Interactive report

How the brain uses time to represent and process visual information[1]

Jonathan D. Victor*

Department of Neurology and Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021, USA

Abstract

Information theory provides a theoretical framework for addressing fundamental questions concerning the nature of neural codes. Harnessing its power is not straightforward, because of the differences between mathematical abstractions and laboratory reality. We describe an approach to the analysis of neural codes that seeks to identify the informative features of neural responses, rather than to estimate the information content of neural responses per se. Our analysis, applied to neurons in primary visual cortex (V1), demonstrates that the informative precision of spike times varies with the stimulus modality being represented. Contrast is represented by spike times on the shortest time scale, and different kinds of pattern information are represented on longer time scales. The interspike interval distribution has a structure that is unanticipated from the firing rate. The significance of this structure is not that it contains additional information, but rather, that it may provide a means for simple synaptic mechanisms to decode the information that is multiplexed within a spike train. Extensions of this analysis to the simultaneous responses of pairs of neurons indicate that neighboring neurons convey largely independent information, if the decoding process is sensitive to the neuron of origin and not just the average firing rate. In summary, stimulus-related information is encoded into the precise times of spikes fired by V1 neurons. Much of this information would be obscured if individual spikes were merely taken to be estimators of the firing rate. Additional information would be lost by averaging across the responses of neurons in a local population. We propose that synaptic mechanisms sensitive to interspike intervals and dendritic processing beyond simple summation exist at least in part to enable the brain to take advantage of this extra information. © 2000 Elsevier Science B.V. All rights reserved.

1. Introduction

How neurons represent and process information is of fundamental interest in neuroscience. It is an intrinsically abstract question, since, at a minimum, it seeks a description of a mapping from events, percepts, and actions to something very different: patterns of neural activity. It is tempting to assume that a common set of principles governs neural coding, but it is yet unclear what these principles are, or at what level of detail they can be applied. It is perhaps more reasonable to anticipate that there is a diversity of biological solutions to the coding problem.

Neural systems need to represent a broad range of domains (objects, odors, intentions, movements), perform a great variety of tasks, and operate under diverse constraints (size and weight is at a premium in insects, ability to learn is more critical to human survival). Ultimately, one would like to understand coding both at the algorithmic level and in terms of the biophysical processes that implement these algorithms. At the sensory and motor peripheries, there is often sufficient knowledge about the cellular and subcellular mechanisms of transduction so that the algorithmic aspects of coding can be inferred. However, within the brain, the combined complexities of synaptic [1,35] and dendritic biophysics [39] and of neural connectivity make this bottom-up approach less straightforward.

Many interesting questions concerning the algorithmic aspects of coding can be phrased in terms of the notion of 'information.' Perhaps the most basic issue is what aspects

of patterns of neural activity carry information. Possibilities include the total number of spikes averaged over a time window or a population of cells, the precise times of individual spikes, the presence of bursts, and patterns of correlated activity across cells, such as oscillations. We would also like to know how this information corresponds to perceptual qualities, whether a neuron's performance reflects its ideal information capacity or rather a compromise reflecting the biological costs of achieving this capacity [32], and how information is transferred from one neuron to another or from one brain area to another.

Addressing these questions experimentally requires that we can quantify the information content in neural activity. Shannon's ground-breaking work in communication theory [56], which provides a formal definition of information, has long been the basis of these efforts [50,59]. The basic idea is that information is equivalent to a reduction in uncertainty. If observation of a symbol (a 'word') at the output of a communication channel allows the observer to refine knowledge of what was present at the input, then information was transmitted. Application of this idea to characterize man-made communication channels is relatively straightforward, because the set of words is known a priori. Difficulties arise in attempting to apply information measures to biologic systems, because the words are unknown. To make full use of information theory (and to avoid assuming answers to the above questions), one would want to begin with as few assumptions as possible. The minimal assumption is that each possible configuration of neural activity (i.e., each arrangement of spikes across time and a set of neurons) is a candidate for a word. Ideally, this would be the starting point, and the formalism of information theory would then determine the actual set of words (and hence, the structure of the neural code). However, this strategy rapidly runs into practical difficulties. Experimental estimates of information are biased by finiteness of datasets, and the extent of this bias is directly proportional to the size of the a priori set of words [14,40,43]. Consequently, progress can only be made by making some compromises, e.g., by making somewhat more restrictive assumptions as to the code structure, which are then checked for internal consistency.

One such strategy is the 'direct method' of information estimation [60]. Here, one assumes limits on the length of each word, and on the meaningful precision of each spike. If the ratio of these two timescales is not too large and that the intrinsic variability of neural responses is sufficiently small, useful estimates of information can be obtained from real datasets. These conditions hold in certain insect sensory systems [8] and allow the structure of the neural code to be identified. Difficulties arise in the mammalian central nervous system, because of the presence of temporal structure at both very short [12,15] and very long [13,52] timescales, and the intrinsic variability of responses that may or may not be a part of the coding scheme. These difficulties are compounded in attempts to examine aspects of coding across even small populations of neurons.

For these reasons, we use a variety of strategies to approach questions of sensory coding in primate visual cortex. First, we can address focused questions such as, with what resolution does the timing of individual spikes carry information, via alternative methods (also based on the principle of uncertainty reduction) to estimate information; we describe these methods below. Second, we use computational approaches to identify aspects of the temporal structure of spike trains that are not expected consequences of their firing rates, and thus represent clues to encoding and decoding strategies. Third, methods similar to the 'direct method' [60] can be used, provided that we restrict consideration to information rates over brief periods of time, and bear in mind that information rates may not be additive [18]. Together, these approaches help to paint a picture of how the temporal structure of spike trains within and across visual cortical neurons represents visual information.

2. Metric-space approach: overview

In view of the above considerations, we recently developed a new strategy for the analysis of neural coding. This strategy, based on the mathematical formalism of metric spaces [66,68], focuses on particular physiologically motivated hypotheses concerning the nature of the neural code, while allowing certain other assumptions, often made implicitly, to be subjected to empirical study. The most fundamental of these assumptions is whether regarding a spike train as an estimator of a time-varying firing probability suffices to understand coding [6]. The alternative is that the richer possibilities that derive from the event-like nature of the spike train must be considered, as has been suggested on theoretical grounds by Hopfield [25]. In order to ask this question, we manipulate spike trains as abstract sequences of events, the action potentials, rather than the more traditional approach, which is based on regularly spaced samples of estimates of firing probability.

From an information-theoretic point of view, if two responses have different likelihoods of occurring in response to a set of external stimuli, then the distinction between these two responses has the potential to convey information. It is impractical to collect sufficient data to determine these probabilities empirically, if one adheres rigidly to the notion that any difference between responses (e.g., even a submillisecond shift in the time of occurrence of a single spike) constitutes a difference. Understanding a neural code requires a more refined notion of when two neural responses should be considered similar or different.

In addition, the choice to work with spike trains as sequences of events rather than as binned estimates of firing probability has major implications for the properties of the 'distances' between responses. A sequence of binned

estimates naturally constitutes a vector and lends itself to operations such as averaging and extraction of Fourier components. Moreover, vectors carry with them an implicit geometrical structure that is embodied in properties of the distances between vectors. For example, there is built-in notion of how responses, considered as vectors, can be 'added': addition of a response C to any pair of responses A and B results in two new responses, $A+C$ and $B+C$. The distance between the new responses $A+C$ and $B+C$ is guaranteed to be identical to the distance between the original responses A and B. That is, the notion of distance is translation-invariant, and respects linearity. Another consequence of the vector space structure is more topological: any pair of responses A and B can be connected by a continuous trajectory $xA+(1-x)B$ of mixtures of the original responses. That is, the space is not only 'connected' but also 'convex.' These and other properties of vector spaces and the constraints they place on the associated distances are mathematically convenient, but they may well be inappropriate for modeling of perceptual domains [25,71].

