were isolated by cluster analysis of spike waveforms initially performed on-line (Autocut, DataWave Technologies), then off-line (custom software). Isolation criteria included stability of principal components of spike waveforms and a 1.2-ms minimum interspike interval consistent with a physiologic refractory period. Spike times for further data analysis were identified off-line to 0.1-ms precision, the accuracy to which the clocks of the recording computer and the stimulus generator were synchronized.

Histology and laminar assignment of recording sites

Experiments lasted for 4–5 days, at the end of which the animal was sacrificed by infusion of a lethal dose of methohexital. After transcardiac perfusion with 4% paraformaldehyde in phosphate-buffered saline, a block of the occipital cortex was processed for histological reconstruction of the electrode track in Nissl-stained sections, aided by electrolytic lesions ($5\ \mu\text{A} \times 5\ \text{s}$, electrode positive) on some tracks and deposition of the lipophilic fluorescent dye DiI (D-282; Molecular Probes, Eugene, OR) on others. This confirmed that all recordings were in V1. The precise laminar position was identified for two-thirds of the recording sites, and the relationship with respect to the granular layer was identified in the remainder.

Optics

The eyes were treated with anti-inflammatory (fluribuprofen) and anti-bacterial (neomycin/polymyxin/dexamethasone) ophthalmic solutions at least daily. Pupils were dilated (1% atropine) and covered with gas-permeable contact lenses. Artificial pupils (2 mm) and corrective lenses were used to focus the stimulus on the retina. Optical correction was estimated by retinoscopy and then refined by optimizing responses of isolated single units to high spatial frequency visual stimuli.

Cell classification

We used the F_1/F_0 modulation ratio to classify cells as simple or complex: if the fundamental of the response to a drifting grating of near optimal spatial parameters was larger than the DC component (after subtraction of the maintained rate of firing), then the cell was considered simple and complex otherwise (De Valois et al. 1982; Movshon et al. 1978a; Skottun et al. 1991). We recognize that the dichotomous nature of this classification has recently been called into question (Mechler and Ringach 2002), but we used it because it is quantitative, widely accepted, and likely related to how responses depend on spatial phase.

Visual stimulation

Visual stimuli were generated by a special purpose stimulus generator (Milkman et al. 1978, 1980) under the control of a PDP-11/93 computer and displayed on a Tektronix 608 monochrome oscilloscope (green phosphor, 150 cd/m² mean luminance, 270.32-Hz frame refresh). The luminance of the display was linearized with lookup tables in the range of 0–300 cd/m². At the 114 cm viewing distance of the animal, the stimuli appeared in a 4° circular aperture on dark background.

Neuronal receptive fields were characterized in a standard way using drifting sine gratings: tuning was measured first for orientation, then for spatial frequency, and finally for temporal frequency, each parameter optimized for subsequent tuning measurements. The contrast response function was measured using the optimal sine grating. Receptive field characterization was always done for the most responsive unit and often for a second unit.

Responses as a function of spatial phase

To assess responses as a function of spatial phase, we presented full-field ($4 \times 4^\circ$) sinusoidal luminance gratings at or near the optimal orientation and spatial frequency at each of 16 equally spaced spatial phases. Stimuli were organized into runs of eight repetitions (236 ms each) of a single spatial phase, separated by 710 ms presentations of a uniform field at the mean luminance. After each of these sets, the spatial phase was changed randomly. Typically four to eight repetitions of the block of unique runs were obtained, with the order of runs within each block randomized, resulting in 32–64 presentations of each of the 16 spatial phases. Runs were aborted if spike discrimination became unreliable or if there was a major change in responsiveness.

Data analysis

Off-line data analysis was performed with custom software written in C and in the Matlab (Mathworks, Natick MA) programming environment. Our goal was to determine the extent to which spatial phase is represented in single-trial responses to static gratings and to characterize the features of spike trains that provide this representation. To this end, we extended the metric-space approach (Victor and Purpura 1996, 1997) from single-cell recordings to responses of simultaneously recorded cells. As described in the following text, this allows us to characterize the role played by spike counts, spike times, and the distribution of spikes across pairs of simultaneously recorded neurons. This approach separates these roles by constructing notions of dissimilarity (“metrics” or rules for calculating distances) between neural responses. Characterization of neural coding via metrics avoids making strong assumptions (such as approximate linearity of the stimulus-response relationship or approximate Poisson-ness of the spike trains) about the neural response.

Single-unit metrics

This basic family of metrics—the “spike time” metrics of Victor and Purpura (1996, 1997), determines whether spike *timing* carries information about the stimulus in addition to the information carried by spike count. If the timing of individual spikes does carry additional information, the metrics determine (via a parameter q , defined in the following text) the temporal precision at which spike timing is relevant. In a “spike time” metric, the distance between two spike trains is defined as the minimal total “cost” of a sequence of elementary steps that transforms one spike train into the other. The following elementary transformation steps are allowed: insertion of a spike (for a cost of 1), deletion of a spike (for a cost of 1), and shifting a spike by an amount of time Δt (for a cost of $q\Delta t$). Because a spike in one train can always be transformed into a spike in another train for a cost of 2 (via its deletion from one train and insertion into the other), spikes are seen as similar only if they occur within $2/q$ of each other. When $q = 0 \text{ s}^{-1}$, spike timing is irrelevant to the distance calculation, and the metric reduces to a comparison of responses based solely on the number of spikes. To use the metrics to identify features of neural coding, we determine the extent to which the metrics induce a “stimulus-dependent clustering” of the responses (see following text). That is, we calculate a quantity H (formally, an information) that indicates the extent to which responses to the same stimulus are close (i.e., similar), while responses to different stimuli are distant (i.e., dissimilar). The calculation of H depends on the pairwise distances between responses to identical stimuli and to distinct stimuli and thus on the parameters of the metric (here, q). If spike timing carries no phase information, then H will be maximal at $q = 0 \text{ s}^{-1}$. This is because alternative metrics with $q > 0$ are influenced by the timing of spikes, which is hypothesized to be irrelevant to the identity of the stimulus. The converse situation is that the timing of spikes, and not just the number of spikes, systematically depends on spatial phase, and this

dependence provides information about spatial phase that is not provided by the spike counts. In this case, H will be maximal for some $q > 0$. This formalizes the notion that a metric that considers spike timing at the appropriate temporal scale will be able to identify the stimulus that elicited a response with greater certainty (i.e., higher H) than a metric that ignored spike timing. In this case, the value of q at which H is maximal determines the precision of spike timing (as $1/q$, in units of seconds) that is relevant to coding. For further details and discussion, see Victor and Purpura (1996, 1997).

Multi-unit metrics

To extend the analysis to multiple neurons, two changes are required. First, we represent a simultaneously recorded response of multiple neurons as a *single* sequence of *labeled* events. That is, a multi-unit response is considered as a “labeled spike train,” where a label assigned to each spike indicates its neuron of origin. For a population of L neurons, we use the integers 1, 2, ..., L as labels. These labels are merely abstract tags—no serial ordering is implied. Second, we introduce a new parameter k into the metric. The new parameter k describes the importance of the neuron of origin of a spike, much as q describes the importance of the time of a spike. Metrics used to characterize coding in such labeled spike trains allow all of the elementary transformations used by the single neuron metric (insertion of a spike, deletion of a spike, and moving a spike). In addition, these metrics include an elementary step in which a spike’s label (i.e., neuron of origin) is changed. This transformation is assigned a cost of k .

The dimensionless parameter k thus indicates the importance of distinguishing spikes that are fired by different neurons. When $k = 0$, the multi-unit metric reduces to the single-unit metric described earlier because there is no cost for reassigning the neuron of origin of a spike. This amounts to considering the neuron of origin to be irrelevant to coding, a situation we designate as a “summed population” code. When $k \geq 2$, spikes from different neurons are never considered to be similar because the cost of matching them (by changing the label of either 1) is no less than the cost of deleting a spike from one neuron’s response and inserting it into another neuron’s response. Thus for $k \geq 2$, the distance between two multi-unit responses becomes equal to the sum of the distances between individual cell responses calculated by the single-unit metric. Because each neuron’s activity is considered independently and noninterchangeably labeled, we designate this situation a “labeled line” code. The range of values of k from 0 to 2 provides a continuum between the extremes of the “summed population code” ($k = 0$) and the “labeled line code” ($k = 2$). For intermediate values, spikes in two multi-unit responses that occur sufficiently close to each other in time can be seen as similar even if they are fired by different neurons.

An efficient algorithm for the calculation of single-unit metric distances has been described before (Victor and Purpura 1996). Its running time is proportional to N^2 , where N is the typical number of spikes fired by a neuron. For this work, we used a straightforward extension of this algorithm to multi-unit metrics (Victor and Purpura 1997). For a population of L neurons, the running time of this extension is on the order of N^{2L} . A more efficient version of the algorithm has been recently developed (Aronov 2003) in which the running time is on the order of N^{L+1} .

Stimulus-dependent clustering

The next stage of our analysis is to determine to what extent our candidate notions of distance provide for discrimination of responses to different spatial phases. That is, for each candidate metric, we determined the extent to which repeated responses to the same spatial phase lie at shorter distances from each other than responses to different spatial phases (i.e., whether the responses to each spatial phase form separate clusters). To quantify such clustering, we utilized

the approach of Victor and Purpura (1996, 1997), which calculates the amount of “transmitted information,” H_{data} , in bits. This information-theoretic quantity is upwardly biased for finite data samples (Carlton 1969; Panzeri and Treves 1996; Treves and Panzeri 1995). We estimated the upward bias by repeating the calculation of transmitted information for data sets that consisted of a random reassignment of responses to the stimulus classes (Panzeri and Treves 1996). We corrected the information estimate by subtracting the mean of 10 such calculations from H_{data} , and we denote the difference $H_{data} - H_{resampled}$ by H .

The net result of the preceding calculations is thus a scalar measure H of the strength of stimulus-dependent clustering derived from each member of a family of notions of similarity. The family of notions of similarity is parametrized by the extent to which spike timing (q) and neuron of origin (k) are significant. This analysis (calculation of the distances between the responses, calculation of H_{data} and $H_{resampled}$ from the pairwise distances, and calculation of the final measure of clustering $H = H_{data} - H_{resampled}$) was performed for all (q, k) pairs in which $q = 0, 1, 2, 4, \dots, 512 \text{ s}^{-1}$, and $k = 0, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0, 1.25, 1.5, 1.75$, and 2 .

Multidimensional scaling

The preceding analysis indicates the extent to which the responses to each of the spatial phases are discriminable and whether this discrimination is best accomplished by consideration of the timing of responses ($q > 0$), the neuron of origin of each response ($k > 0$), or both. However, it ignores the size of the difference between pairs of stimuli. Among our stimuli, some pairs were quite similar to each other (e.g., 22.5° apart in phase), whereas others were diametrically opposite (180° apart in phase). Features of spike trains that convey information about spatial phase should not only support discrimination of stimuli that are of different spatial phases but, also, should provide a “representation” of spatial phase. That is, responses to similar spatial phases should be more nearly alike than responses to very different spatial phases. Moreover, spatial phase is intrinsically a cyclic quantity. It is therefore natural to ask whether the neural responses, when compared by a candidate metric, have a corresponding behavior.

To determine whether response features indeed provide for representation (and not merely discrimination) of spatial phase, we need to go beyond the information-theoretic quantity H , which indicates discriminability but not relationships. That is, we need to consider the geometry of the response clusters elicited by gratings at each spatial phase and not only whether the clusters are distinct.

To determine this geometry, we used a standard technique, multidimensional scaling (Kruskal and Wish 1978). In general, multidimensional scaling embeds a set of point in a Euclidean space so that the distances between the points correspond to prespecified numbers. In the present application, the “points” correspond to the individual responses. The “distances” correspond to the distances provided by a metric with specified values of q and k . Thus multidimensional scaling assigns coordinates to each spike train so that the standard Euclidean distances between spike trains are the best possible approximations of the distances given by a metric. According to the (standard) procedure of Kruskal and Wish (1978), this best fitting Euclidean embedding is found by determining the eigenvectors of a matrix A_{jk} whose entries are given by

$$A_{jk} = -\frac{1}{2} \left(d_{jk}^2 - \frac{1}{M} \sum_r d_{rk}^2 - \frac{1}{M} \sum_s d_{js}^2 + \frac{1}{M^2} \sum_{rs} d_{rs}^2 \right) \quad (1)$$

In Eq. 1, d_{jk} indicates the metric distance between spike trains j and k in the set of M spike trains. The m th eigenvector of A is a vector of length M . After division by the square root of the m th eigenvalue, its M entries yield the m th coordinate of the embedded spike trains. Note that this procedure is carried out on all M responses, without regard as to which stimulus elicits each response.

Fitting ellipses

Each of the 16 stimuli elicits a subset (of size $M/16$) of the M responses and thus corresponds to a subset of the M embedded points found by the preceding multidimensional scaling procedure. For each of the 16 spatial phases, the location of this corresponding subset of embedded points can be summarized by its centroid. (In contrast, H summarizes the extent to which these subsets overlap but ignores their relative locations.) The coordinates of each centroid are simply the average of the coordinates of the corresponding set of spike trains as provided by multidimensional scaling. One would expect the 16 centroids of these clusters in our experiment to lie along a closed curve due to the cyclic nature of spatial phase. Because linear mechanisms are a fundamental ingredient of receptive field models and because a linear response must fall on or near an ellipse (see APPENDIX) (see also Movshon et al. 1978a), we sought to characterize the positions of these centroids in terms of best-fitting ellipses.

We therefore fitted ellipses to the 16 centroids of responses to each of the spatial phases. We found best-fit ellipses by minimizing the mean squared distance between the centroids and 16 points located at constant phase intervals around an ellipse whose shape, position, and orientation were allowed to vary freely. (This step can be carried out as a linear regression.) To characterize the arrangement of response clusters, we quantified the shapes of the best-fit ellipses by their axis ratios. The shape of the ellipse can vary from a doubly covered line segment (an axis ratio of 0) to a circle (an axis ratio of 1). The former extreme indicates that only one mechanism effectively contributes to the response. The latter extreme corresponds to two spatiotemporal mechanisms in quadrature (Emerson 1997; Emerson and Huang 1997; Heeger 1992; Marcelja 1980; Pollen et al. 1985), a situation in which the cycle of spatial phase is faithfully represented by the circular trajectory of responses.

We quantified the goodness of fit of these ellipses by the variance in the layout of the centroids left unexplained by the 16 corresponding points on the best-fit ellipse.

Temporal profiles

Multidimensional scaling of the pairwise metric distances creates an arrangement of responses in an abstract space that depends on their temporal structure, but it does not identify which aspects of the neural responses contribute to this arrangement. The analysis described in the following text seeks a simple temporal interpretation for the coordinates identified by multidimensional scaling. In particular, if the coordinates of the embedded responses could be derived by linear operations applied to the responses (see APPENDIX), then this procedure will identify them. If the embedded responses do not form an ellipse but nevertheless may be construed as a combination of a small number of factors (that do not vary sinusoidally with spatial phase but combine linearly), this procedure will also identify these factors and their time courses.

Let us assume that multidimensional scaling has embedded the M responses into a D -dimensional response space. We seek up to D distinct factors (“temporal profiles”) whose linear superposition accounts for the observed embedding coordinates. We choose a desired temporal resolution for the calculation of the temporal profiles and divide the response period into K time bins, corresponding to the desired resolution. Because we postulate that these D factors operate linearly on the response trains, our analysis seeks a $K \times D$ matrix \mathbf{P} whose entries are the weight of each of the D mechanisms in each of the K bins.