Notions of distance that arise when working with spike trains as sequences of events are more general than distances based on vectors, and are not subject to the above constraints. Since the neural code is 'read' by other neurons, it is natural to consider notions of distance that are based on the biophysics of how neurons operate. Neurons can act as simple integrators, but they can also act as coincidence detectors [2,10,35,39,57], and they can be sensitive to the temporal pattern of their individual inputs (e.g., synaptic depression and post-tetanic potentiation). These behaviors suggest a variety of candidate notions of distance (formally, 'metrics'). Our strategy is to evaluate these candidates experimentally by determining the extent to which they separate neural responses into classes that correspond to the stimuli that elicited them. By quantifying the stimulus-dependent clustering induced by each candidate metric, we can determine whether various stimulus attributes can be signaled by the precise time of occurrence of a spike or by the pattern of interspike intervals and if so, at what resolution. In combination with resampling techniques, we then can determine whether this signaling is merely a consequence of underlying firing rate modulation.

3. Which aspects of a spike train are reliably stimulus-dependent?

To address this question, we introduce two families of metrics. The first set of metrics assesses the extent to which individual spike times (in addition to the number of spikes) are reliably stimulus-dependent, and thus have the potential to carry information. Individual spike times would be expected to be crucial for information transmission if neurons acted as coincidence detectors. This is because the frequency of coincidences between two spike trains depends on the timing of the individual spikes, and not just on the overall firing rates.

To determine whether individual spike times carry information, we use a family of metrics denoted by $D^{\text{spike}}[q]$. The distance $D^{\text{spike}}[q]$ between two spike trains is defined as the minimum total 'cost' to transform one spike train into the other via any sequence of insertions, deletions, and time-shifts of spikes. The cost of moving a spike by an amount of time t is set at qt, and the cost of inserting a spike or deleting it is set at unity. Thus, spike trains are considered similar (in the sense of $D^{\text{spike}}[q]$) only if they have approximately the same number of spikes, and these spikes occur at approximately the same times, i.e., within $1/q$ or less. For neurons that can be considered to behave like coincidence detectors, the value of q for which stimulus-dependent clustering is highest describes the precision of the coincidence-detection. This is because shifting a spike by more than $1/q$ makes just as much a difference as deleting the spike altogether — either maneuver moves the spike train by one unit of distance.

Another general feature of neuronal biophysics is that the postsynaptic response may depend on the pattern of intervals of an arriving signal, independent of any activity at other synapses [1,9,31,35,51,54,58,64]. This motivates another family of metrics, denoted by $D^{\text{interval}}[q]$, which is sensitive to the interval structure of a spike train. Like $D^{\text{spike}}[q]$, the distance $D^{\text{interval}}[q]$ is defined on the basis of the minimal total cost of transforming one spike train into another via a prescribed sequence of steps. However, these elementary steps operate on intervals rather than individual spikes. That is, the parameter q specifies the cost qt of expanding or contracting an interspike interval by an amount t. Note that expanding or contracting a single interval by an amount qt necessarily shifts the time of all subsequent spikes, since the subsequent *intervals* are unchanged. Consequently, the metrics $D^{\text{interval}}[q]$ and $D^{\text{spike}}[q]$ are quite distinct. Closeness in the sense of $D^{\text{spike}}[q]$ does not imply closeness in the sense of $D^{\text{interval}}[q]$, since the former depends on absolute times and the latter only on relative times. For example, spike trains that consist of similar stereotyped bursts are similar in the sense of $D^{\text{interval}}[q]$ whether or not the bursts occur at similar times. The same spike trains would be similar in the sense of $D^{\text{spike}}[q]$ only if the times of occurrence of the bursts were similar as well.

For both families of metrics, the parameter q indicates the relative importance of timing (either of individual spikes or of interspike intervals) compared to spike count. In the limit of $q=0$ (very low temporal resolution), both $D^{\text{spike}}[q]$ and $D^{\text{interval}}[q]$ become independent of the timing of the individual spikes, and are sensitive only to the total number of spikes. We denote these reduced metrics $D^{\text{spike}}[0]$ and $D^{\text{interval}}[0]$ by D^{count}. D^{count} is identical to a vector-space distance, but $D^{\text{spike}}[q]$ and $D^{\text{interval}}[q]$ (for $q>0$) are qualitatively different.

$D^{interval}[q]$ does not respect linearity, and neither metric corresponds to a connected topology.

For both D^{spike} and $D^{interval}$, we examined a wide range of values for q, since neural coincidence-detectors with precisions ranging from milliseconds to seconds have been identified [10], and the range of timescales for which firing rates influence synaptic efficacy is also large. Fortunately, there are highly efficient algorithms [68] to calculate these distances, based on the dynamic programming techniques [53] to compare nucleotide sequences.

If the temporal features that govern a candidate metric can signal a stimulus attribute (such as contrast or orientation), then responses to stimuli that share the same value of attribute will tend to be close to each other in the sense of this metric, while responses to stimuli that differ along this attribute should be separated by a greater distance. That is, the position of the response in the space whose geometry is determined by the metric will reduce the uncertainty about this stimulus attribute. To quantify the extent of this stimulus-dependent clustering and consequent uncertainty reduction, we use a dimensionless quantity, the transinformation H [5].

If the largest value of H is achieved for D^{count}, then individual spike times or interspike intervals have no systematic dependence on the stimulus beyond what is implied by overall spike count. On the other hand, if larger values of H are achieved for $D^{spike}[q]$ or $D^{interval}[q]$, then the corresponding features (spike times or spike intervals) of the spike train are stimulus-dependent, in a manner that is not predictable on the basis of the spike count alone. Thus, these temporal features have the potential to represent a stimulus domain with greater fidelity than the representation provided by spike counts.

The value of q that maximizes H is denoted q_{max}. According to the definitions of $D^{spike}[q]$ and $D^{interval}[q]$, shifting spikes (or intervals) by more than $1/q$ results in a spike train that is just as 'different' as adding or deleting a spike. Consequently, $1/q_{max}$ is the width of the temporal window within which the fine structure of spike times contributes to signaling. That is, q_{max} can be viewed as the informative precision of spike timing, and is thus a key descriptor of the neural code. In principle, q_{max} is limited not by the temporal resolution of sensory processing (e.g., the frequency response of the front-end neurons), but rather by the intrinsic precision of spike generation. q_{max} is also not directly related to the temporal characteristics of the stimuli; rather, it is the time resolution that best distinguishes the dynamics of the responses that they elicit.

It is important to note that our inferences will be drawn from the relative size of the transinformation H and its dependence on the metric. That is, although H is motivated by information-theoretic considerations, it cannot be equated with the actual information contained in a spike train, which is typically larger. There are many reasons for this: (i) the brain does not use our decoding scheme; (ii) our decoding is almost certainly not optimal; and (iii) only a limited number of stimuli have been explored. This is merely an honest statement about the difficulties of understanding how the brain represents information, not a failure specific to our approach. These difficulties notwithstanding, comparison of H values across a range of metrics is a rigorous way to determine which aspects of a spike train depend systematically on the stimulus, and thus have the potential to carry usable information.

4. Application to coding of visual attributes by single neurons

We applied the above analysis to responses of visual cortical neurons in primary (V1) and secondary (V2) visual cortex to stimuli that varied across several attributes, including contrast, orientation, size, spatial frequency, and texture (Fig. 1). These recordings were carried out in alert monkeys trained to fixate; further experimental details are provided in [66]. For each dataset, we determined the extent to which the two families of metrics $D^{spike}[q]$ and $D^{interval}[q]$ led to consistent stimulus-dependent clustering, and we quantified this clustering by the corresponding transinformation H. The scattergram in Fig. 2 shows a comparison of this measure across all neurons and the five stimulus attributes. In most cases, the spike time metrics $D^{spike}[q]$ do a better job of segregating the responses than the spike interval metrics $D^{interval}[q]$, as seen by the preponderance of points below the diagonal. This scattergram gives hints that the several stimulus attributes may be coded in distinct fashions. For example, the relative advantage of $D^{spike}[q]$ over $D^{interval}[q]$ is greatest for coding of contrast and check size, as seen by the scattering of red and green points that lie substantially below the diagonal. On the other hand, in some neurons, $D^{interval}[q]$ has a slight advantage over $D^{spike}[q]$, but this is almost exclusively for coding of texture (pink) and spatial frequency (yellow).

Fig. 3 provides a closer look at the dependence of coding on stimulus attribute. Here, we focus on $D^{spike}[q]$ and examine the value q_{max} that provides the optimal stimulus-dependent clustering. In primary visual cortex (V1), q_{max} is somewhat higher for coding of contrast than for coding of the other attributes, but this difference is modest. In V2, there is a more dramatic difference. Coding of contrast is characterized by q_{max} of approximately 100 s^{-1}, corresponding to a temporal precision ($1/q_{max}$) of 10 ms. Coding of orientation and check size is characterized by an intermediate value of q_{max} of approximately 30 s^{-1}, corresponding to a temporal precision of 30 ms. Coding of spatial frequency and texture is characterized by the lowest temporal precisions (q_{max} ca. 10 s^{-1}), corresponding to a temporal precision of 100 ms. This means that contrast and pattern information are multiplexed. The same spike train primarily signals contrast when examined with high tem-
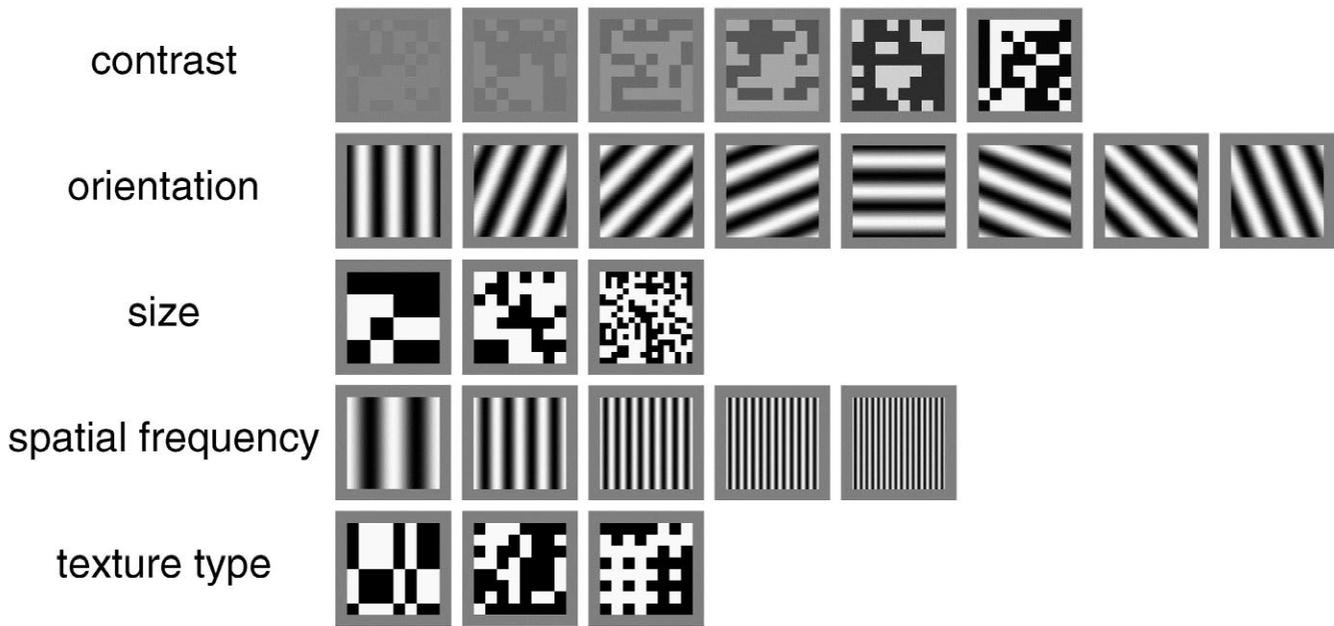
Fig. 1. Examples of visual stimuli used to characterize the informative precision of spike trains.

poral resolution but primarily signals spatial information when examined at lower temporal resolution.

These characterizations of coding and temporal precision represent averages over the entire response duration. Typically, cortical neurons' responses to the abrupt appearance of a pattern consist of a transient response at stimulus onset followed by a tonic or sustained component. We have recently analyzed the coding characteristics of these components separately [45], in recordings of V1 neurons elicited by stimuli varying in contrast. During the initial transient, temporal precision can be as high as 1 ms Conversely, the contribution of temporal pattern during the sustained component is often minimal, with no advantage for $D^{\text{spike}}[q]$ over $D^{\text{count}}[q]$. Thus, it appears that the ultimate biophysical precision of cortical neurons observed in vitro [34] is indeed available for signaling.

However, the notion that temporal structure in spike times relative to stimulus onset can be used for signaling only makes sense if the cortex somehow 'knows' when stimulus onset occurs. In the laboratory, transient presentation of a stimulus (the abrupt appearance of contrast) triggers an excitatory burst of population activity. Under natural viewing conditions, the eyes make approximately three saccades per second, and the visual scene is more or less fixed during each intersaccadic interval [70]. Each saccade results in abrupt contrast changes that likely trigger a burst of population activity. In either case, the population burst marks stimulus onset, and hence provides a reference for the times of subsequent spikes (both for decoding and encoding). If this view is correct, then a transient in the visual stimulus might facilitate the appearance of informative temporal structure in the spike train. Indeed, this is what we found [37] in a metric-space

analysis of neural responses in V1 to drifting sine gratings and edges. Responses to drifting sinusoidal gratings showed essentially no evidence of temporal coding, but drifting edges resulted in temporal coding of contrast that was comparable to what we have seen [69] for transient stimuli. This suggests that the burst of excitatory activity associated with a transient (whether resulting from abrupt onset of a static pattern, the transit of an edge across the RF, or an eye movement) provides a resetting or synchronizing event for its intrinsic circuitry [44]. It is noteworthy that the typical intersaccadic interval [70] is well-matched to the period over which spike train dynamics are informative [66].

5. Individual spikes: merely estimators of a firing rate, or something more?

The above analysis shows that the temporal pattern of a cortical sensory neuron's activity contains information about the contrast and spatial attributes of a visual stimulus. However, it stops short of showing that the timing of individual spikes is significant. The alternative possibility is that the information conveyed solely by the time-course of the average firing rate, which varies in a stimulus-dependent manner as a function of time. In the brain, this average firing rate might be extracted by averaging over a population of similar neurons [55]. In the laboratory, the average firing rate can be determined by averaging responses to many replications of the same stimulus.