We initially consider single-unit responses. The coordinates of the M responses (as provided by multidimensional scaling) constitute a coordinate matrix \mathbf{C} , of size $M \times D$. Corresponding to the desired number of bins K in the temporal profiles, we create a second representation of the M single-unit responses by an $M \times K$ matrix \mathbf{R} . A row of \mathbf{R} contains the number of spikes of a particular response in

each of the K bins. To find a correspondence between linear operations on the spike trains and the embedded coordinates \mathbf{C} , we seek a $K \times D$ matrix \mathbf{P} such that $\mathbf{R}\mathbf{P}$ best approximates \mathbf{C} in the least-squares sense, other than an arbitrary translation. Allowance for an arbitrary translation is necessary since the origin of the embedding identified by multidimensional scaling does not necessarily correspond to the null spike train, and the addition of an arbitrary translation to all embedded coordinates does not change their relative distances. To include the effects of an arbitrary translation, we append a column of 1's to \mathbf{R} to form an $M \times (K + 1)$ matrix \mathbf{R}' . We then solve for the $(K + 1) \times D$ matrix \mathbf{P}' such that $\mathbf{R}'\mathbf{P}'$ best approximates \mathbf{C} in the least squares sense (a standard linear regression). The last row of \mathbf{P}' corresponds to the arbitrary translation, i.e., the coordinates of the null response. With the last row deleted, the n th column of \mathbf{P}' is the *temporal profile* $P_n(t)$ of the n th dimension in the response space, with the understanding that time t is discretized into K bins. That is, $P_n(t)$ indicates the weight with which the response at time t contributes to the n th coordinate in the multidimensional response space. Explicitly, it provides an approximation for the n th coordinate of the m th spike train

$$c_{mn} = \int_a^b P_n(t) \cdot R_m(t) dt + \sigma_n \quad (2)$$

where c_{mn} is an entry of \mathbf{C} , $R_m(t)$ describes the firing rate of the m th response, σ_n is the translation constant for the n th dimension, and the response is considered on the time interval $a \leq t \leq b$.

The notion of temporal profiles is readily extended to multi-unit responses, as is this method of calculating them. For a response containing spikes from L neurons, we construct an $M \times KL$ matrix \mathbf{R} by horizontally concatenating L matrices (each $M \times K$) constructed by the preceding method from the individual units' responses. As in the preceding text, we augment the matrix \mathbf{R} by appending a column of 1's, to form an $M \times (KL + 1)$ matrix \mathbf{R}' . We then solve for a $(KL + 1) \times D$ matrix \mathbf{P}' such that $\mathbf{R}'\mathbf{P}'$ best approximates \mathbf{C} in the least squares sense. The n th column of this matrix \mathbf{P}' contains L segments of length K . The l th segment of the n th column of \mathbf{P}' is the temporal profile of the l th neuron along the n th dimension. We denote this segment, a list of K numbers, by $P_{n,l}(t)$, with the understanding that time t is discretized into K bins. The n th coordinate of the m th multi-unit response can then be approximated by

$$c_{mn} = \sum_{l=1}^L \left(\int_a^b P_{n,l}(t) \cdot R_{m,l}(t) dt \right) + \sigma_n \quad (3)$$

where $R_{m,l}$ is the firing rate of the l th neuron in the m th multi-unit response and σ_n is the translation constant for the n th dimension, obtained from the final column of \mathbf{P}' .

To the extent that Eqs. 2 and 3 provide a good match for the embedding identified by multidimensional scaling, they imply that the pairwise distances between spike trains can be recovered from linear operations on the responses. Because the metric distances are not created from linear operations on the spike responses, there is no a priori guarantee that such approximations will be accurate—so the

very existence of temporal profiles makes a nontrivial statement about the temporal representation of phase. Moreover, because the temporal profiles indicate how the coordinates are derived from the temporal structure of individual responses, they provide an interpretation of the coordinates and a way to compare coding across populations of neurons.

Portions of this material were presented at the annual meetings of the Society for Neuroscience (2000) and The Association for Research in Vision and Ophthalmology (2001).

RESULTS

Altogether, we recorded responses of 140 units (Table 1). This database was pared down to 70 units based on the presence of stable responses to the gratings (no evident change in mean rate over time, and no evident change in phase tuning over time) that differed from the maintained firing rate. Spatial phase tuning or selectivity was not a criterion for cell selection. Examples of rasters from cells that met these selection criteria are shown in Fig. 1. Note that the phase preference of each cell remains constant over the data-collection period. Moreover, the phases at which there are qualitative transitions in the nature of the response (e.g., from an ON response to an OFF response) do not change across the trials. Examples of such transitions include: Fig. 1A, both sets of rasters, 90–112.5°; Fig. 1B, *top rasters*, 337.5–0°, *bottom rasters*, 270–292.5°; Fig. 1C, *bottom rasters*, 270–292.5°. This kind of stability would not have been present had the eyes been moving.

Data from all cells that met these criteria (38 simple cells, 32 complex cells) were submitted to metric space clustering analysis, ellipse fitting, and the derivation of temporal profiles as single units (see METHODS). We used all pairs of simultaneously recorded units obtained from these 70 neurons for analysis of ellipse fitting and temporal profiles by neuron pairs. These paired recordings consisted of 29 simple-simple pairs, 11 simple-complex pairs, and 17 complex-complex pairs, many of which were overlapping. Because of the computational burden required by existing algorithms, not all of these paired data sets were subjected to the metric space clustering analysis. Instead, this analysis was restricted to pairs of simultaneously recorded cells of the same class (either simple or complex). At sites at which three or more units of the same class were recorded, cells were randomly paired for analysis (and, if the number of cells was odd, a randomly chosen unpaired cell was not analyzed). This resulted in a total of 22 simple cells and 18 complex cells (Table 1) analyzed in disjoint simultaneously recorded pairs (11 simple-simple pairs, 9 complex-complex pairs). For each analysis that follows, we examine representative neuron pairs before presenting the results of the population.

TABLE 1. Summary of experiments performed and data sets analyzed

	Individual Cells			Pairs of Cells			
	S	C	Total	S-S	S-C	C-C	Total
Number recorded	60	80	140	88	45	117	250
Number used for measurement of clustering	22	18	40	11	0	9	20
Number used for ellipse fitting and temporal profile analysis	38	32	70	29	11	17	57

The entries in the right half of the table tally pairs of units. These units are also included in the single-unit analysis (left half of the table). Criteria for data set selection and for simple/complex classification are in the text.




FIG. 1. Raster plots of responses of 3 pairs of neurons to grating stimuli at each of 16 spatial phases. Spatial phases differed in steps of 22.5° ; every fourth spatial phase is labeled ($0, 90, 180, 270^\circ$). *A*: 2 simple cells (*units 430908.st*) with similar phase preferences, 50 trials each. *B*: 2 simple cells (*units 410106.st*) with dissimilar phase preferences, 64 trials each. *C*: 2 complex cells (*units 431108.st*), 48 trials each. The 1st trial is the bottom raster. Sets of rasters are labeled by circles and solid squares, to correspond to the labeling in subsequent figures. The stimulus presentation period (237 ms) and analysis period (473 ms) are indicated by the scale adjacent to each set of rasters.

Stimulus-dependent clustering

The goal of this analysis was to determine how spatial phase information is coded across a local population of neurons by comparing joint responses of pairs of neurons and responses of individual neurons considered in isolation. We compared transmitted information in single-unit and pair responses and explored their dependence on the precision of spike timing (measured by the parameter q) and the importance of the neuron of origin of each spike (measured by the parameter k).

Sample data set 1: pair of simple cells with similar phase preferences

Analysis of responses of two simultaneously recorded simple cells to gratings at 16 spatial phases are shown in Fig. 2*A* (*left*). The two cells clearly have phase-dependent responses and similar spatial phase preferences. Responses of both cells are the largest at spatial phases $135\text{--}225^\circ$ and are small outside of this range. In the range of the preferred phases, both cells respond with a transient peak at the onset of the stimulus followed by a sustained ON response, both of which are phase dependent. Both cells show negligible dependence of the OFF response on spatial phase.

The two cells were first analyzed independently with the single-unit metric (curves in Fig. 2*A*, *middle*). Maximal clustering for these cells is achieved for q in the range of 16 to 32 s^{-1} . This indicates that stimulus-dependent clustering is stronger when the temporal structure of responses is taken into account ($q > 0\text{ s}^{-1}$) than when only the number of spikes is considered ($q = 0\text{ s}^{-1}$). As q increases beyond the optimal values, H decreases, eventually to chance level ($H \approx 0$). This decline in H for values of $q > 32\text{ s}^{-1}$ indicates that comparing responses with a metric that is sensitive to very small temporal shifts of spikes degrades the relationship between spatial phase and response cluster. In other words, taking into account the timing of spikes improves the sharpness of the dependence of responses on spatial phase but only down to a particular timing resolution. This timing resolution is measured by $1/q$. In typical data sets, such as this one, the temporal precision $1/q$ is in the range 30–60 ms. For many cells, such as the one denoted by squares in Fig. 2*A*, H shows a substantial rise even for small values of q . That is, analysis of the responses with only a coarse sensitivity to spike times leads to a substantial improvement of stimulus-dependent clustering.

The two cells were then analyzed jointly with the multi-unit metric, for a mesh of ordered pairs (q, k) , with the same values

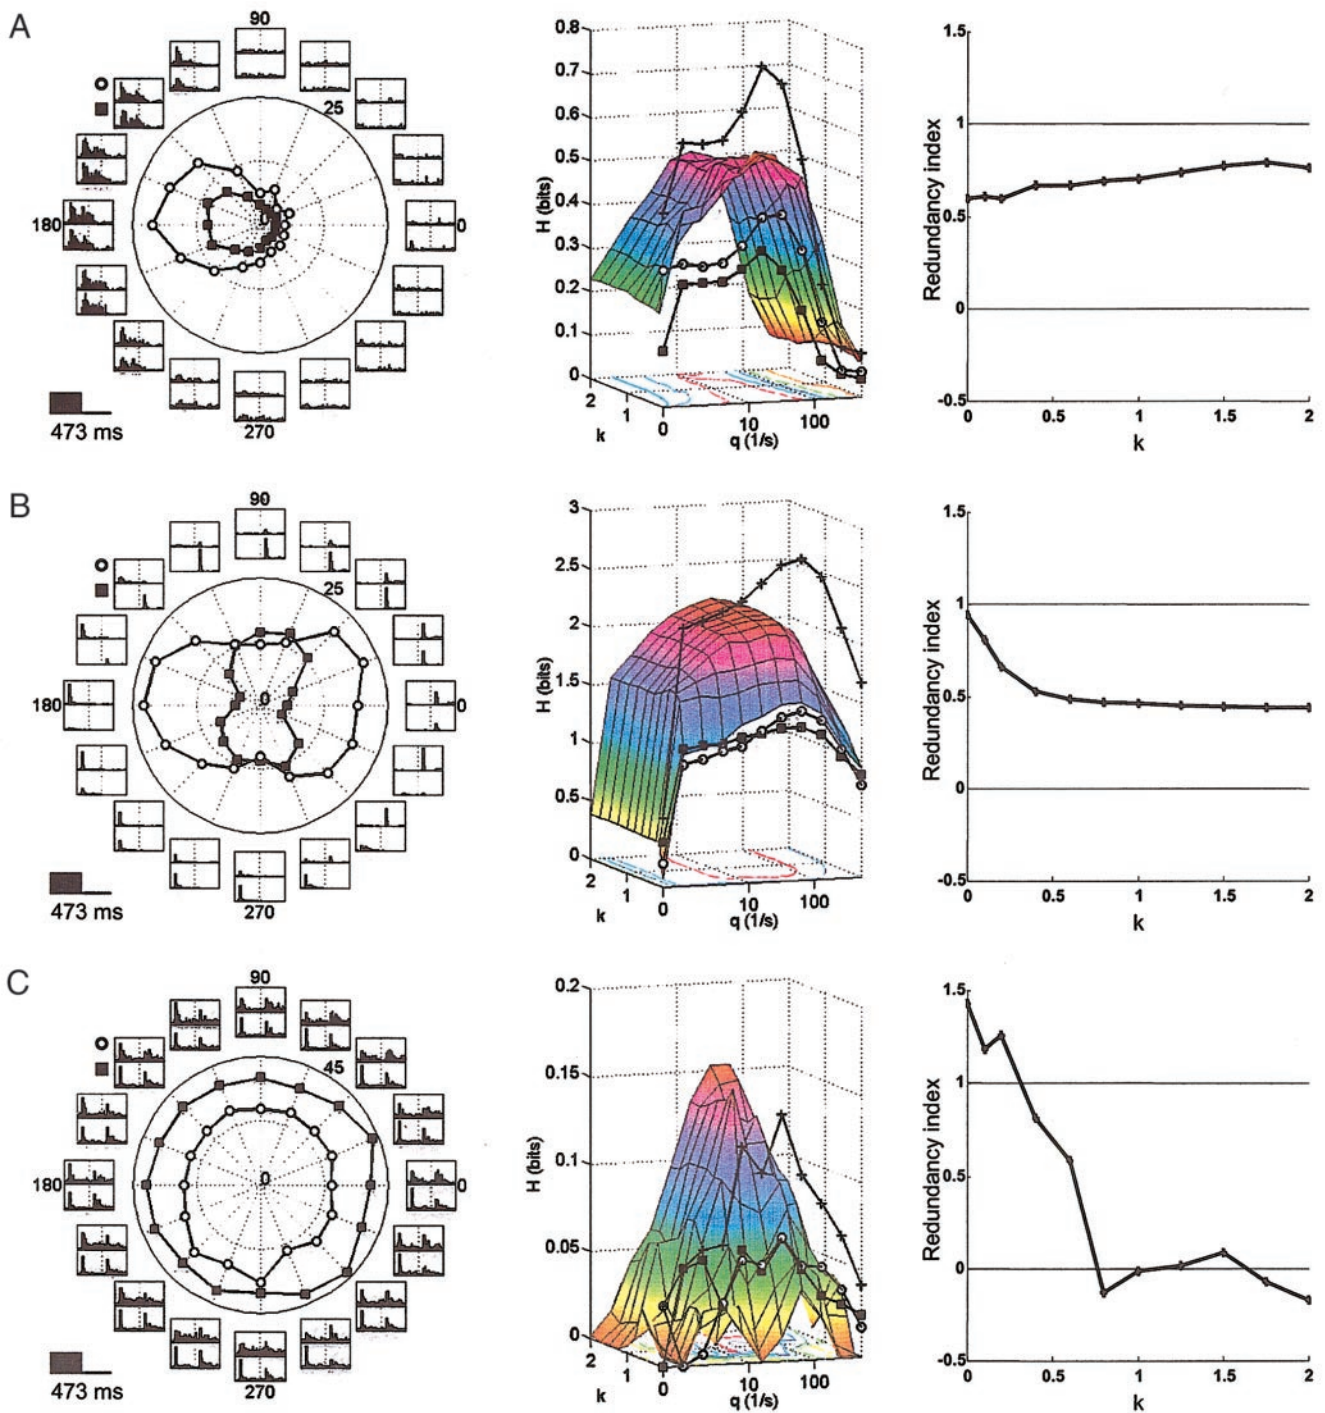


FIG. 2. Stimulus-dependent clustering analysis of the pairs of neurons of Fig. 1. *Left*: average responses of the 2 cells to gratings at 16 spatial phases. Histograms are scaled independently for the 2 cells. Symbols (circles and solid squares) to the left of the histograms identify each histogram with the corresponding polar plot and correspond to the symbols used to identify the units in Fig. 1. Vertical dotted lines at 237 ms mark the disappearance of the stimulus. Full scales on the histograms represent *A*, 76 spikes/s (circle) and 43 spikes/s (solid square); *B*, 353 spikes/s (circle) and 271 spikes/s (solid square); *C*, 115 spikes/s (circle) and 237 spikes/s (solid square). Polar plot: average firing rates of the 2 cells throughout the entire 473 ms response; rate corresponding to outer circle is 25 spikes/s (*A* and *B*), 45 spikes/s (*C*). *Middle*: analysis of stimulus-dependent clustering. Circles and solid squares: H for the individual cells (plotted with the corresponding symbol in *left*) as a function of q ; +: sum of the individual values of H . Surface: H_{joint} for the pair of cells, plotted as a function of q and k . *Right*: redundancy index (Eq. 4), plotted as a function of k .

of q as in the preceding text and k ranging from 0 to 2, sampled more finely at lower values. The index of joint response clustering (denoted by H_{joint} to distinguish it from calculations based on single units) is plotted as a function of these param-

eters as the surface in Fig. 2*A*. For constant values of k , H_{joint} depends on the timing parameter q in a way similar to H for the single-unit responses, achieving a maximum for q in the range 16–32 s⁻¹. In contrast, H_{joint} has very little dependence on k

for any value of q . This indicates that, in this data set, distinguishing spikes that are fired by different neurons has no effect on the strength of stimulus-dependent clustering. That is, for the purpose of determining spatial phase from these neurons' responses, it suffices to consider them indistinguishable members of a population.