The idea that the time-course of firing can convey spatial information is firmly grounded in sensory physi-
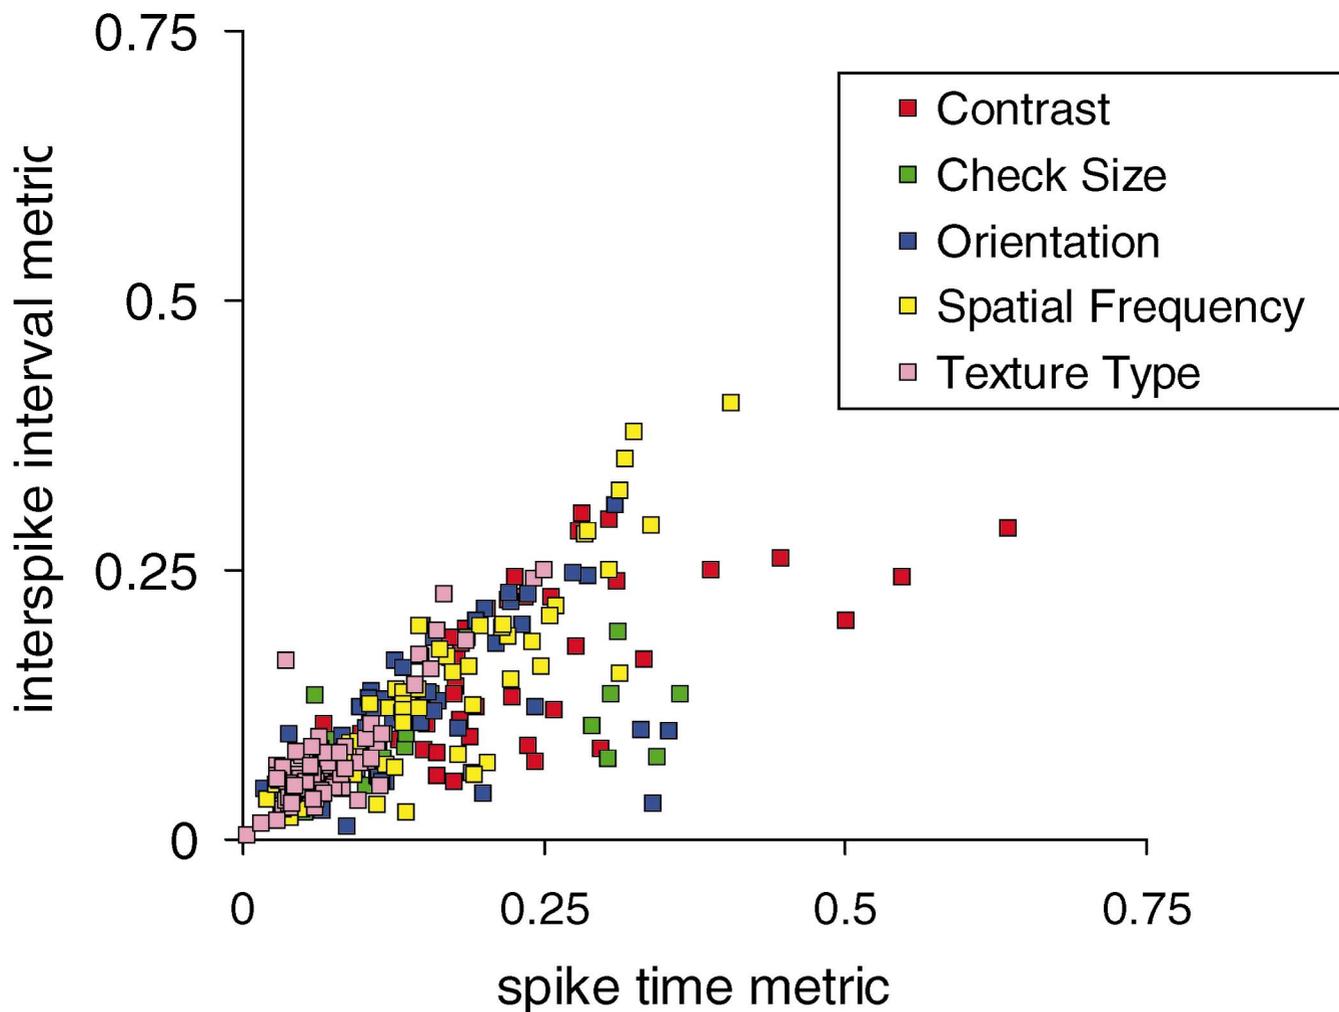
Fig. 2. The extent to which information concerning five visual attributes (contrast, size, orientation, spatial frequency, and texture type) is carried by spike times and interspike intervals, as determined from recordings in V1 and V2. The abscissa indicates the maximal information achieved by $D^{spike}[q]$, and the ordinate indicates the maximal information achieved by $D^{interval}[q]$. A correction for chance clustering has been subtracted, and information values have been normalized by the maximum achievable value, $\log_2$(number of stimuli). For further details, see [66]. Adapted from Fig. 5 of [66] with permission of the Publisher.

ology [11,21,65], and the computational elements required for this to occur are clearly present in neural circuitry. For example, to a first approximation, the firing rate of a retinal ganglion cell can be considered to be a linear function of its input. This linear function has two components: a fast component from the receptive field center, and a slower, antagonistic component from surrounding space. Consequently, stimuli that are restricted to the receptive field center tend to produce a sustained discharge. Stimuli that are more spatially extended recruit the delayed, antagonistic surround component as well, and consequently result in a more transient response. More generally, any combination of inputs that differ both in their spatial and temporal characteristics will result in a coupling of the spatial aspects of the input into the temporal aspects of the average response.

But the basic biophysics of neurons suggest that the temporal structure of individual responses, and not just the average firing rate, can also contribute to the transmission of information. A much-reduced version of the Hodgkin–Huxley neuron, integrate-and-fire neuron, behaves in this way [29,30] since stimulus intensity is determined by the interval between two successive spikes, so that it is not necessary to average across replicate runs or neurons.

The reason for this behavior is that spike trains produced by an integrate-and-fire neuron (and more realistic model neurons) are more structured than what would be produced by neurons that generate spikes at uncorrelated and random times, constrained only by an overall firing rate that may vary with time. The latter kind of spike train, a time-dependent Poisson process, is guaranteed to carry information only in its (possibly time-varying) average firing rate.

Conversely, spike trains that are not time-dependent Poisson processes cannot be described solely by their average time-course, and the deviations from Poisson
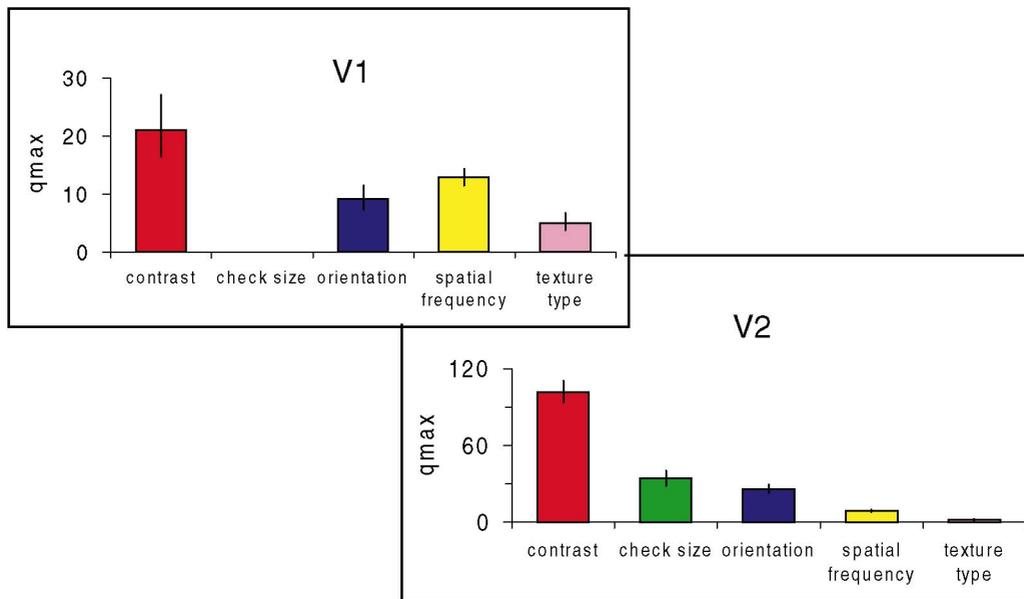
Fig. 3. Characterization of temporal representation of five visual attributes in V1 and V2. Plotted values are the geometric means of q_{max}, the value of the cost parameter q at which the information extracted by $D^{spike}[q]$ is maximal. q_{max} determines the informative precision of spike times, which is high for contrast, and low for pattern attributes. The difference in coding across attributes is more marked in V2 than in V1. For further details, see [66]. Adapted from Fig. 6 of [66] with permission of the Publisher.