The preceding observation does not mean, however, that the two neurons are redundant—merely that they are no less informative if one ignores which neuron fired which spike. Indeed, if the contributions of cells to stimulus-dependent clustering were completely redundant, the values of H_{joint} would be no greater than the single-unit values of H . Figure 2A thus shows that the responses are not completely redundant: at constant q , H_{joint} is greater than the either of the single-unit values of H across the entire range of values for k . On the other hand, if there were *no* redundancy, the values of H_{joint} would be equal to the *sum* of the single-unit values of H . It does not achieve this value (curve indicated by '+'s) at any value of q and thus the contributions of the two cells to stimulus-dependent clustering are *partially* redundant.

Redundancy index

To quantify the degree of independence of the single neuron responses, we used a redundancy index (Reich 2001b). The redundancy index is given by

$$\text{RI} = \left(1 - \frac{H_{\text{joint}}}{H_1 + H_2} \right) / \left(1 - \frac{\max\{H_1, H_2\}}{H_1 + H_2} \right) \quad (4)$$

where H_1 and H_2 are the clustering indices H derived from the two neurons considered separately. For each value of k , clustering indices in Eq. 4 are measured at the optimal value of q . The redundancy index is 0 when cells contribute entirely independent information ($H_{\text{joint}} = H_1 + H_2$) and 1 when the cells are completely redundant ($H_{\text{joint}} = \max\{H_1, H_2\}$). The index can be >1 if individual contributions of stimulus-related information are contradictory or confusing or <0 if the neurons code synergistically.

The redundancy index for this data set is plotted as a function of k in Fig. 2A (right). When analysis is restricted to $k = 0$, the redundancy index is 0.60. At the optimal value of H_{joint} , the redundancy index is the same (0.60 at $k = 0.2$). Thus for this pair of cells, the two neurons together provide more information about which spatial phase was present than either neuron alone, but there is no increase in the fidelity of the phase representation associated with keeping track of which cell fired which spike.

Sample data set 2: pair of simple cells with dissimilar phase preferences

Responses of two simple cells with dissimilar phase preferences are shown in Fig. 2B (left). Each cell discharges a transient ON response to gratings in a range of spatial phases and a transient OFF response to gratings at the opposite phases. The cell denoted by circles has the largest ON responses at spatial phases 135–225° and the largest OFF responses at 315–45°. The cell denoted by squares has an ON and OFF response at spatial phases 225–315° and 45–135°, respectively, which are nearly orthogonal to the preferred phases of the first neuron.

Independent analyses of the two cells with the single-unit

metric yield results (curves in Fig. 2B, middle) similar to those described for the data set in Fig. 2A. Stimulus-dependent clustering of both cells' responses is greatly improved when their temporal structure is taken into consideration. Maximal clustering is achieved for q in the region of 16–64 s⁻¹. The maximal values of H are 1.32 bits (the cell denoted by squares) and 1.46 bits (the cell denoted by circles). These values are atypically high, being more than three times greater than the average across simple cells. Nevertheless, they are only $\sim 1/3$ the value expected of perfect clustering ($\log_2 16 = 4$).

The information surface from joint analysis of the same cells (Fig. 2B, middle) shows a similar dependence of response clustering on q as the surface from the data set in Fig. 2A. However, in contrast to measurements from pairs of cells with similar phase preferences, the values of H_{joint} in the present data set are strongly dependent on k . At $k = 0$, the values of H_{joint} are not significantly higher than the values of H derived from single-unit responses, whereas at near-optimal values of q , H_{joint} increases greatly with k . The maximum value of H_{joint} (as a function of q) levels off at approximately $k = 0.5$ but continues to rise slightly until $k = 2$. When the joint analysis is limited to $k = 0$, the redundancy index for this pair of cells is 0.94 (Fig. 2B, right). For $k > 0$, however, it decreases to a much lower value of 0.44 (Fig. 2B, right), indicating that distinguishing spikes fired by different neurons improves stimulus-dependent clustering in the present data set. If neurons are not distinguished ($k = 0$), the addition of the second cell does not significantly increase the amount of transmitted information, corresponding to a redundancy index of 1. That is, if responses are simply pooled, the benefits that might result from reduced noise (due to independent contributions from each neuron) are offset by the penalty of combining responses with distinguishable phase preferences.

Sample data set 3: pair of complex cells

Responses of two simultaneously recorded complex cells are shown in Fig. 2C (left). Each cell has prominent ON and OFF discharges with both transient and sustained components at all spatial phases. The average firing rates of both cells do not vary significantly with phase.

As seen from the values of H in Fig. 2C (middle), stimulus-dependent clustering of single-unit responses is considerably weaker for the present data set than for simple cells, as would be expected from the qualitative phase-insensitivity seen in the response histograms (Fig. 2C, left). For both neurons, the amount of transmitted information is near chance (0 bits) at $q = 0$ s⁻¹, corresponding to the observation that the number of spikes in the cells' responses did not vary systematically with spatial phase. At positive values of q , however, H becomes significantly higher than chance and reaches a maximum at q in the region of 32 s⁻¹. Thus a modest amount of information about the spatial phase of the stimulus is encoded in the temporal structure of responses, even though the total number of spikes carries no information. As in the previous examples, the peak value of q indicates that the spike times are informative on a time scale of $\sim 1/q = 30$ ms.

In contrast to the previous data sets, analysis of joint responses reveals a strong dependence of clustering on k (Fig. 2C, middle). At $k = 0$, the values of H_{joint} are lower than the values of H derived from individual responses. Thus simple

addition of the two responses confounds the coding of spatial phase. This corresponds to a redundancy index >1 (Fig. 2C, right). At near-optimal values of q , however, H_{joint} increases with k and reaches values that are even higher than the sum of the individual measurements. This corresponds to synergistic coding of spatial phase, and a redundancy index of <0 (Fig. 2C, right). The generally low values of H , especially at $q = 0 \text{ s}^{-1}$, and the increase in H_{joint} for $k > 0$ were typical of the recordings of pairs of complex cells. However, synergistic coding, clearly demonstrated by these two cells, was not typical in our data sets.

Summary across data sets

Figure 3 summarizes the clustering analysis across the data sets. For individual cells that constituted the analyzed pairs (22 simple and 18 complex), the average behavior of H is shown in Fig. 3A. The dependence of the peak value of H on the F_1/F_0 modulation ratio for these individual neurons is shown in Fig. 3B, and characteristics of the stimulus-dependent clustering analysis are summarized in the first two columns of Table 2. At all values of the metric parameter q , individual responses of simple cells exhibited stronger stimulus-dependent clustering than did those of complex cells (Fig. 3A). At $q = 0 \text{ s}^{-1}$ (the spike count metric), the values of H for complex cells were usually not above chance level, while for simple cells they were usually highly significant (Table 2). This is consistent with the notion that response magnitude is phase dependent in simple cells and phase independent in complex cells. However, most simple and complex cells yielded values of H that were above chance level when spike timing was taken into consideration, consistent with our previous findings. Maximal clustering occurred at $q > 0 \text{ s}^{-1}$ for 20 of the simple cells (91%) and 16 of the complex cells (89%). The average value of H at optimal q was 2.2 times higher than the value at $q = 0 \text{ s}^{-1}$ for simple cells and 2.9 times higher for complex cells. Thus temporal coding can allow both simple and complex cells to transmit more than twice the amount of information about spatial phase than is contained in spike counts. Qualitatively, H depended on q in a similar way for the two classes of cells. Maximal clustering was achieved, on average (geometric mean), at $q = 24 \text{ s}^{-1}$ for simple and 26 s^{-1} for complex cells, corresponding to a temporal precision of $\sim 40 \text{ ms}$.

It might appear surprising that modulation ratio and H were not tightly linked. In principle, eye movements might artifactually reduce the measured value of H . However, the rasters of Fig. 1 show that phase tuning was stable over the course of an experiment. This rules out the possibility that eye movements affect our measurements. (This was a data selection criterion.) More likely, modulation ratio (or phase tuning) and H are not closely linked because H depends critically on the signal-to-noise ratio, whereas modulation ratio and phase tuning depend primarily on signal. There are additional reasons that information values can be low even for a narrowly tuned cell with high a signal-to-noise ratio. If a threshold limits responses to only one portion of the spatial phase gamut, it will necessarily reduce the number of spatial phases that can be distinguished on the basis of the neuronal response. Thus paradoxically, responses of a more narrowly tuned neuron can contain less information about spatial phase. Comparison of Fig. 2, A and B, show an example of this. The neurons of Fig. 2A are

narrowly tuned, responding well to only five or six of the stimuli, and have values of H of ~ 0.4 bits. The neurons of Fig. 2B respond to a broader range of phases and have values of H of >1 bit.

Although there was a dramatic difference between the average values of H across simple and complex cells (Fig. 3A and Table 2), this difference does not necessarily mean that simple cells, as a class, signal spatial phase, while complex cells do not. As seen in Fig. 3B, the range of peak values of H for simple and complex cells was overlapping. Two other aspects of this figure indicate that neurons cannot be classified as simple and complex based on the amount of transmitted information about spatial phase: some simple cells have low values of H and the maximal value of H covaries with the modulation ratio, not the classification (simple vs. complex) of the neuron per se. That is, within the cells of either category, higher modulation ratios are associated with larger values of H ; the classification cutoff at a modulation ratio of 1 plays no special role. These observations are in keeping with findings concerning the selectivity of responses of simple and complex cells to the relative spatial phases of compound gratings.

The average behaviors of H_{joint} for the 11 pairs of simple cells and 9 pairs of complex cells are shown in Fig. 3, C and D, respectively. For the joint responses, the dependence of clustering on temporal resolution was similar to that for the individual responses. That is, the qualitative dependence of H_{joint} on q was the same at all values of k , indicating that the temporal structure of responses contributes to stimulus-dependent clustering independently of the neuron of origin. For a pair of simple cells, the information transmitted when decoded as a summed population (0.56, “ q optimal, $k = 0$ ” in Table 2) was less than when decoded in a manner that was sensitive to which neuron fired which spike (0.69, “ q optimal, k optimal”). This improvement corresponds to a drop in the average redundancy index (Fig. 3E) from ~ 0.8 at $k = 0$ to 0.6 at optimal k .

For complex cells, the improvement of stimulus-dependent clustering for optimal k (typically near 1) was greater than for simple cells (Fig. 3D). While clearly less than the amount of information transmitted by pairs of simple cells about spatial phase, the amount of information (0.13, “ q optimal, k optimal” in Table 2) transmitted by a pair of complex cells is not negligible. To put this quantity in perspective, this amount of information supports a performance level of 70% correct on a two-alternative forced choice task.

The behavior of the redundancy index in complex cells indicates diversity across this subpopulation. On average (Fig. 3E), the redundancy index declines from 1.70 at $k = 0$ to just under 1 at optimal k . This reflects an admixture of two kinds of behavior: pairs such as the one in Fig. 2C, in which there is very little redundancy at sufficiently high k , and other pairs, for which the redundancy index remains high even for large k . The latter group includes pairs in which the maximal value of H for one of the neurons is near zero. This leads to high values of the redundancy index as long as $H_{\text{joint}} < H_1 + H_2$.

One would expect that the importance of the neuron of origin of individual spikes would be larger for pairs of cells with dissimilar phase preferences than for pairs of cells with similar phase preferences. For simple cells, this intuition is confirmed (Fig. 4A). We quantified the importance of neuron of origin (i.e., of which neuron fired which spike) by the drop in the redundancy index, calculated between $k = 0$ and optimal k .

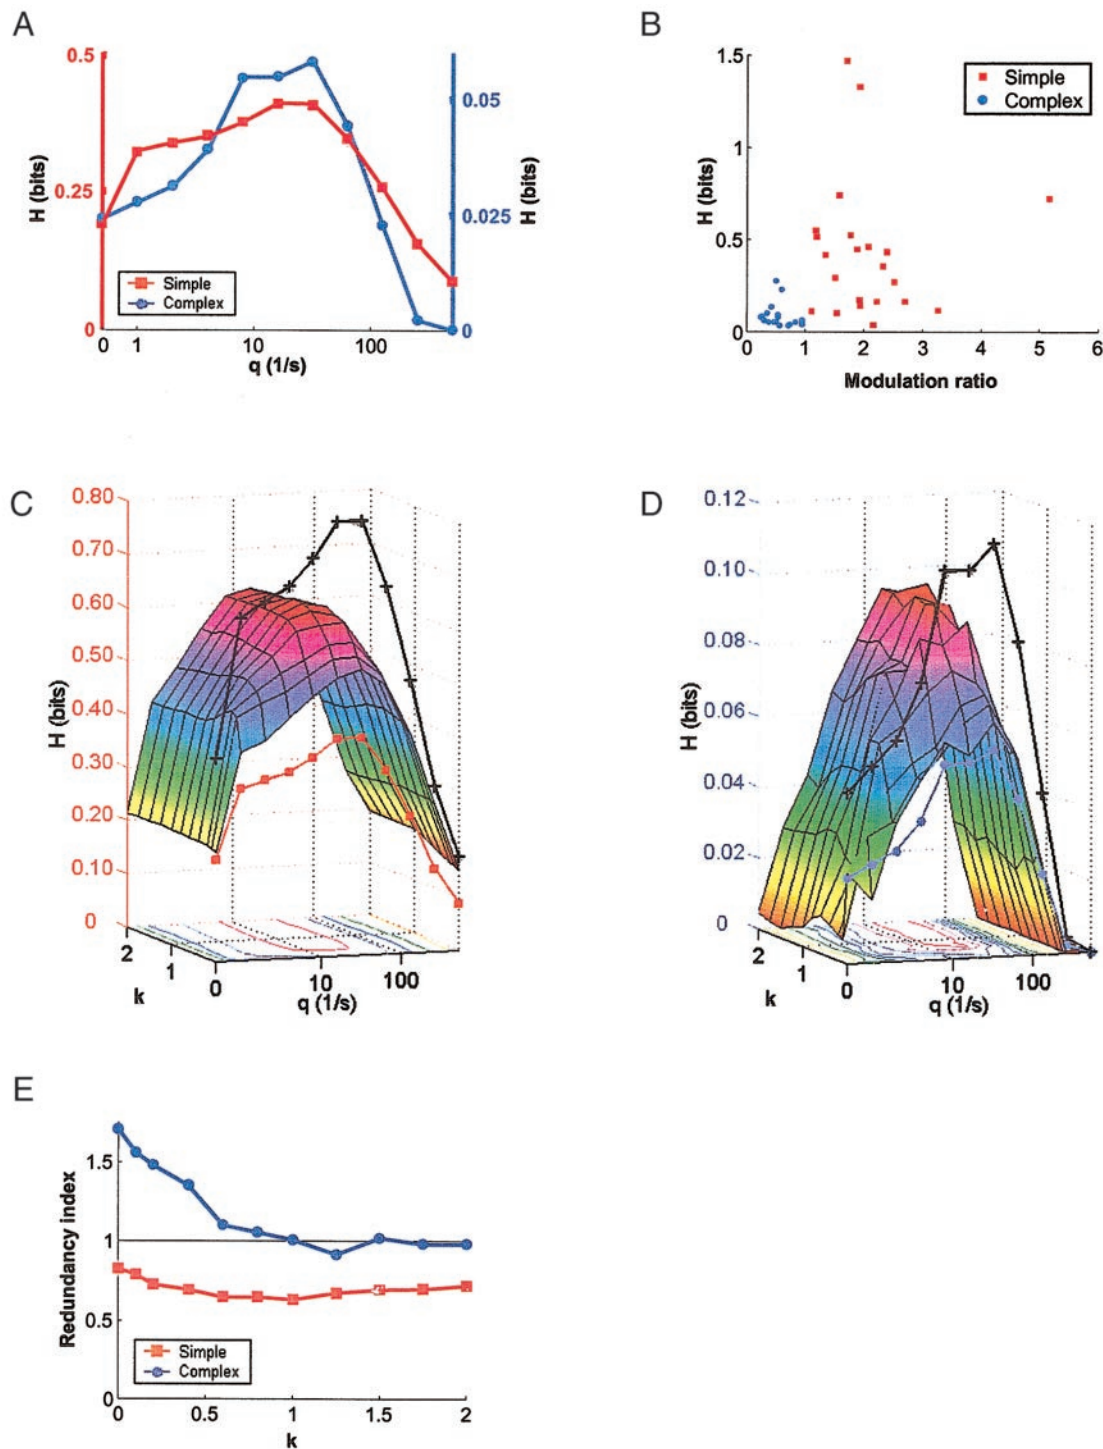


FIG. 3. Summary of stimulus-dependent clustering across data sets. *A*: average values of H across all simple ($n = 22$, vertical scale on the left) and complex cells ($n = 18$, vertical scale on the right), plotted as functions of q . *B*: scatter plot of H vs. the modulation ratio (F_1/F_0 , see METHODS) for all individual neurons ($n = 40$) included in the stimulus-dependent clustering analysis. *C*: ■: H averaged across individual simple cells, plotted as a function of q ; +: average sum of individual values of H for pairs of simple cells (twice the average values of H). Surface: H_{joint} averaged across pairs of simple cells ($n = 11$), plotted as a function of q and k . *D*: average graphs for complex cells plotted as in *C* ($n = 9$). *E*: redundancy indices averaged across all simple and complex cells, plotted as functions of k .

We correlated this value (on the ordinate) with a single parameter r_{tuning} (on the abscissa) that quantifies the similarity of phase preferences. r_{tuning} was obtained as follows. We took the average firing rate over the entire 473-ms interval as the

response measure at each spatial phase and set r_{tuning} equal to the standard (Pearson) correlation coefficient of the cells' responses across the 16 spatial phases. Three pairs of simple cells had negative values of r_{tuning} , indicating largely opposite phase

TABLE 2. Summary of stimulus-dependent clustering

Class	Individual Cells		Pairs of Cells			
	$q = 0$	q optimal	$q = 0, k = 0$	$q = 0, k$ optimal	q optimal, $k = 0$	q optimal, k optimal
S	0.19 (91%)	0.43 (100%)	0.20 (82%)	0.26 (91%)	0.56 (100%)	0.69 (100%)
C	0.02 (6%)	0.08 (72%)	0.03 (33%)	0.03 (33%)	0.08 (89%)	0.13 (100%)

Entries specify the average clustering index (amount of transmitted information in bits) for individual cells, H , and for pairs of cells, H_{joint} . Percentages in parentheses indicate the fraction of these data sets for which the measurement was above chance level. Note that the values under the headings “ q optimal” and “ k optimal” reflect the average of the peak values of the H curves and the H_{joint} surfaces obtained from each cell (or cell pair) individually, not the peak of the average values (as displayed in Fig. 3, A, C, and D).

preferences. These pairs were also the ones with the largest differences (>0.5) between the two redundancy indices (e.g., data set in Fig. 2B). Conversely, for the four pairs of simple cells with the most similar phase preferences ($r_{\text{tuning}} > 0.97$),

the difference between the two redundancy indices was <0.005 (e.g., data set in Fig. 2A). For complex cells, there was a greater range of the redundancy index, but no clear correlation with phase tuning similarity—most likely because spike counts poorly reflect the phase tuning of complex cells, and because measurement of the redundancy index is less reliable when values of H are small.

Figure 4A suggests that simultaneously recorded neurons tend to have similar phase preferences (i.e., r_{tuning} tends to be close to 1 rather than randomly distributed between -1 and 1) but is limited to the randomly selected 20 pairs of cells (11 simple-simple, 9 complex-complex) for which we performed the stimulus-dependent clustering analysis parametric in q and k . Figure 4B shows r_{tuning} for all pairs of simultaneously recorded neurons that met the selection criteria (29 simple-simple pairs, 11 simple-complex pairs, and 17 complex-complex pairs; see Table 1). The tendency for the phase tuning of neighboring neurons to be similar ($r_{\text{tuning}} > 0$) is more clearly evident, especially for simple-simple pairs (mean $r_{\text{tuning}} = 0.68$) but also for simple-complex (mean $r_{\text{tuning}} = 0.29$) and complex-complex (mean $r_{\text{tuning}} = 0.30$) pairs. Also, there is no evident tendency for a quadrature relationship between neighboring neurons, which would correspond to $r_{\text{tuning}} = 0$.

Geometry of response clusters

The analysis so far focused on the extent to which the responses to the 16 different spatial phases fell into distinguishable clusters but provided no insight into the relationship between these clusters. For example, this analysis could yield identical information (H) values for a neuron that primarily confounded each spatial phase with its opposite (an idealized ON-OFF neuron with little noise) and for a neuron that confounded each spatial phase with its neighbor (an idealized linear neuron with some noise). Yet these distinctions are important because they can suggest the kinds of receptive-field mechanisms that distinguish spatial phases. Additionally, even a high degree of clustering would be useless in representing spatial phase unless the responses varied with spatial phase in some systematic manner.

The first step in characterizing the geometry of the response clusters was multidimensional scaling of metric distances between single-trial responses. For single neurons, we used the single-unit metric with $q = 32 \text{ s}^{-1}$; for joint responses of simultaneously recorded neurons, we used the multi-unit metric with $q = 32 \text{ s}^{-1}$ and $k = 1$. These choices were close to maximal across the entire population of analyzed neurons. We chose a uniform value of q and k for this analysis so that we would not confound differences between recordings with the dependence of the analysis of individual recordings on q and k .

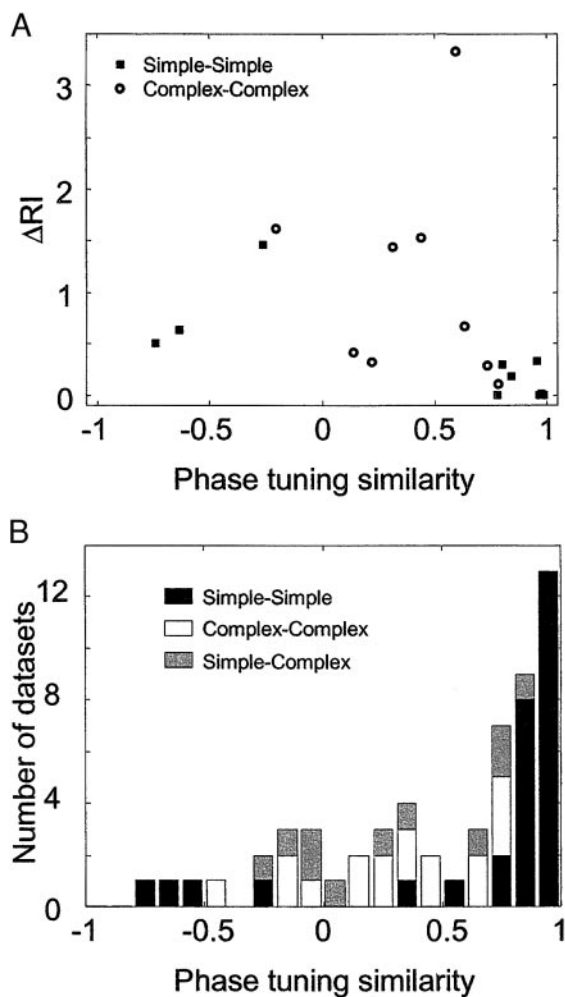


FIG. 4. A: The relationship between the extent to which neuron of origin reduces redundancy (ordinate), and similarity of phase tuning of nearby neurons, as calculated from 20 randomly selected pairs of neurons (11 simple-simple, 9 complex-complex) that met the selection criterion. The ordinate indicates the fractional reduction in the redundancy index calculated at $k = 0$ and the redundancy index calculated at the maximum value of H_{joint} . The abscissa indicates the correlation coefficient r_{tuning} of 2 neighboring cells' tuning functions, determined from spike counts as described in the text. Data from 4 separate simple-simple pairs coincidentally overlap as a single filled rectangle in the lower right-hand corner of the plot. B: the distribution of similarity of phase tuning among all 57 pairs of simultaneously recorded neurons that met the selection criterion, as measured by the correlation coefficient r_{tuning} . The 20 pairs of A are included in this histogram.

The details of the results of multidimensional scaling do change with the metric parameters, but this change is gradual, and our main conclusions do not depend on the specific choices of q and k within the range $q = 8\text{--}128\text{ s}^{-1}$ and all $k > 0.5$.

After embedding of the responses in a 10-dimensional Euclidean space by multidimensional scaling, we characterized the geometry of the response clustering by fitting ellipses to the centroids of responses to each spatial phase. Separable receptive fields predict a response trajectory that doubly covers a line segment (moving sinusoidally as a function of spatial phase), while linear but inseparable receptive fields predict an elliptical trajectory with a nonzero minor axis (see METHODS and APPENDIX). Therefore it is important to determine whether the fitting of a response trajectory by an ellipse is significantly better than a fit with a doubly covered line segment, i.e., an ellipse with an axis ratio of zero.

Our approach to estimating significance of the minor axis was the following. The null hypothesis is that the apparent minor axis is no longer than what might be expected due to chance. We generated surrogate trajectories by reflecting randomly chosen subsets of the 16 centroids over the major axis of the ellipse—i.e., negating their coordinate along the direction of the minor axis but leaving their coordinate along the direction of the major axis unchanged. (All transformations were made parallel to the plane of the best-fit ellipse and only involved coordinate changes along the minor axis.) Under the null hypothesis that the centroids were positioned at random along the *minor* axis but perhaps systematically along the *major* axis, these surrogates would be just as probable as the

observed data. The best-fit ellipse was then recalculated for each of 1,000 such sets of surrogate centroids, and the in-plane variance explained by the best-fit ellipse was tabulated. If the original centroids were positioned at random along the minor axis, then random reflection would be just as likely to improve the goodness of the elliptical fit as to worsen it. If, on the other hand, centroids were positioned systematically in two dimensions around an ellipse, reflections would typically decrease the goodness of fit. We defined $p_{\text{line}}^{\text{seg}}$ as the fraction of the surrogate sets for which ellipses explained more of the variance than for the original trajectory. Data sets with $p_{\text{line}}^{\text{seg}} < 0.05$ were thus fit by ellipses significantly better than by a doubly covered line segment. For some data sets, ellipses accounted for the centroid positions significantly better than doubly covered line segments (by this measure), yet both of these fits accounted for only a small fraction of the variance. Because of this possibility, we considered an elliptical fit to be significant only if it passed the above test for systematic arrangement in two dimensions and it explained $>50\%$ of the total variance.

Sample data set 1

In Fig. 5, we present this analysis of the same three data sets shown in Fig. 2. Figure 5A, *top* and *middle*, shows the first two dimensions of the embedding provided by multidimensional scaling of the individual responses illustrated in Fig. 2A (simple cells with similar phase tuning). Figure 5A, *bottom*, shows the first two dimensions of the embedding of the joint responses. In each case, along the first dimension, the projections

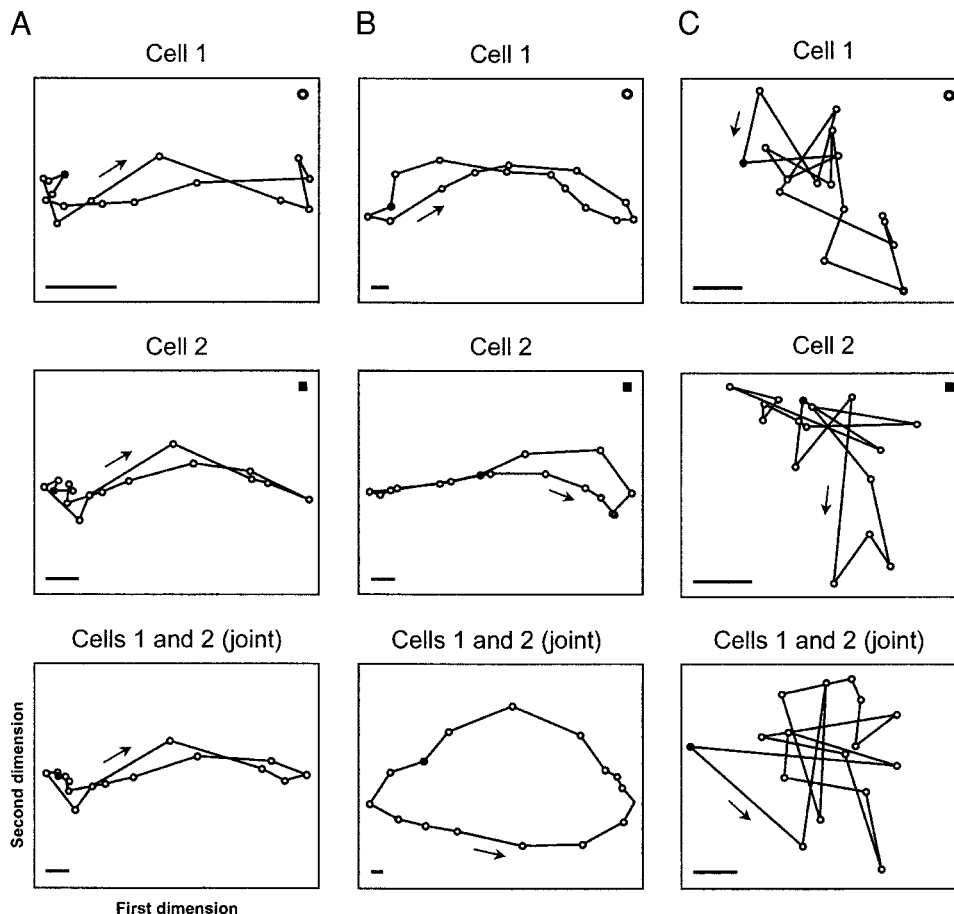


FIG. 5. Multidimensional scaling of responses for the 3 data sets in Fig. 2. Closed curves are projections of the individual response trajectories onto the plane of the 1st 2 dimensions. Variation along the horizontal axis corresponds to the 1st dimension and the vertical to the 2nd. Circles indicate centroids of responses corresponding to each of the 16 spatial phases; the single solid circle in each trajectory indicates the centroid corresponding to 0° . Arrows indicate the direction of increasing phase. Scale bars: 1 unit of distance in the metric spaces. *Top* and *middle*: single-unit data; symbols in the top-right corners of the plots identify each trajectory with the corresponding cell in Fig. 2. *Bottom*: joint responses.

TABLE 3. Summary of response trajectory shapes as analyzed by ellipse fitting

	Individual Cells		Pairs of Cells		
	S	C	S-S	S-C	C-C
Axis ratio	0.14 (39)	0.39 (38)	0.18 (52)	0.33 (55)	0.46 (41)
Variance explained, %	82	66	81	81	65

Entries are averages across the data sets indicated by the column header. For axis ratio, the number in parentheses indicates the percentage of these data sets for which the elliptical fit was significantly different from a doubly covered line segment and explained >50% of the variance.

appear to be similar to doubly covered line segments. Correspondingly, all best-fit ellipses are highly eccentric, with axis ratios equal to 0.036 for the cell denoted by a circle, 0.049 for the cell denoted by a solid square, and 0.026 for the pair of cells analyzed jointly. These ellipses account for 81, 79, and 83% of the variance, respectively. All three ellipses yield high values of the preceding statistic p_{lineseg} , indicating that along the second dimension, the trajectories are no more consistent with ellipses than would be expected by chance variation from a doubly covered line segment.

Sample data set 2

Figure 5B shows the same analysis for the cells of Fig. 2B (simple cells with distinct phase preferences). Unlike the single-unit response trajectories in Fig. 5A, the trajectories in B appear to vary systematically along the axis of the second dimension. This is even more evident for the joint response trajectory, which is nearly an ellipse. In the plane of the first two dimensions, single-unit trajectories appear to be ellipses that have been twisted around the major axis. Best-fit ellipses to these trajectories have axis ratios equal to 0.075 for the cell denoted by a circle and 0.127 for the cell denoted by a solid square. These ellipses explain 92 and 85% of the variance and are both significant, with $p_{\text{lineseg}} < 0.01$ and $p_{\text{lineseg}} < 0.001$, respectively. The joint response trajectory has a much higher axis ratio of 0.425 and explains 88% of the variance with $p_{\text{lineseg}} < 0.001$.

As we will see in the following text, the first dimension in the multidimensional space indicates the difference between the ON and OFF responses (see *Temporal profiles*). Centroids move in one direction along the first dimension as the ON response increases and in the other direction as the OFF response increases. If the number of spikes in the ON and OFF discharges were the only attribute of the response that varied systematically with spatial phase, one would expect significant systematic variation of the responses only along this dimension, resulting in doubly covered line segments as response trajectories. The coding analysis (Figs. 2 and 3) demonstrates that temporal structure, and not just response magnitude, varies systematically with spatial phase. In multidimensional scaling of the second data set, centroids of responses with equal magnitudes of the ON and OFF components are separated along the second dimension. This must be due to differences in the temporal structure of these responses and not just changes in response magnitude. That is, temporal coding must underlie a systematic arrangement of the responses in a second dimension of the response space.