behavior may be useful for carrying information or decoding. Neural simulations based on reduced or extended Hodgkin–Huxley models [20,36,57] characteristically lead to spike trains with non-Poisson characteristics.

Thus, there is ample reason to believe that the firing rate can carry information about two attributes, such as contrast and pattern. However, there is also reason to believe that the information in the time-course of the response is not solely carried by average firing rate. To determine whether a cortical neuron's firing rate is the sole carrier of such information, we created surrogate datasets that preserved the time-course of the average firing rate but disrupted the temporal structure of the individual responses. If information is contained solely in the time-course of the average firing rate, then this maneuver will leave the amount of transmitted information unchanged. Conversely, if the surrogate datasets carry less information, then the arrangement of spikes in the individual responses must be significant.

To create these surrogate datasets, we used an 'exchange resampling' procedure [66]. In exchange resampling, pairs of spikes are randomly and repeatedly swapped between different examples of responses to the same stimulus. For each stimulus, the resulting spike trains have the same firing rate time-course as the original dataset, since they are built out of the same spikes. The only difference between the surrogate dataset and the original dataset is the correlation structure within individual responses. When these surrogate datasets are analyzed by the spike metric method, we find (e.g., Fig. 7C and Table 3 of [66]) that clustering is reduced, on average, by about 20% in

comparison to an analysis of the original spike trains. This gap represents the information contained in the temporal pattern of each response beyond what is present in the average firing rate elicited by each stimulus.

We conclude that the spike trains of cortical neurons are not examples of a Poisson-like process with a time-dependent firing rate. Non-Poisson spike trains allow more efficient decoding schemes than population averaging (even if the goal is to extract average firing rate), and they make available a much richer set of possibilities for the manner in which spike trains to represent visual stimuli.

The non-Poisson nature of the spike trains generated by neurons in primary visual cortex can readily be demonstrated without recourse to information-theoretic calculations. Variability in spike counts is typically in excess of the Poisson prediction [16,24,66]. This can have surprising ramifications, including an apparent excess of precisely timed spike patterns [42]. Another non-Poisson feature, the presence of bursts [33], has been postulated to play an important role in information transmission. The analysis we present below provides both another direct demonstration of the non-Poisson nature of a V1 neuron's responses, and some insight into how this temporal structure might be used for signaling (further discussed in [46]).

As diagrammed in Fig. 4, we [46] recorded responses of V1 neurons elicited by a sequence of pseudorandom checkerboards (m-sequences [61–63]). The sequence of checkerboards was long enough to characterize the receptive field, yet short enough to permit recording of responses to multiple identical trials. The first analysis applied to these spike trains is a tabulation of the dis-

M-sequence analysis: firing pattern
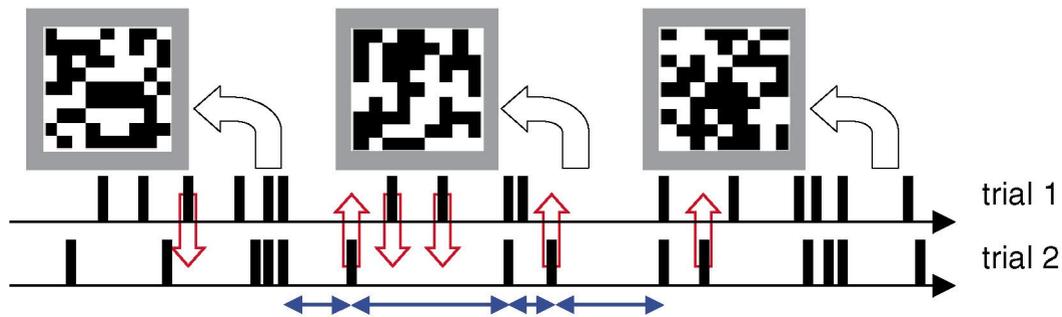
Fig. 4. Scheme for the use of m-sequences to analyze interspike interval structure and its relationship to receptive field organization. An interspike interval histogram (Fig. 5) is created from the responses elicited by a sequence of pseudorandom checkerboards. The interspike interval histogram anticipated from rate modulation is obtained from the interspike intervals in a surrogate dataset created by randomly exchanging spikes between trials (hollow red arrows). Responses to the m-sequence are also cross-correlated against the stimulus (hollow black arrows) to obtain an estimate of the spatiotemporal receptive field (Fig. 7) [49,61–63].

tribution of interspike intervals (the interspike interval histogram, ISIH), shown in Fig. 5. Note that the binning on the abscissa is logarithmically spaced, to facilitate visualization of the extremes of the interspike interval distribution. There are three modes in this distribution: a population of short intervals with a mode at approximately 2 ms, a population of long intervals with a mode at approximately 150 ms, and population of intermediate intervals (between 3 and 38 ms).

Because these spike trains were elicited by stimulation with an m-sequence that strongly modulated the neuron's firing rate, one might wonder if the shape of this ISIH can be explained solely by rate modulation. To test this, we also calculated the ISIH of exchange-resampled responses (red curve in Fig. 5). Were the original responses rate-modulated Poisson processes, the two ISIHs would be identical. However, a marked difference between the two ISIHs is evident.

Of the 66 neurons studied [46], 11 neurons showed this trimodal structure, and across these neurons, the positions of the three modes were conserved. In another 29 neurons, the ISIH was bimodal. In the bimodal ISIHs, the partitions between the modes coincided with the partitions observed in the trimodal histograms. In eight of these neurons, the short interspike interval mode remained distinct, and the medium and long modes appeared merged. In the remaining 21, the long mode remained distinct, and the short and medium modes were merged. The consistency of these mode boundaries across neurons is shown in Fig. 6.

As in Fig. 5, the Poisson prediction failed to account for the ISIH distribution in more than half of the neurons,

including neurons with bimodal and even unimodal ISIH distributions. The difference between the ISIH predicted by rate modulation alone and the experimentally-observed ISIH cannot be due to the absolute refractory period, since the differences between the two ISIHs are prominent for long intervals as well as short intervals. But other intrinsic neuronal properties may contribute to the observed ISIHs, and their conserved features across neurons. For example, a slow calcium conductance is known to underlie the bursts seen in the spike trains of thalamic neurons [27,28] and may play a role in defining the short interspike intervals. However, it is unclear what mechanisms might contribute to the two longer modes.