Sample data set 3

Figure 5C shows multidimensional scaling of the cells of Fig. 2C (2 complex cells). Although responses to certain spatial phases seem to lie apart from the bulk of the responses, the trajectories of both single- and multi-unit responses are not similar to ellipses. Indeed, the best-fit ellipses to these trajectories explain 49 and 50% of the variance for the cells denoted by a circle and a solid square, respectively, and 48% of the variance for the joint responses.

Summary across data sets

Results of elliptical fitting are summarized in Table 3. The first two columns indicate that there were significant quantitative differences between simple and complex cells. For complex cells, ellipses explained a smaller fraction of the variance than for simple cells, but their axis ratios tended to be larger. The former difference is expected from the fact that responses of complex cells vary less systematically with spatial phase than responses of simple cells. The latter difference does not have a clear explanation on this basis. As pointed out in the preceding text, any two-dimensional trajectory requires two (or more) spatial mechanisms, each coupled to distinct temporal responses. The greater axis ratio in complex cells suggests that the spatial and temporal distinctions between these mechanisms were greater in complex than in simple cells.

The last three columns focus on the joint responses of pairs of cells. The fraction of data sets with response trajectories that were significantly elliptical was larger for pairs of cells analyzed jointly than for individual cells. The axis ratios of the ellipses were also higher for the joint responses. That is, their trajectories were more nearly circular. This difference was substantial for data sets such as the one in Fig. 5B in which the phase tunings of the individual neurons were very different. However, for most data sets, the average axis ratios were only slightly higher for pairs of cells than for individual neurons.

The goodness of the elliptical fit was similar for single neurons' responses and for joint responses. As seen in Fig. 6, in both cases the goodness of fit was strongly correlated with the amount of spatial phase information in the responses as measured by stimulus-dependent clustering. Response trajectories that strongly deviated from ellipses were produced pri-

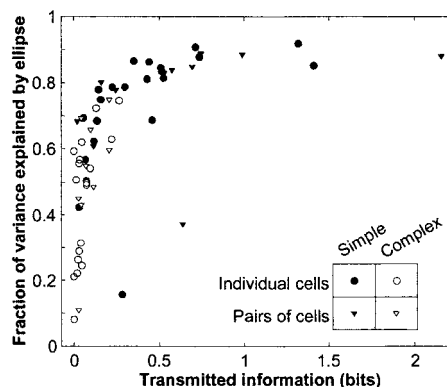


FIG. 6. Correlation of the goodness of the elliptical fit with the amount of stimulus-dependent clustering. Transmitted information is 1 of the 2 indices of stimulus-dependent clustering: H for single-unit responses or H_{joint} for joint responses of pairs of cells. Measurements of the variance explained by the elliptical fit were made in the plane of the best-fit ellipse.

marily by those data sets that exhibited weak stimulus-dependent clustering. As H and H_{joint} increased, the fraction of response variance explained by ellipses increased almost to 1. This indicates that, for the neurons that produced the most discriminable responses to spatial phases, the ellipse is indeed the geometry that expresses how these responses depend on spatial phase.

Temporal profiles

So far, we have shown that both the number of spikes and the temporal structure of responses contribute significantly to stimulus-dependent clustering and that these response clusters often result in a systematic geometric representation of spatial phase in two dimensions of a response space. Inspection of response waveforms suggests that these two coordinates represent response size and some aspects of the temporal structure. Here we use temporal profiles (see METHODS) to make this observation more precise, by determining relative contributions of various sections of spike trains to the separation of responses along each dimension.

To further clarify the rationale behind this analysis, we need to consider what the analysis so far has, and has not, revealed. The dependence of the strength of stimulus-dependent clustering (as measured by H) on the metric parameter q indicates how much a spike can be shifted in time ($1/q$) before the response is less recognizable as being elicited by the same stimulus. For analysis of pair recordings, the dependence of the strength of stimulus-dependent clustering H_{joint} on the metric parameter k indicates how important neuron of identity is in determining which spatial phase elicited a particular response. The response trajectories, calculated for a specified choice of q and k , indicate that pattern of “closeness” of responses forms a roughly elliptical response space, recapitulating the cyclic nature of spatial phase. But what this analysis does *not* reveal is *how* the temporal structure of the response carries phase information. One possibility consistent with our findings is that response latency carries phase information, much as it is known to carry contrast information (Gawne 2000; Gawne et al. 1996; Reich et al. 2001a). Another possibility is that the phase information is coded by a comparison between the number of spikes in one portion of the response and the number of spikes in another portion of the response. More generally, phase information might correspond to the similarity of the response profile to one or more templates. Without further analysis, the possibility also exists that two neurons might provide very similar response trajectories, yet use very different temporal representations of phase.

The extraction of temporal profiles sorts out these possibilities by providing an interpretation of the coordinates (or axes) of the response trajectories of Fig. 5. If temporal profiles can be found, then the similarity of a response to the temporal profiles (i.e., the dot-product of the response with the profiles) yields the coordinates of that response. In this sense, extraction of the templates, which we call temporal profiles, is similar to principal components analysis. However, there are important differences between our procedures and the principal components analysis of neural responses (McClurkin et al. 1991). Principal components analysis is predicated on the assignment of coordinates to each response via a *linear* process such as binning or Fourier analysis. Temporal profile analysis begins with the

distances between *pairs* of responses, calculated from the metrics, a procedure that does not assume linearity. Moreover, temporal information need not be carried by similarity to a response template. In such circumstances [e.g., involving the phase of a semi-regular spike train, shown in Fig. 9 of Victor and Purpura 1997)], the metric space approach will identify coding schemes that lead to lawful response trajectories, but no temporal profiles will exist (and principal components analysis of the responses will also fail). Thus the application of temporal profiles to the preceding response trajectories is intended to address several issues. Can the temporal information indeed be characterized by similarity to particular templates? If so, are these templates consistent from cell to cell. If so, what relationships do they bear to recognizable response features?

We continue to use the same three data sets as examples and base the analysis on the same metric parameters as for elliptical fitting ($q = 32 \text{ s}^{-1}$, $k = 1$). For the derivation of temporal profiles, we separate responses into 32 equal bins, each 14.8 ms long. This temporal resolution is roughly twice as fine as the temporal precision specified by the metric at $q = 32 \text{ s}^{-1}$.

Sample data set 1

Temporal profile analysis of the data set presented in Figs. 2A and 5A is shown in Fig. 7A. For the two cells analyzed separately (Fig. 7A, *top 4 graphs*), the temporal profiles associated with the first two dimensions account for nearly all of the variance within this plane (98 and 97%, respectively, for the cells denoted by \circ and \blacksquare). Moreover, the temporal profiles are strikingly similar in shape for the two cells. In each case, the first temporal profile (Fig. 7A, *left, top 2 graphs*) increases to a plateau that starts ~ 75 ms after the beginning of the response and stays approximately constant until the end of the ON response. This indicates that all spikes in the ON response beyond 75 ms are weighted equally (as if they were simply counted) in determining the first coordinate. After the disappearance of the stimulus, the first temporal profile decreases over a period of ~ 75 ms and then remains at approximately the same negative value. This indicates that spikes in the OFF response also contribute to the first dimension but have an opposing effect on the first coordinate compared with spikes in the ON response. The magnitude of the first temporal profile is larger in the ON response than in the OFF response, indicating that spikes in these two sections of the response do not contribute equally to the first dimension. In summary, the temporal profile analysis reveals that the first coordinate in the response space is essentially a weighted difference of the spike counts in the ON and OFF responses.

As we now detail, the second temporal profile (Fig. 7A, *right, top 2 graphs*) indicates that the second coordinate of a response is determined by its temporal structure rather than by spike counts. The second temporal profile starts at relatively high positive values. After an initial latency, it begins to decrease and continues to do so throughout the entire ON response, crossing zero at ~ 150 ms. This time corresponds approximately to the transition between the transient and sustained components of the response (see Fig. 2A). This sign reversal of the temporal profile indicates that spikes in the transient and sustained components of the ON response have opposing effects on the second coordinate. For the duration of the OFF response, the second temporal profile remains negative,

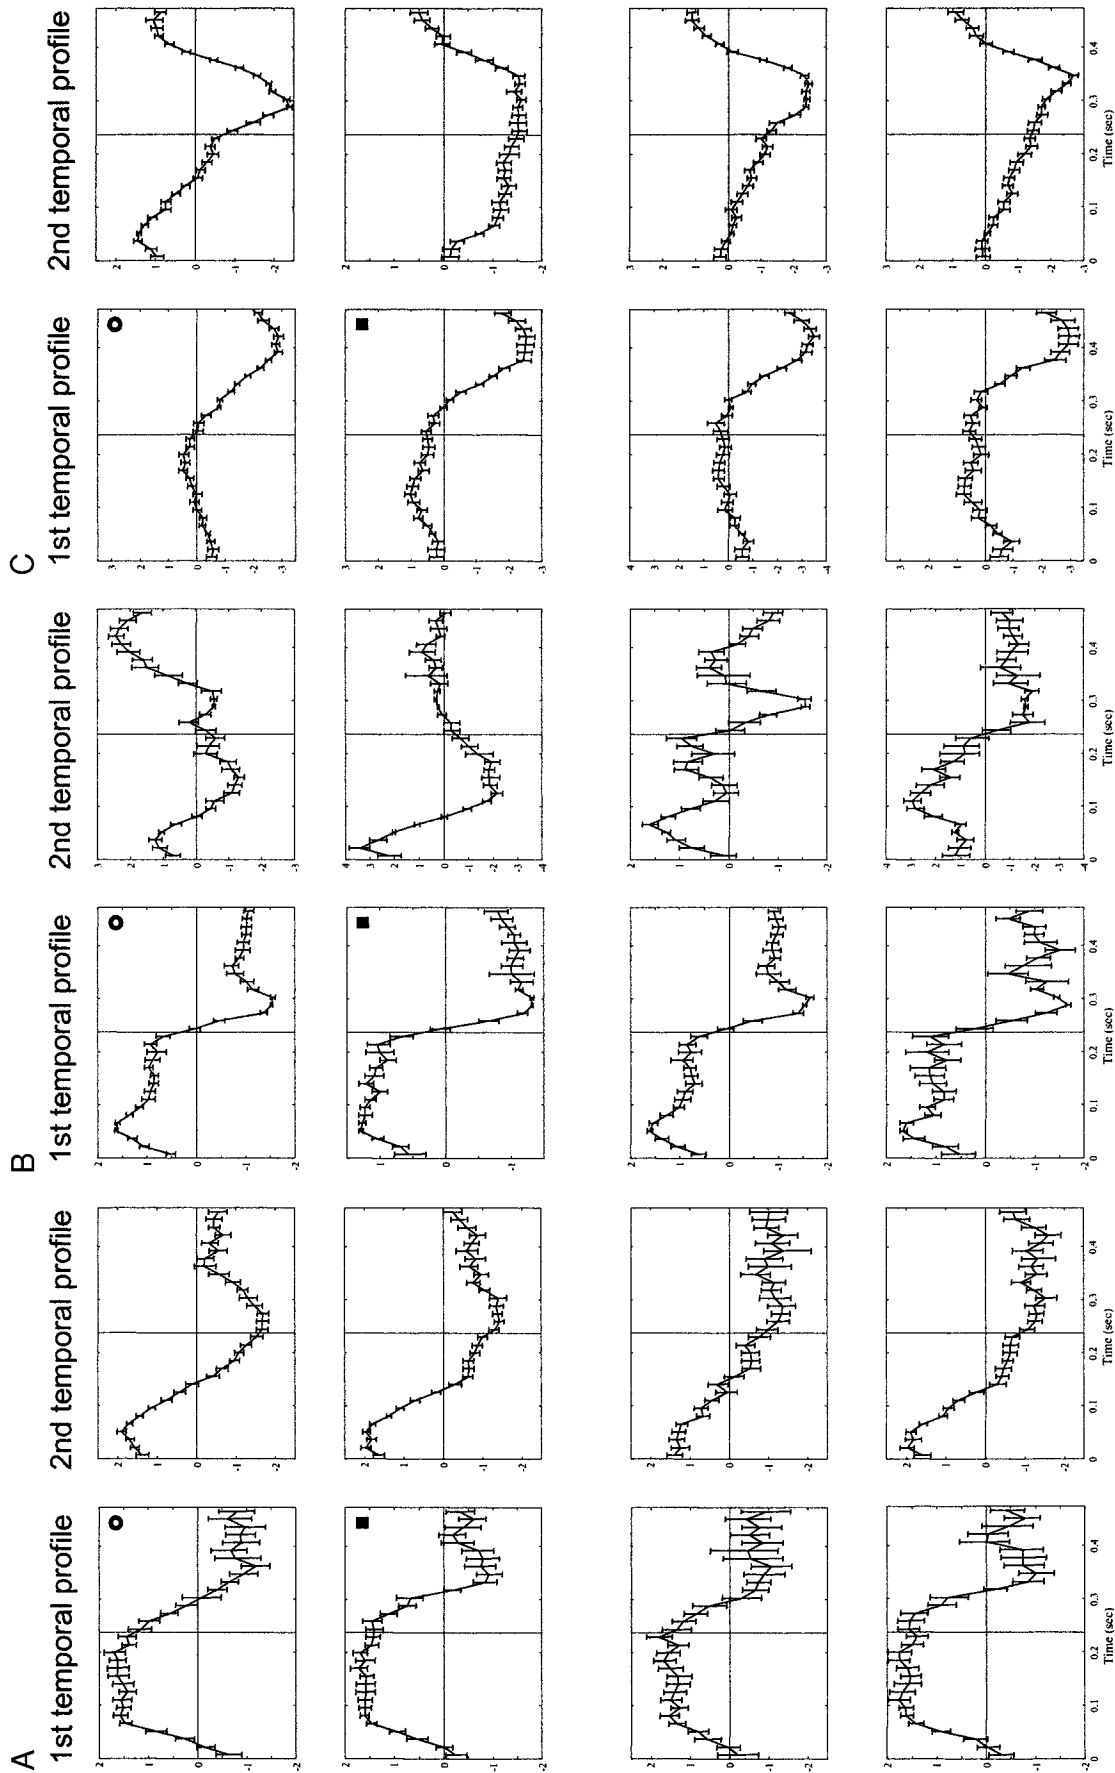


FIG. 7. Temporal profiles of responses for the 3 data sets in Fig. 2 and 5. Each temporal profile is normalized by its average magnitude. Error bars indicate 95% bootstrap confidence intervals from 1,000 resamplings. Vertical lines at 237-ms mark the disappearance of the stimulus. Symbols in the top-right corners of the plots identify rows of temporal profiles with the corresponding cells in Fig. 2 and 5. *Top 2 rows:* 1st and 2nd temporal profiles of the individual cells' responses. *Bottom 2 rows:* 1st and 2nd temporal profiles of the same cells analyzed jointly.

indicating that spikes during this period reinforce the sustained portion of the ON response in opposing the contributions of the ON transient.

Temporal profiles derived from the joint responses of the two cells are shown in Fig. 7A (*bottom 4 graphs*). These profiles explain 98% of the variance in the plane of the first two dimensions of the joint response space. The profiles are similar for the two cells, indicating that both cells contribute similarly to the geometric separation of the responses. Temporal profiles associated with each dimension are also similar for the individual and joint analyses. This means that the temporal structure of responses contributes to phase coding in the same way, whether responses are decoded one neuron at a time or jointly.

Sample data set 2

Temporal profile analysis of the data set presented in Figs. 2B and 5B is shown in Fig. 7B. For the two cells analyzed individually (Fig. 7B, *top 4 graphs*), the two temporal profiles again account for nearly all of the variance (99% for the cell denoted by ○ and 98% for the cell denoted by ■).

The first temporal profiles for both cells bear some similarities in shape to those in the previous data set (Fig. 7A). However, in the present data set, the ON and the OFF responses are weighted equally in determining the first coordinate. In addition, for the cell denoted by a circle, spikes in the response transient are weighted differently from spikes in the sustained response. This nonuniform weighting was uncommon for the first dimension across data sets.