Although the biophysical basis for these modes is at present unclear, we can nevertheless ask about the roles that the interspike intervals play in signaling. As diagrammed in Fig. 4 and described in detail elsewhere [49,61–63], reverse correlation of the observed spike trains with the m-sequence stimulus leads to a spatiotemporal map of the receptive field. This map has dual interpretations: (i) the average stimulus that most likely precedes a spike, and (ii) the net excitatory or inhibitory contribution of each point in space-time to the firing probability. (These are approximate statements, rigorously valid only in the limit of a linear receptive field). The first row of Fig. 7 shows the receptive field map of a typical cortical complex cell (with a unimodal ISIH), demonstrating a vertically oriented excitatory region peaking at 67 ms prior to spike onset (red) flanked in both space and time by a much stronger inhibitory region (blue). That is, spikes are generated by this neuron preferentially in response to
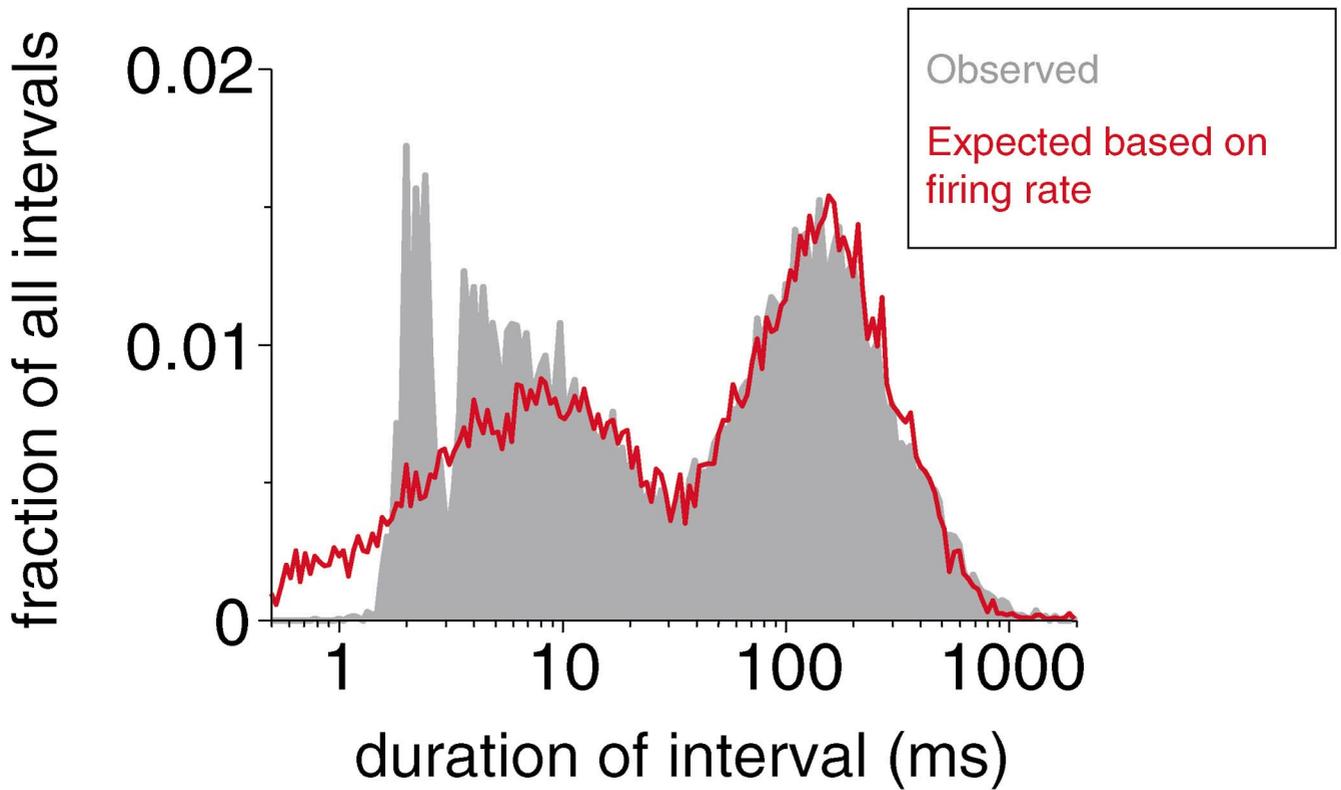
Fig. 5. The interspike interval histogram of a cell in V1 stimulated by an m-sequence, as diagrammed in Fig. 4 (gray profile), and the prediction based on rate modulation (red). Note the presence of three modes in the histogram obtained from the neural response, and the qualitative discrepancy between this histogram and the prediction based on rate modulation. Unit 35/1, a simple cell. Adapted from Fig. 3 of [46] with permission of the Publisher.
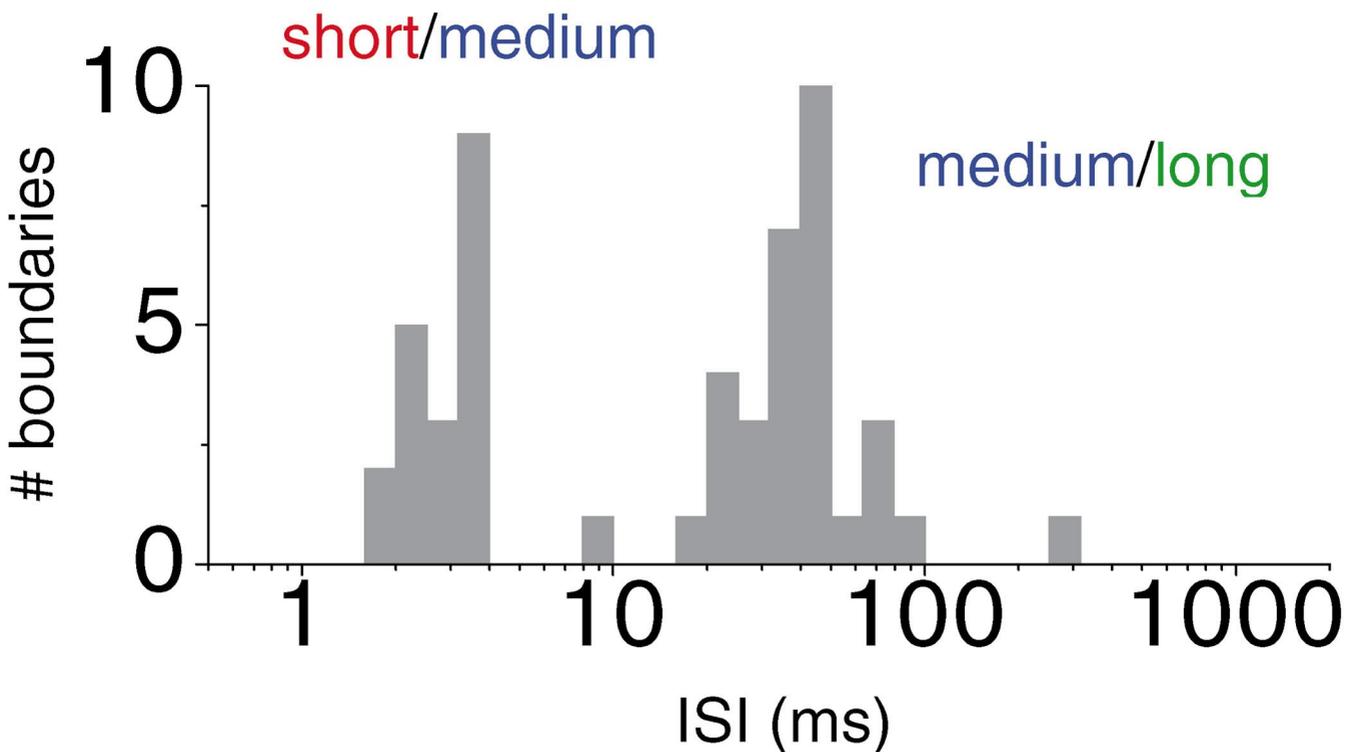

Fig. 6. The boundaries between the modes of interspike interval histograms in 66 V1 neurons, including 11 neurons with trimodal structure (as in Fig. 5) and 37 neurons with bimodal structure. Adapted from Fig. 2 of [46] with permission of the Publisher.

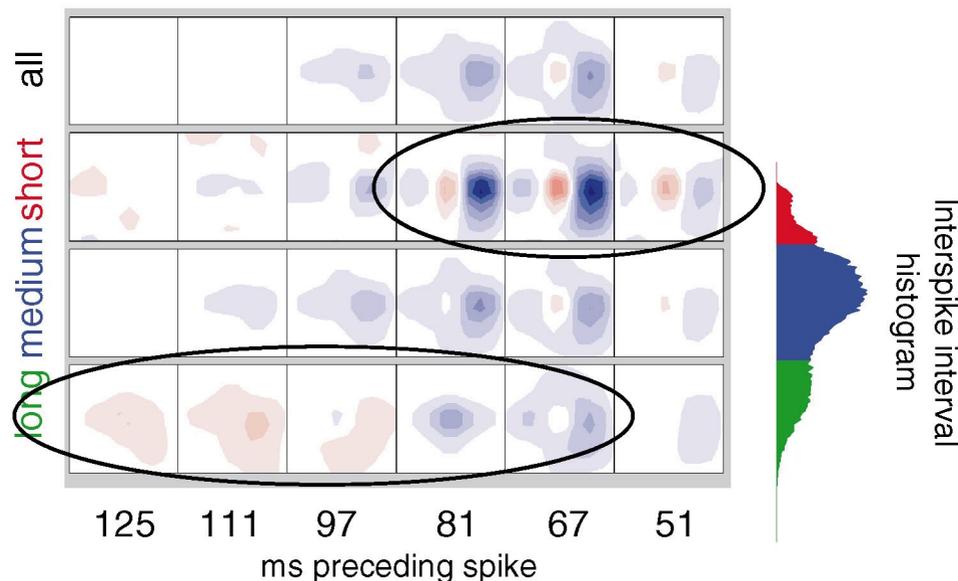# Dependence of receptive field characteristics on spike subset

Fig. 7. Top row: the spatiotemporal receptive field of a V1 neuron estimated by reverse correlation of all spikes with an m-sequence (as diagrammed in Fig. 4). Bottom three rows: the spatiotemporal receptive field determined by reverse correlation of subsets of spikes, selected on the basis of their preceding interspike interval. The spike subsets are defined by partitioning the interspike interval histogram of this neuron, illustrated on the right, according to the boundaries between modes identified across all neurons, as shown in Fig. 6. Note that each spike subset defines a different spatiotemporal receptive field. Complex cell 33/1. Adapted from Fig. 4A of [46] with permission of the Publisher.

patterns that have a vertical orientation, and also at a particular time. Despite the clearly-defined receptive field map, the message carried by any single spike is ambiguous: even if no noise were present: spikes can be generated by stimuli that are suboptimal spatially but optimal temporally, or vice-versa, or of high contrast but not particularly optimal either in space or in time.

Using a modification of the reverse correlation technique, we can determine whether the spikes belonging to each of the three modes help to parse this information. In this analysis, we partition the interspike interval distribution along the lines of the conserved mode boundaries (Fig. 6), even though these modes are not clearly apparent in this neuron's responses. The rationale for this partitioning is an assumption that there is a common synaptic or post-synaptic neural machinery for decoding this information. Such machinery might rely on the common dynamics of synaptic depression and facilitation [1,54], rather than on neural measurement of the (possibly idiosyncratic) statistics of individual spike trains. The spike train is thus partitioned into three subsets: spikes that follow the short, medium, and long interspike intervals observed in the entire population. Reverse correlation of each of these three subsets against the m-sequence stimuli yields the maps shown in the lower three rows of Fig. 7. The receptive field map for the short interspike intervals (second row) has a more pronounced spatial structure, with

more elongated and more sharply defined subregions, consistent over a range of time. That is, spikes that follow short interspike intervals preferentially signal a vertical orientation, and are relatively less affected by the dynamics of the stimulus. Conversely (bottom row), the receptive field map for spikes that follow long interspike intervals does not reveal spatially antagonistic subregions, but rather a temporal transition from excitatory to inhibitory between 97 and 81 ms prior to stimulus onset. That is, spikes that follow long interspike intervals indicate that there has been a change in overall luminance at a particular time, and they are relatively independent of the spatial structure of the stimulus. Spikes that follow the intermediate interspike intervals (third row) show a mixture of these features, with some spatial antagonism and intermediate temporal selectivity; indeed, the receptive field map of these spikes is very similar to the overall receptive field map (top row).

These general features were typical of the neurons we examined: spikes that followed brief interspike intervals tended to yield receptive field maps that were more spatially selective, while spikes that followed long interspike intervals tended to yield receptive field maps that were more temporally selective.

In summary, the interspike interval structure of a spike train has structure that is not merely a consequence of its time-varying firing probability, and this structure appears to be useful in decoding the intrinsic ambiguity of a single

neuron's response. This kind of information is readily accessible to known mechanisms that act at the level of individual synapses (e.g., synaptic depression and facilitation), and that modulate sensitivity to spikes based on the length of the preceding interspike interval.

6. Firing patterns across neurons

A single spike is intrinsically ambiguous, and the temporal structure of the spike train may help to disambiguate the information it carries. In particular, we showed above that the interspike interval structure allows postsynaptic neurons to extract different messages from the same spike train, via preferential sensitivity to intervals of particular durations. However, it is unlikely that temporal structure within single spike trains is the sole clue to disambiguate neural information. Indeed, one might expect that the pattern of activity across a local cluster of neurons plays a greater role in visual coding [38].

One can imagine a range of decoding schemes for reading out this pattern of activity. Perhaps the simplest such scheme is the notion [55] that the average firing rate of a cluster of functionally similar neurons is the primary signal. At the other extreme is the notion [3,4,7] that the detailed relationship of spike times across neurons is crucial.

Summation of activity within a local cluster to estimate an average firing rate provides a signal that reflects the average tuning properties of the local cluster. To the extent that nearby neurons have similar tuning properties [26,41] and that differences in individual neurons' responses reflect only 'noise' [55], this strategy is both simple and complete. However, when analyzed in quantitative detail, adjacent neurons can differ significantly in their tuning [17]. Such differences can lead to a more efficient representation of visual information, provided that the population activity is not decoded merely by summation.

We recently made multi-unit recordings that allow us to address this issue [47,67], with an extension of the spike metric analysis described above. The responses from a cluster of neurons can be considered to be a sequence of spikes labeled by the time and neuron of origin. Our focus is whether the labeling each spike by its neuron of origin provides additional information, and if so, how the additional information relates to the informative timing precision of individual spikes. We generalized the spike time family of metrics D^{spike} to sequences of labeled events by introducing another kind of elementary step — changing the label. The cost associated with changing the label is a parameter k, analogous to the cost q of shifting a spike in time. The efficient computational algorithms for $D^{\text{spike}}[q]$ can be generalized to schemes to calculate $D^{\text{spike}}[q,k]$. The generalizations are 'efficient' (i.e., run in polynomial time), but computationally intensive.