Temporal profiles of the second dimension are also shown in Fig. 7B (*right, top 2 graphs*). As in Fig. 7A, they indicate opposing contributions of the transient and sustained components of the ON response. During the OFF discharge, however, these profiles behave differently. For the cell denoted by ○, transient and sustained components of the OFF response have opposing effects on the second coordinate with most of the weight carried by the sustained response. For the cell denoted by ■, the second profile stays essentially at the zero level, indicating that the OFF response of this cell has no effect on the second coordinate of the responses.

Temporal profiles of the joint responses of the two cells are shown in Fig. 7B (*bottom 4 graphs*). They explain 99% of the variance in the first two dimensions. The first temporal profiles of both cells are similar in shape to each other and to the profiles of the individual cells. The shapes of the second temporal profiles, however, have features that were not typical across data sets. For both cells, the ON and OFF components of the responses have opposing effects on the second dimension. Various sections of both components are weighted differently, but the transient and sustained parts of the responses do not have opposite effects, as they do for the ON responses in Fig. 7A.

Sample data set 3

Temporal profile analysis of the data set presented in Figs. 2C and 5C is shown in Fig. 7C. For the two cells analyzed individually (Fig. 7C, *top 4 graphs*), the two temporal profiles again account for nearly all of the variance (96% for the cell denoted by ○ and 98% for the cell denoted by ■). The first temporal profiles for both cells are very similar to each other.

As in the other two data sets, ON and OFF components contribute in opposing fashions, but in contrast to the other data sets, the weighting of the spikes during the ON response is much smaller than during the OFF response.

Temporal profiles for the second dimension are also shown in Fig. 7C (*right, top 2 graphs*). Unlike the profiles for the first dimension, the temporal profiles for the second dimension are not very similar. For the cell denoted by ○, there is a change in sign of the weights at the transition between the ON and OFF responses. Spikes within the transient and sustained components of the OFF response have opposing and unequally weighted effects. For the cell denoted by ■, there is no change in the weights at the transition between the ON and OFF components, but the first and second halves of the OFF response are weighted differently.

Temporal profiles of the joint responses of the two cells are shown in Fig. 7C (*bottom 4 graphs*). These profiles explain 98% of the variance in the first two dimensions. The first temporal profiles are similar to each other as well as to the single-unit temporal profiles of the constituent cells. The second temporal profiles are also very similar to each other despite differences in the second temporal profiles seen when cells are analyzed separately.

Summary across data sets

For most data sets, temporal profiles provided an accurate approximation of the multidimensional scaling coordinates. On average, temporal profiles explained 97% of the variance in the plane of the first two dimensions for responses of individual simple cells and nearly the same fraction for responses of individual complex cells. In the joint response space, temporal profiles accounted for 98% of the variance for simple-simple pairs, 97% of the variance for simple-complex pairs, and 98% of the variance for complex-complex pairs. Thus the structure of the response space could be accurately recovered from linear operations on individual responses. As discussed in METHODS, this is far from trivial because the metrics need not yield distances that correspond to any Euclidean distance derived from linear functionals on the spike trains.

Figure 8 shows average temporal profiles for all simple and complex cells analyzed individually (Fig. 8A) and as components of pairs (Fig. 8B), along with the range that includes 50% of the profiles encountered. The common features of the simple cell examples shown in the preceding text are clearly demonstrated. The first temporal profile for simple cells essentially consists of a weighted difference in the average firing rate of the ON and OFF responses. The second temporal profile for simple cells represents antagonism between the transient component of the ON response and an additive combination of the sustained component of the ON response and the transient component of the OFF response. For complex cells, the first temporal profile represents a weighted sum of the ON response and the OFF response, whereas the second component represents the difference between the ON and OFF responses. The existence of these prototypical temporal profiles suggests a natural mechanism for reading out the spatial phase information contained in the temporal structure of the response. We also note that the sharpest temporal features in these profiles (e.g., the rise time of ~40–50 ms in the profiles for the simple cells) corresponds closely to the temporal precision of phase

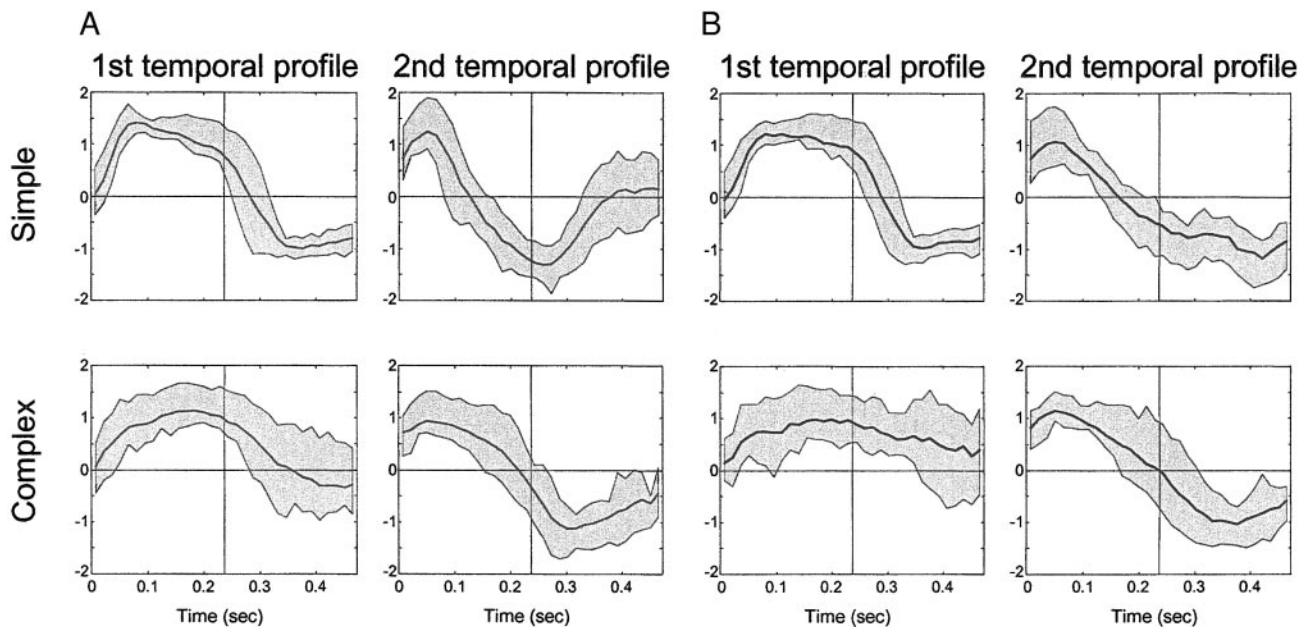


FIG. 8. Average temporal profiles across simple cells (*top*) and complex cells (*bottom*). Shaded areas indicate 50% confidence intervals. Vertical lines at 237-ms mark the disappearance of the stimulus. *A*: temporal profiles of individual responses for the 2 classes of cells. *B*: temporal profiles of joint responses. *Top*: temporal profiles of simple cells analyzed as members of simple-simple and simple-complex pairs. *Bottom*: temporal profiles of complex cells analyzed as members of complex-complex and simple-complex pairs.

coding as determined from the dependence of stimulus-dependent clustering on q ($q_{max} \approx 25 \text{ s}^{-1}$, $1/q_{max} \approx 40 \text{ ms}$).

Although temporal profiles clearly demonstrated commonality across data sets, the confidence bands of Fig. 8 indicated that the shapes of many profiles, especially those of complex cells, were considerably different from the average profiles. In the preceding text, in Fig. 3*B*, we showed that there was also variability in the degree to which the response clusters themselves were distinct. We therefore considered the possibility that the temporal profiles that deviated most from the population means corresponded to cells for which the response clusters themselves were relatively indistinct (i.e., for which H was small). To test this notion, we defined prototypical first and second temporal profiles as the averages of these profiles across all individual cells, and we measured the correlation coefficients of each cell's profiles with these prototypes (r_{profile}). Figure 9 compares these correlation coefficients with the amount of stimulus-dependent clustering H . The first temporal profile in cells with strong stimulus-dependent clustering ($H > 0.3$), a group that included only simple cells, was highly similar to the prototypes ($r_{\text{profile}} > 0.7$). Among cells with

weaker stimulus-dependent clustering, which included both simple and complex cells, the first temporal profiles were more diverse. Some were highly similar to the prototypes, but others were virtually uncorrelated ($r_{\text{profile}} = 0.2\text{--}0.9$). Examination of the second temporal profile yielded similar results for complex cells but somewhat surprising results for the subset of simple cells that showed strong stimulus-dependent clustering. While most of the latter cells conformed closely to the prototype ($r_{\text{profile}} > 0.8$), four of the recorded cells with strong stimulus-dependent clustering had temporal profiles that were distinctly different from the prototype ($r_{\text{profile}} < 0.3$).

To compare consistency of the profiles across groups of cells, and in particular to compare consistency of the profiles between cell pairs recorded simultaneously and cell pairs drawn randomly from the population, we proceeded as follows. For each pair of cells within the population of interest, we calculated Pearson correlation coefficients for their corresponding temporal profiles rather than between the individual profiles and the population average. Because the polarity of temporal profiles is arbitrary, our population measure of overall similarity r_{profile} was equal to the average of the absolute




FIG. 9. Correlation of the amount of stimulus dependent clustering H with the shape of the 1st temporal profile (*left*) and 2nd temporal profile (*right*). The correlation coefficients of the temporal profiles with the prototypical profiles are plotted along the vertical axis. Prototypes were taken to be the averages of temporal profiles across all data sets.

value of the correlation coefficient derived from each pair of cells. A summary of these average correlation coefficients is shown in Table 4. For simple cells, consistency was high ($r_{\text{profile}} = 0.72$) for the first profile but lower for the second profile ($r_{\text{profile}} = 0.49$), reflecting the presence of a subset of cells whose profiles differed substantially from the prototype. For complex cells, the first temporal profile showed somewhat less consistency ($r_{\text{profile}} = 0.51$) than among simple cells, but the second temporal profile was as consistent among complex cells ($r_{\text{profile}} = 0.50$) as among simple cells.

For both simple and complex cells, the consistency of profiles extracted from pair responses was as strong as it was for individual responses (line 2 of Table 4). That is, even though the identity of which neuron fires which spike improves the fidelity of coding, customization of the temporal profiles used to extract phase information from each pair is not required. Moreover, the consistency of profiles extracted from a joint response with the profiles extracted from the individual neurons that constituted the pair (line 3 of Table 4) was very strong ($r_{\text{profile}} > 0.8$ for 1st temporal profile, $r_{\text{profile}} > 0.7$ for 2nd temporal profile). This indicates that temporal structure contributed to the response space of an individual cell similarly to the way it contributed to the joint response space of two nearby cells, and it suggests a detailed similarity of response properties between neighboring neurons (see following text).

Similarity of spatial phase coding in neighboring cells

The simultaneous pair recordings allowed us to examine the similarity of spatial phase representation in neighboring neurons. We compared each measurement (e.g., H , the eccentricity of the fitted ellipse, the shape of the temporal profiles) across pairs of cells recorded simultaneously at a single site and across pairs of cells recorded (not simultaneously) at different sites. For scalar measurements, such as H and the eccentricity of the fitted ellipse, dissimilarity was defined simply by the absolute difference of values at the two neurons. For tuning functions and temporal profiles, similarity was defined by correlation coefficients r_{tuning} and r_{profile} , with averages performed either on the signed (r_{tuning}) or unsigned (r_{profile}) values. A summary of this analysis is shown in Table 5.

To test whether differences between neighboring and non-neighboring pairs of cells was significant, we randomly re-assigned cells to recording sites, keeping the number of cells at

TABLE 4. Similarity of temporal profiles across cells

	S		C	
	1st Profile	2nd Profile	1st Profile	2nd Profile
Comparison across individual cells	0.72	0.49	0.51	0.50
Comparison across pairs of cells	0.67	0.50	0.51	0.53
Comparison of pairs to constituent cells	0.92	0.73	0.87	0.72

Entries indicate average correlation coefficient r_{profile} between the indicated temporal profiles, determined as described in text. Line 1: pairwise comparisons of temporal profiles determined from individual responses of simple cells (S) and complex cells (C). Line 2: pairwise comparisons of temporal profiles determined from joint responses. Line 3: pairwise comparisons of temporal profiles determined from an individual response within a pair, with a temporal profile determined from the pair.

TABLE 5. Characteristics of spatial phase coding in pairs of neighboring cells

Measurement	Neighboring Pairs	Non-Neighboring Pairs
Fraction of S-C pairs, %	19	51
Correlation coefficient of tuning functions (r_{tuning})	0.50	-0.02
Difference of clustering indices (H) in bits	0.15	0.31
Difference of axis ratios of best-fit ellipses	0.10	0.24
Difference of the goodness of fit of the best-fit ellipses, %	15	26
Similarity of temporal profiles (r_{profile})		
1st	0.76	0.58
2nd	0.57	0.46

Each entry represents the average of the indicated quantity, calculated from neighboring cell pairs that were simultaneously recorded (1st column) or exhaustively from all data sets (2nd column).

each recording site the same. We performed 1,000 such reassignments for each comparison and performed the same similarity calculations. We calculated the P value for each measurement as the fraction of reassignments for which the average similarity of neighboring cells was greater than the original value.

Average measurements for neighboring and non-neighboring pairs of cells are shown in Table 5. The fraction of simple-complex pairs was significantly lower ($P < 0.001$) among neighboring pairs, indicating that simple and complex cells were anatomically clustered. Spatial phase tuning functions were strongly correlated (r_{tuning}) within recording sites, indicating that spatial phase preferences were also similar for neighboring cells. (Phase tuning functions across randomly chosen non-neighboring pairs of cells are expected to be completely uncorrelated because all spatial phases are equally represented.) In addition to these classical receptive field characteristics, all other measurements related to spatial phase representation were significantly more similar for neighboring than for non-neighboring pairs ($P < 0.01$ for every measurement). That is, cells that represented spatial phase with high fidelity tended to be grouped together, and cells in which that representation was consistently two-dimensional (large axis ratio and high goodness of fit for the ellipses) were also grouped together.

Significant differences between neighboring and non-neighboring pairs could be due to a sampling bias between cells that were recorded in groups and cells that were recorded individually. Artfactual differences might also have resulted from the fact that the stimulus parameters such as spatial frequency and orientation chosen for studying a group of cells might be compromises between the optimum choice for each cell individually. To rule out these possibilities, we also compared the various measures listed in the first column of Table 5 for cells recorded as components of pairs and cells recorded individually. However, we found no substantial variation between these groups for any of the measurement. Therefore the significant differences between neighboring and non-neighboring cells discussed in the preceding text indicate the dependence of measurements on the proximity of neurons rather than a dif-

ference in sampling or stimulation strategy for cells recorded as part of a group versus cells recorded individually.

DISCUSSION

Our results provide new insights into the way spatial phase is represented in individual responses and joint responses of neighboring cells. It is common to consider simple cells as linear and complex cells as nonlinear phase-insensitive energy operators (Adelson and Bergen 1985; Burr and Morrone 1992; Movshon et al. 1978a,b; Skottun et al. 1991; Spitzer and Hochstein 1985). We show that these generalizations overlook significant detail especially when the temporal structure of responses is taken into account. Furthermore our results demonstrate that the temporal details of the response are conserved across cells and thus provide a means for spatial phase representation.

Temporal coding plays a crucial role in representation of spatial phase

We found that stimulus-dependent clustering of responses to spatial phase is greatly enhanced when the temporal structure is taken into account. Typically, the temporal resolution that optimized clustering (15–60 ms) was more precise than the resolution necessary to distinguish the ON response from the OFF response or to distinguish the transient from the sustained components. Thus our results do not simply concern magnitudes of these major response components but also reveal a systematic variation of responses with spatial phase at higher temporal resolutions.

The number of spikes in a response tended to be phase sensitive for simple cells and phase insensitive for complex cells—consistent with qualitative characterizations of simple cells as linear and complex cells as energy operators. In our analysis, this is seen by the amount of stimulus-dependent clustering at $q = 0 \text{ s}^{-1}$ (Victor and Purpura 1998) and is expected from the above caricature of simple and complex cells. However, when the temporal structure was taken into account, the amount of stimulus-dependent clustering was above chance level for complex cells as well. Qualitatively, the dependence of clustering on temporal resolution was identical for the two classes of cells (H peaked at similar values of q for optimal separation of the responses), indicating that both classes may employ similar temporal coding mechanisms for information about spatial phase.

Consistent stimulus-dependent temporal structure in the response not only increases the amount of information that can be extracted in a formal sense but also supports a qualitative difference in the nature of the neural representation of spatial phase. Without temporal coding, unambiguous representation of the stimulus by individual neurons is not possible because spatial phase requires a two-dimensional response space. Our analysis of response trajectories indicates that temporal coding can be used to represent the entire stimulus space of spatial phases. The trajectories of most neurons that exhibited relatively strong stimulus-dependent clustering were elliptical, indicating that responses of these cells represented spatial phase unambiguously. We found that, although the major dimension of the ellipse was roughly a measure of the number of spikes in the ON and OFF responses, the minor axis reflected the

temporal structure—in particular, the difference between the transient and sustained components of the response. The extracted temporal profiles show that this was consistent across most neurons. Thus relatively simple linear mechanisms can be used to recover the second coordinate of the phase representation and thus to relate individual responses of most cells to the spatial phase of the stimulus. Some neurons had response-space trajectories and temporal profiles that deviated from the prototypical ones. As seen in Fig. 9, a fraction of these cells also had distinct response clusters (large values of H). Nevertheless, stereotyped temporal patterns were consistent across most of the cells with large values of H , i.e., the cells for which the information-theoretic analysis indicated were likely to represent spatial phase.

Responses from as few as 10 independent simple cells, each with $H = 0.4$ (typical of our data, see Fig. 3), carry enough information (4 bits) to reliably distinguish which of 16 equally spaced spatial phases is present. This corresponds to discriminating 2 min of an arc for a 2 c° grating—far better than human performance. Twenty complex cells with $H = 0.2$ (the upper end of our measurements) or 100 complex cells with $H = 0.04$ (the lower end of our measurements) would yield the same performance. Thus there are enough neurons within either class to account for performance, but for complex cells, this requires sensitivity to temporal structure (since stimulus-dependent clustering for complex cells was not significantly above chance for $q = 0$). It is not clear whether counting spikes over some time window, or sensitivity to temporal pattern, is a more “natural” neural operation—there are good arguments on either side (Mel 1993; Sen et al. 1996; Shadlen and Newsome 1998; Softky and Koch 1993). Thus the significance of our results does not lie in a determination of whether a particular class of neurons supports phase representation—clearly there is an adequate number of neurons in either class, with or without temporal coding (in simple cells) or with temporal coding (in complex cells). Rather we emphasize that temporal information contributes to phase coding in two *qualitative* ways by disambiguating contrast from phase (see following text) and providing the substrate for a cyclic representation of a cyclic quantity.

The temporal representation of spatial phase is distinct from that of contrast, another stimulus attribute that has a consistent representation in temporal features of the response. As previously reported by others (Gawne 2000; Gawne et al. 1996) and by our laboratory (Reich et al. 2001a), response latency primarily represents stimulus contrast, whereas response size is controlled both by contrast and by certain aspects of form (e.g., orientation). In the present experiments, all stimuli had the same contrast. The temporal profiles show that phase representation relies on a comparison of relative sizes of the transient and the sustained portions of the response. Had spatial phase been coded by latency, the temporal profiles would have shown monotonic growth during the initial transient of the response, and no contributions from the sustained portion. Thus the temporal representation for spatial phase is distinct from, and can co-exist with, a latency code for contrast.

The similarity of the temporal profiles across neurons is all the more impressive when one considers the analytical steps that were taken to derive them. The distances between individual responses, as determined by the metrics, need not have corresponded to any set of distances in a low-dimensional

space. Moreover, even given that there is such a correspondence, there was no guarantee that these distances could be recapitulated from linear operations on the spike train (for a counterexample, see Fig. 9 of Victor and Pupura 1997). However, in this analysis, these linear operations account for nearly all (e.g., $\geq 97\%$) of the shape of the response trajectories, even though the elliptical fits themselves only account for 60–80% of the shape of the trajectories. That is, the decoding process suggested by the response profiles accounts not only for the phase representation implied by a stereotyped linear receptive field response model (in which trajectories would be exactly elliptical) but for whatever information is conveyed by deviations from response linearity as well.

It is also worthwhile to compare our analysis of temporal aspects of receptive field structure to that of DeAngelis et al. (1999) in the cat and Hawken et al. (1996) in the macaque. The DeAngelis study, which used temporally shifted Gabor functions to parameterize response dynamics, found substantial variability (~ 3 octaves) of the preferred temporal frequency of cortical neurons and at least a quarter of a cycle variation of spatial phase. Preferred temporal frequency, but not temporal phase, was strongly correlated among nearby cells. Hawken et al. (1996) found a similar range of variation of the best temporal frequency for macaque V1 simple and complex cells. Here we examine not the overall temporal tuning of a neuron but, rather, the temporal aspects of a response that are informative about spatial phase. Given the variability of temporal tuning across neurons in these studies, it is striking that we find that the temporal aspects of the response that are relevant to phase coding to be so stereotyped.

Continuum of phase coding behavior

For each of the ways we characterized responses, we encountered a wide and continuous distribution of measurements across data sets. Most simple cells exhibited strong stimulus-dependent clustering and had response trajectories that tended to be elliptical and temporal profiles that were similar to a single prototype. The population also included cells that exhibited weak clustering (both simple and complex), which had less regular response trajectories and temporal profiles that were less similar to the prototypes, as well as a few simple cells with strong clustering but atypical temporal profiles (Fig. 9). The indices of spatial phase coding covaried with the modulation ratio (Fig. 3B) but did not show suggestions of a bimodal distribution (cf. Mechler and Ringach 2002) corresponding to simple and complex cells per se. In particular many simple cells showed weak clustering, despite having a modulation ratio that was clearly within the “simple” range (Fig. 3B). Thus neither the modulation ratio nor a dichotomous partition of it determines how a cell represents spatial phase. As mentioned in the preceding text in conjunction with Fig. 3 (RESULTS), modulation ratio and information transmitted about spatial phase need not be tightly linked.

Spatial phase representation is locally similar

In all of our analyses, we found very significant similarities among neighboring neurons. This included classical receptive field characteristics, such as spatial phase tuning and the modulation ratio. In addition, all measurements used in our analy-

ses, such as those that characterized response clustering, response trajectories, and the pattern of temporal coding, tended to be similar in neighboring cells.

In a previous study of 45 pairs of neurons in primary visual cortex of the adult cat, DeAngelis et al. (1999) reported results that could be interpreted as indicating that spatial phase tuning is not clustered within the cortex. This apparent disagreement with the present findings likely reflects two different ways that spatial phase can be defined, and it can be resolved in a manner that lends insight into cortical functional anatomy. DeAngelis et al. (1999) measured spatial phase relative to the peak of the envelope of a Gabor function for the cells' spatiotemporal receptive fields, where the Gabor was chosen individually for each cell. This relative measure of spatial phase preference describes whether the cell's spatial profile is even-symmetric, odd-symmetric, or intermediate. They found that spatial phase tuning in nearby neurons, defined in this fashion, is uncorrelated. In contrast, the present study defines spatial phase with respect to a fixed point in space and a fixed spatial frequency. Figure 4 shows a substantial tendency for nearby neurons to have similar phase tuning (defined in this manner), especially for simple-simple pairs cells. Had phase tuning of nearby neurons been uncorrelated, then the 20 points in Fig. 4A would have symmetrically distributed about zero as would the distributions of Fig. 4B. Instead, values of r_{tuning} are substantially skewed toward 1. These data are clearly inconsistent with uncorrelated phase tuning, but the limited sample size prevents us from making precise statements about how strong the correlation is.

In sum, though both studies measured “spatial phase” tuning, DeAngelis et al. (1999) showed that receptive field *symmetry* is uncorrelated among neighboring cells [a result that is consistent with our studies of compound gratings (Mechler et al. 2002)], while the present study shows that neurons in a local population have similar spatial phase tunings when referred to a fixed point in space. Similarity of spatial phase tuning of nearby neurons (when spatial phase is measured in the present absolute sense) likely has its anatomical basis in the finding that the basic spatial structure of cortical receptive fields is shaped by thalamic input (Reid and Alonso 1996). If nearby cortical neurons make use of the same geniculate afferents (monosynaptically or via intervening circuitry) to generate the region of peak sensitivity, then the locations of the peaks of their receptive field sensitivity profiles necessarily will match, and similar spatial phase tunings will result. We speculate that, superimposed on this shared input, cortical neurons also receive more variable inputs from within the cortex that dominate their responses in the flanking regions of the receptive field. This would allow the overall envelopes of nearby neurons' sensitivity profiles to vary considerably, and thus nearby cells might vary substantially in local symmetry and other details [accounting for the results of DeAngelis et al. (1999) in the cat and our results (Mechler et al. 2002; Reich et al. 2001b) in the macaque].

Our finding that nearby neurons tend to have similar phase tuning, and the preceding interpretation, might at first seem at odds with earlier studies of the topography of the representation of the visual field in primary visual cortex of the cat (Albus 1975) and monkey (Hubel and Wiesel 1974), which emphasized local scatter of receptive fields. In those studies, receptive fields were compared between successively recorded neurons

along a penetration. Distances between neurons were often on the order of 80 μm on a tangential (Albus 1975) or vertical (Hubel and Wiesel 1974) penetration, but occasionally within 30 μm of each other or simultaneously recorded on a vertical penetration (Hubel and Wiesel 1974). In our study, neuron pairs were always within $\sim 30 \mu\text{m}$ of each other, typically on a near-vertical penetration. This separation is inferred from the three considerations. The contacts on the tetrode form an approximate tetrahedron, with intercontact distances of 30 μm . A typical spike waveform has a substantial amplitude on one or more contacts but not all four. Finally, moving the electrode by 30 μm can result in recording from an entirely new group of cells, whereas moving 10 or 15 μm does not.

While the preceding studies found [either quantitatively (Albus 1975) or qualitatively (Hubel and Wiesel 1974)] that local scatter in receptive field position was comparable to the average receptive field diameter, they also found that overlap of successive receptive fields by 50% or more of their area (70% of their width) was typical. This amount of overlap is also consistent with the findings in the cat of DeAngelis et al. (1999), who found that typically, neighboring receptive fields are offset by not more than one-quarter of their width or one quarter of a cycle of their preferred grating stimulus. This less-than-perfect regularity is sufficient to bias the correlation of spatial phase tuning of nearby neurons toward 1.

We interpret local correlation of phase tuning in terms of shared subcortical input. If such shared input is the anatomical basis of such overlap, we predict that these regions of overlap should correspond to subregions of identical polarity (either ON or OFF). The preceding studies did not provide sufficient detail to determine that receptive field subregions (and not just receptive fields) overlap extensively, but it is entirely consistent with the studies of Reid and Alonso (1996).

Joint coding of spatial phase by pairs of neurons

Although neighboring cells coded spatial phase similarly, analyzing responses jointly improved the coding of spatial phase. For neighboring cells such as those of Fig. 2A that had very similar phase tunings, this improvement could be realized simply by pooling the responses together (i.e., H_{joint} did not increase as a function of the metric parameter k). For neighboring cells such as those of Fig. 2B whose phase tunings differed, the benefit of joint coding depended strongly on k . However, complete separation of the responses (a labeled line code, $k = 2$) was no more effective than partial separation of the individual neurons' responses ($k = 0.5$). Such partial distinction of neighboring neurons is, in fact, more physiologically plausible than the extreme of a labeled line code.

The fact that the benefit of temporal coding (the dependence of H_{joint} on q) was similar for a summed population code ($k = 0$) and for a code with the neurons distinguished ($k > 0$) provides a nonparametric corroboration for the conclusion that temporal coding was consistent across local populations. The elliptical geometry of the response space had a higher axis ratio for pairs of cells than for single cells (Table 3). Thus in addition to exhibiting more distinct clustering, the joint representation of spatial phase was closer to the circular nature of the domain of spatial phase, although this change was modest.

Although tracking which neuron fires which spike ($k > 0$) reduces the redundancy between neighboring neurons com-

pared with a summed population code ($k = 0$), redundancy remained substantial. This contrasts with the results of Reich et al. (2001b), in which a labeled line code resulted in elimination of essentially all redundancy among clusters of up to six neurons in macaque V1. The earlier study, which was based on a population of neurons that overlapped substantially with those presented here, measured information via the direct method of Strong and coworkers (1998) and m -sequence stimuli. The direct method allowed for analysis of joint coding by multiple neurons (only pairs were considered here) but did not allow for consideration of codes intermediate between the extremes of the summed population code and the labeled line code. Statistical considerations related to the direct method restricted that analysis to information *rates* over brief periods of time (ca. 15 m). This difference is likely to be very important: neurons that convey independent information over a brief time period may become redundant over longer periods. A lack of redundancy over time has been reported (Reinagel and Reid 2000) for continuously modulated stimuli. It is thus notable that Reich et al. (2001b) used m sequences—a continuously modulated stimulus—while the stimuli in the present study only changed at onset and offset. Another difference between the two studies is that here, we only examined response variation across spatial phase, while the m -sequence stimuli included a much wider spatial variety. It is likely that some of the redundancy observed here might also be eliminated had we used a larger stimulus set. In summary, finding substantial redundancy for a restricted stimulus set over long time periods (as we do here) is fully consistent with a finding of no redundancy for a rich spatiotemporal stimulus set and brief time periods (Reich et al. 2001b). In particular, the two studies overlap in their conclusion that tracking the neuron of origin of each spike reduces redundancy across neurons.

Summary

This work applies the metric-space approach to analyze spatial phase coding in single and pair-recorded neurons and introduces a novel method, extraction of temporal profiles, to interpret this analysis. Our main finding is that the temporal profile of a response plays a crucial role in transmitting information about spatial phase. The pattern of temporal coding, as revealed by the temporal profiles we have extracted, is remarkably consistent across a large subset of neurons and has features that provide a simple means for decoding this information. All attributes of spatial phase representation are generally shared by a local population. Nevertheless, mere pooling of signals from nearby neurons does not provide for a representation of spatial phase as effective as joint coding.

APPENDIX

Elliptical response trajectories in linear neurons

It is well known that if a linear neuron's response to a sinusoidally modulated spatial sine grating is quantified in terms of the real and imaginary (or sine and cosine) Fourier components, then its response forms an elliptical trajectory as spatial phase varies (Movshon et al. 1978a). However, this observation can be broadened. The coordinates need not be defined by Fourier analysis, and the stimulus time course need not be sinusoidal. To see this, consider first a linear neuron that contains only a single spatiotemporally separable mechanism (i.e., a

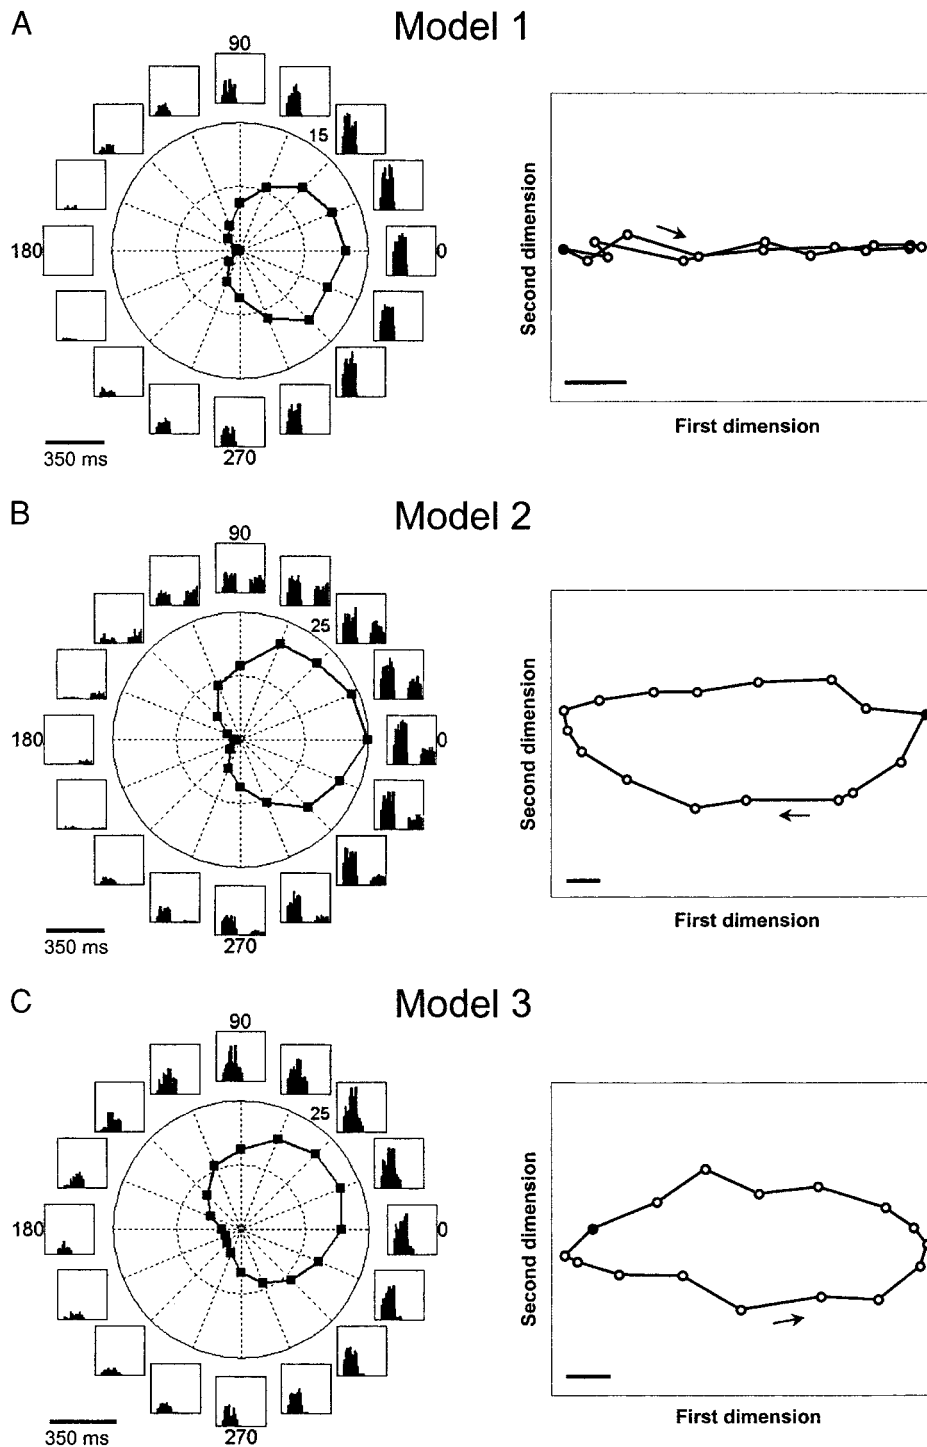


FIG. 10. Average responses (*left*) and multidimensional scaling results (*right*) for 3 linear Poisson models. Conventions as in Fig. 2 for average responses and Fig. 5 for multidimensional scaling. *A*: response trajectory for a model neuron with a separable receptive field. *B*: response trajectory for a model neuron whose receptive field has 2 temporally separated components with overlapping phase preferences. *C*: response trajectory for a model neuron whose receptive field has 2 temporally overlapping components with a maximal phase tuning difference.

mechanism whose spatiotemporal impulse response is a product of a function of space—a sensitivity—and a function of time—an impulse response). A single spatiotemporally separable linear receptive field mechanism necessarily produces a response of constant shape, whose magnitude depends sinusoidally on spatial phase. Consequently, any quantification of the response via a linear procedure will also depend sinusoidally on spatial phase.

Now consider the response of a linear neuron that receives contributions of two or more such mechanisms. The response can be expressed as a linear combination of components, one from each mechanism. As in the preceding text, each component has a fixed

shape (i.e., time course) and a magnitude that varies sinusoidally with spatial phase. As spatial phase varies, at most a two-dimensional subspace of response waveforms is swept out. This is because the “sine” components of all the mechanisms can be considered as a single contribution whose magnitude varies as the sine of the spatial phase, and the cosine components form a second such contribution. Because the size of these two contributions varies as sine and cosine (though their shapes, i.e., directions in response space, are not necessarily orthogonal), the trajectory of the response (in the vector space of all possible response waveforms) is an ellipse (Movshon et al. 1978a).

Elliptical response trajectories for spike metrics: Poisson model neurons

Our method of embedding responses into a vector space does not fall within the confines of the preceding linear analysis because it uses multidimensional scaling applied to metrics, which are generically non-Euclidean. Moreover, for many cells, such as the ones in Fig. 5C, the observed response trajectories deviate substantially from ellipses. Because the metrics are nonlinear, significant distortion of elliptical trajectories might be an artifact of our method rather than indicative of physiology. To rule out this possibility, we performed multidimensional scaling based on spike metrics for linear model neurons. The synthesized spike responses were generated by an inhomogeneous Poisson process whose rate depended on the stimulus as described in the following text.

Model 1 simulates a linear neuron with a separable receptive field—namely, one composed of a single spatial mechanism associated with a single temporal time course. In the simulation, it fires during the period from 50 to 150 ms after stimulus onset, and the average firing rate varies sinusoidally with spatial phase (mean of 30 spikes/s, depth of modulation 30 spikes/s, maximum firing rate at 0°). The other two models simulate linear neurons that have receptive fields composed of two separable components. In Model 2, the two components are temporally separated but have a substantial phase preference overlap. One of the components contributes spikes on the interval from 50 to 150 ms, whereas the other component contributes spikes on the interval from 250 to 350 ms. The firing rates of both components vary sinusoidally (mean of 40 spikes/s, depth of modulation 40 spikes/s for the 1st component, mean of 20 spikes/s, depth of modulation 20 spikes/s for the 2nd component). The maximum firing rates for the two components occur at 0 and 45°. In model 3, components are temporally overlapping but have a maximum phase tuning difference. The firing rates for the two components vary in the same way as for the model 2 neuron but have maxima at 0 and 90°. The timing of the second response component is shifted to the interval from 90 to 190 ms so that it largely overlaps with the first component (50–150 ms).

For each of the 16 spatial phases, we generated 64 responses (Fig. A1, A–C, left) and performed the same multidimensional scaling analysis as for real neurons. The resulting response trajectory of the Model 1 neuron (Fig. A1A, right) is essentially a doubly covered line segment. The best-fit ellipse to this trajectory has an axis ratio of 0.032, which is not significantly different from 0 ($P > 0.15$). The elliptical fit explains 96% of the variance. The response trajectories of the other two model neurons are nearly elliptical. For the model 2 neuron (Fig. A1B, right), the best-fit ellipse has an axis ratio of 0.36 and explains 94% of the variance. For the Model 3 neuron (Fig. 10C, left), the best-fit ellipse has an axis ratio of 0.34 and explains 96% of the variance. In both cases, the axis ratio is significantly different from 0 with $P < 0.001$.

Response trajectories in Fig. 10, A–C (right), constructed with a single linear receptive field component (model 1) or two linear receptive field elements (models 2 or 3), deviate only slightly from ellipses. Because any linear projection of the response of linear model neurons must vary sinusoidally with spatial phase, a model neuron with more than two components is equivalent to a two-component model at any single spatial frequency. This result shows that the elliptical nature of the response space holds not only for the standard Euclidean embedding but is also a useful approximation for the metric spaces. Thus any significant distortions of response trajectories can be attributed to neural response properties that deviate from linearity, not possible nonlinear effects of the metric space analysis or the embedding process.

For example, one commonly observed distortion from an elliptical trajectory observed in real responses was a twisting or pinching as in Fig. 5B, top 2 plots. These areas of pinching generally correspond to stimulus phases in which responses were small. It is likely that these

responses were affected by a threshold because there was no maintained discharge. Because of the nature of the metrics, distances between responses would be reduced by thresholds because there were fewer spikes in the response. This neural nonlinearity would thus cause the centroids of the smallest responses to be closer to each other than the centroids of the larger responses, a fact that appears to account for the main distortion observed in many data sets. This phenomenon, and the interpretation in terms of a response nonlinearity, is closely related to the “wasp-waisting” of the elliptical response trajectories reported in Movshon and Tolhurst’s classical study (1978a) of grating responses in of neurons in cat primary visual cortex. This underscores the fact that the deviation of the trajectories from true ellipses is a manifestation of their response properties rather than an artifact of the metric-space approach.

This work is supported by National Institutes of Health Grants EY-09314 to J. D. Victor and GM-07739 and EY-07138 to D. S. Reich.

REFERENCES

- Adelson EH and Bergen JR.** Spatiotemporal energy models for the perception of motion. *J Opt Soc Am [A]* 2: 284–299, 1985.
- Ainsworth A, Dostrovsky JO, Merrill EG, and Millar J.** An improved method for insulating tungsten micro-electrodes with glass. *J Physiol* 269: 4P–5P, 1977.
- Albus K.** A quantitative study of the projection area of the central and the paracentral visual field in area 17 of the cat. II. The spatial organization of the orientation domain. *Exp Brain Res* 24: 181–202, 1975.
- Aronov D.** Fast algorithm for metric-space analysis of simultaneous responses of multiple single neurons. *J Neurosci Meth* 124: 175–179, 2003.
- Berry MJ, Warland DK, and Meister M.** The structure and precision of retinal spike trains. *Proc Natl Acad Sci USA* 94: 5411–5416, 1997.
- Bullock TH.** Signals and signs in the nervous system: the dynamic anatomy of electrical activity is probably information-rich. *Proc Natl Acad Sci USA* 94: 1–6, 1997.
- Burr D and Morrone C.** Nonlinear model of feature detection. In: *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks*, edited by Nabet B. Boca Raton, FL: CRC, 1992, p. 309–327.
- Burr DC, Morrone MC, and Spinelli D.** Evidence for edge and bar detectors in human vision. *Vision Res* 29: 419–431, 1989.
- Carlton AG.** On the bias of information estimates. *Psychol Bull* 71: 108–109, 1969.
- Dan Y, Alonso JM, Usrey WM, and Reid RC.** Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat Neurosci* 1: 501–507, 1998.
- DeAngelis GC, Ghose GM, Ohzawa I, and Freeman RD.** Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19: 4046–4064, 1999.
- De Valois RL, Albrecht DG, and Thorell LG.** Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res* 22: 545–559, 1982.
- Emerson RC.** Quadrature subunits in directionally selective simple cells: spatiotemporal interactions. *Vis Neurosci* 14: 357–371, 1997.
- Emerson RC and Huang MC.** Quadrature subunits in directionally selective simple cells: counterphase and drifting grating responses. *Vis Neurosci* 14: 373–385, 1997.
- Gawne TJ.** The simultaneous coding of orientation and contrast in the responses of V1 complex cells. *Exp Brain Res* 133: 293–302, 2000.
- Gawne TJ, Kjaer TW, and Richmond BJ.** Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* 76: 1356–1360, 1996.
- Gray CM, Konig P, Engel AK, and Singer W.** Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338: 334–337, 1989.
- Hawken MJ, Shapley RM, and Grosof DH.** Temporal-frequency selectivity in monkey visual cortex. *Vis Neurosci* 13: 477–492, 1996.
- Heeger DJ.** Half-squaring in responses of cat striate cells. *Vis Neurosci* 9: 427–443, 1992.
- Hubel DH and Wiesel TN.** Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158: 295–305, 1974.
- Kruskal JB and Wish M.** *Multidimensional Scaling*. Beverly Hills, CA: Sage, 1978.
- Mainen ZF, and Sejnowski TJ.** Reliability of spike timing in neocortical neurons. *Science* 268: 1503–1506, 1995.

- Marcelja S.** Mathematical description of the responses of simple cortical cells. *J Opt Soc Am* 70: 1297–1300, 1980.
- McClurkin JW, Gawne TJ, Optican LM, and Richmond BJ.** Lateral geniculate neurons in behaving primates. II. Encoding of visual information in the temporal shape of the response. *J Neurophysiol* 66: 794–808, 1991.
- Mechler F, Reich DS, and Victor JD.** Detection and discrimination of relative spatial phase by V1 neurons. *J Neurosci* 22: 6129–6157, 2002.
- Mechler F and Ringach DL.** On the classification of simple and complex cells. *Vision Res* 42: 1017–1033, 2002.
- Mechler F, Victor JD, Purpura KP, and Shapley R.** Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. *J Neurosci* 18: 6583–6598, 1998.
- Meister M, Lagnado L, and Baylor DA.** Concerted signaling by retinal ganglion cells. *Science* 270: 1207–1210, 1995.
- Mel BW.** Synaptic integration in an excitable dendritic tree. *J Neurophysiol* 70: 1086–1101, 1993.
- Merrill EG and Ainsworth A.** Glass-coated platinum-plated tungsten microelectrodes. *Med Biol Eng* 10: 662–672, 1972.
- Milkman N, Schick G, Rossetto M, Ratliff F, Shapley R, and Victor JD.** A two-dimensional computer-controlled visual stimulator. *Behav Res Methods Instrum* 12: 283–292, 1980.
- Milkman N, Shapley R, and Schick G.** Experimental applications. A micro-computer-based visual stimulator. *Behav Res Meth Instrum* 10: 539–545, 1978.
- Movshon JA, Thompson ID, and Tolhurst DJ.** Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J Physiol* 283: 53–77, 1978a.
- Movshon JA, Thompson ID, and Tolhurst DJ.** Receptive field organization of complex cells in the cat's striate cortex. *J Physiol* 283: 79–99, 1978b.
- Oppenheim AV and Lim JS.** The importance of phase in signals. *Proc IEEE* 69: 529–541, 1981.
- Panzeri S and Treves A.** Analytical estimates of limited sampling biases in different information measures. *Network* 7: 87–100, 1996.
- Pollen D, Foster K, and Gaska J.** Phase-dependent response characteristics of visual cortical neurons. In: *Models of Visual Cortex*, edited by Rose D and Dobson V. Chichester, UK: Wiley, 1985, p. 281–291.
- Reich DS.** *Information Encoding by Individual Neurons and Groups of Neurons in the Primary Visual Cortex*. New York: Rockefeller University, 2000.
- Reich DS, Mechler F, and Victor JD.** Temporal coding of contrast in primary visual cortex: when, what, and why? *J Neurophysiol* 85: 1039–1055, 2001a.
- Reich DS, Mechler F, and Victor JD.** Independent and redundant information in nearby cortical neurons. *Science* 294: 2566–2568, 2001b.
- Reich DS, Victor JD, Knight BW, Ozaki T, and Kaplan E.** Response variability and timing precision of neuronal spike trains in vivo. *J Neurophysiol* 77: 2836–2841, 1997.
- Reid RC and Alonso JM.** The processing and encoding of information in the visual cortex. *Curr Opin Neurobiol* 6: 475–480, 1996.
- Reinagel P and Reid RC.** Temporal coding of visual information in the thalamus. *J Neurosci* 20: 5392–5400, 2000.
- Sen K, Jorge-Rivera JC, Marder E, and Abbott LF.** Decoding synapses. *J Neurosci* 16: 6307–6318, 1996.
- Shadlen MN and Newsome WT.** The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18: 3870–3896, 1998.
- Skottun BC, De Valois RL, Grosf DH, Movshon JA, Albrecht DG, and Bonds AB.** Classifying simple and complex cells on the basis of response modulation. *Vision Res* 31: 1079–1086, 1991.
- Softky W.** Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience* 58: 13–41, 1994.
- Softky WR and Koch C.** The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13: 334–350, 1993.
- Spitzer H and Hochstein S.** Simple- and complex-cell response dependences on stimulation parameters. *J Neurophysiol* 53: 1244–1265, 1985.
- Strong SP, Koberle R, Ruyter van Steveninck RR, and Bialek W.** Entropy and information in neural spike trains. *Phys Rev Lett* 80: 197–200, 1998.
- Théunissen F, Roddey JC, Stufflebeam S, Clague H, and Miller JP.** Information theoretic analysis of dynamical encoding by four identified primary sensory interneurons in the cricket cercal system. *J Neurophysiol* 75: 1345–1364, 1996.
- Treves A and Panzeri S.** The upward bias in measures of information derived from limited data samples. *Neural Comput* 7: 399–407, 1995.
- Victor JD and Purpura KP.** Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol* 76: 1310–1326, 1996.
- Victor JD and Purpura KP.** Metric-space analysis of spike trains: theory, algorithms and application. *Network* 8: 127–164, 1997.
- Victor JD and Purpura KP.** Spatial phase and the temporal structure of the response to gratings in V1. *J Neurophysiol* 80: 554–571, 1998.