The useful range of the new parameter k is from 0 to 2.

For $k=0$, the metric $D^{\text{spike}}[q,k]$ ignores the neuron of origin. That is, responses are compared after all the component spike trains are merged into a population response. With increasing k, the neuron of origin plays an increasingly larger role in determining the similarity of two multineuronal responses. At $k=2$, the metric is maximally sensitive to the neuron of origin. In this case, spikes from each neuron are kept completely segregated, since moving a spike from one neuron to another has the same cost as deleting it from one neuron's response and inserting it into another neuron's response. Thus, when $k=2$, the distance between two multineuronal responses is the sum of the distances between each of the corresponding single-unit responses. At intermediate values of k, the metric has an intermediate sensitivity to the neuron of origin. Changing the neuron of origin of a spike from one neuron to another has the same effect as changing the number of spikes in each neuron's response by $k/2$, or, of keeping the neuron of origin the same, but shifting the spike in time by k/q.

Recordings were made with tetrodes [23] in V1 of anesthetized, paralyzed monkeys, and spikes were classified using in-house modifications of spike-sorting software [19]. Two-hundred and eighty-seven neurons at 57 distinct sites were recorded, in clusters of two to eight neurons. The analysis described here is limited to pairs of neurons that were cleanly isolated and had stable, robust responses. Stimuli consisted of gratings whose orientation and spatial frequency was either optimal or close to optimal for both neurons of the pair. Stationary gratings were presented abruptly at each of 16 spatial phases (positions), and typically 64 responses of the neuron pair to each position were collected and analyzed.

Fig. 8 analyzes responses obtained from two representative neuron pairs (rows), via an information measure of the clustering induced by the various metrics. For each recording, responses were analyzed for the first 100 ms (the on transient only), the first 256 ms (the entire on response), and the first 473 ms (the on response and the off response). First, we focus on the coding within each of the two neurons, as characterized by $D^{\text{spike}}[q]$. This analysis is summarized by the foreground curves in each panel. The height of the curves at $q=0$ corresponds to the amount of information that each neuron transmits about spatial phase, based solely on the number of spikes in its response. As indicated by the ascent of the curves for increasing values of q, sensitivity to the timing of individual spikes typically increases the amount of information that can be extracted from these individual responses. For optimal sensitivity to spike times, the size of this increase is typically 20–30% over the amount of information that can be extracted when spike times are ignored ($q=0$). The maximal increase occurs for q in the range 30–100, corresponding to a temporal precision of $1/q=10–30$ ms. At even higher values of q, information is reduced, indicating that, on average, sensitivity to spike times at even finer precision does not provide additional information about spatial
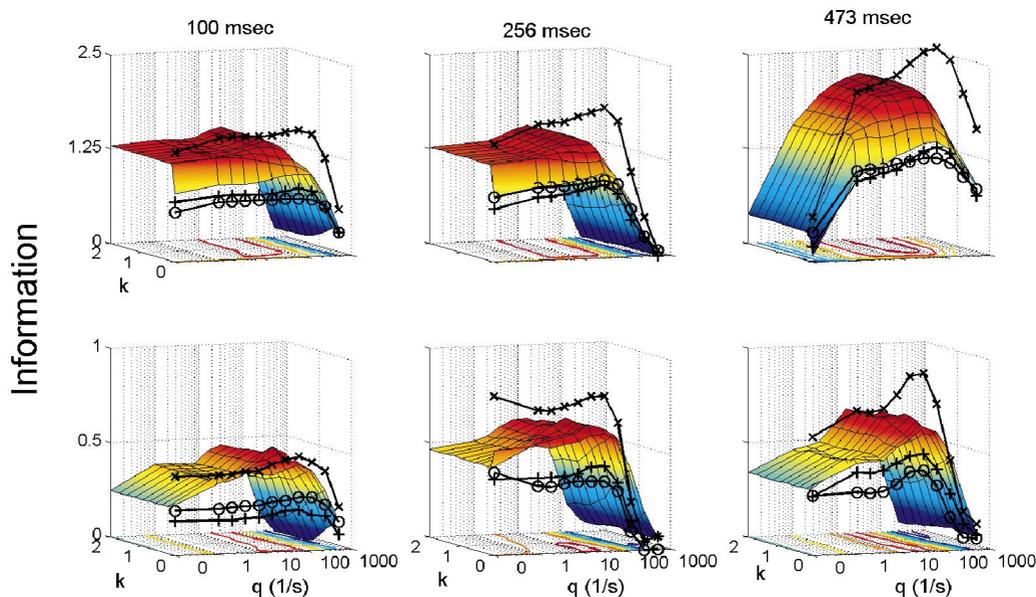
Fig. 8. Comparison of coding by pairs of V1 neurons (surface) and individual neurons (lines), in response to transiently presented gratings at a range of spatial phases. Responses are analyzed for 100, 256, and 473 ms after stimulus onset (columns). The height of each line graph is an information-theoretic measure of the extent of stimulus-dependent clustering (see [66]) as determined for a range of candidate codes $D^{\text{spike}}[q]$, where q is the cost/s to shift a spike. This measure of information, applied to responses from each of the neurons individually, is indicated by the symbols ($\bigcirc$) and ($+$). The sum of these values, which corresponds to the maximum information that could be transmitted by the pair if their responses were independent, is indicated by the line with symbols ($\times$). For neuron pairs, the surface indicates the extent of stimulus-dependent clustering of the joint responses as a function of q and a second parameter, k. The parameter k is the cost to change the neuron of origin of a spike. The increase in information for $q>0$ indicates the informative value of spike timing, independent of the neuron of origin. The increase in information for $k>0$ indicates the informative value of the neuron of origin of each spike. An estimate of the upward bias of information due to chance clustering has been subtracted. Units 41/1/6/s and 41/1/6/t (top); 43/9/8/t and 43/9/8/u (bottom), all simple cells.

phase. (Note that this is an average measure of precision across the entire spike train. Had we analyzed only the initial transient, the useful temporal precision would have been higher [22,48]).

The surfaces in each panel characterize the joint coding by the two neurons via the amount of clustering induced by $D^{\text{spike}}[q,k]$. At $k=0$, the neuron of origin of each spike is ignored, and thus, the response of the two neurons together is interpreted as if it were simply summed into a single spike train. Not surprisingly, the information carried by this composite train of 'unlabeled' spikes is greater than the information carried by either train in isolation, indicating that the responses are at least not completely redundant. However, the information is also much less than the sum of the information of the two separate responses (curves marked by X), suggesting a large amount of redundancy. This could represent real redundancy across the responses, but it could also indicate an apparent redundancy, resulting from decoding the responses without reference to the neuron of origin. For $k>0$, the metrics become increasingly sensitive to which neuron fired each spike. As shown here, this sensitivity increases the amount of information that can be extracted from the joint response, typically to levels approaching the total information in the two individual responses. In the case of the

neuron pair in the top row, the extent of the increase in information as k increases is dramatic, approximately 40%. More typically, as seen in the neuron pair in the bottom row, it is relatively subtle, approximately 10%. For both recordings, the importance of the neuron-of-origin label is greater when the entire response is analyzed (second and third columns of Fig. 8) than when just the initial response is analyzed (first column). This suggests that the initial transient response is relatively stereotyped, while differences in the late response dynamics across neurons are more informative.

The overall conclusion is that nearby neurons' responses are much less redundant, provided that their responses are not simply summed together. Even when the benefit of sensitivity to the neuron of origin is large, only a modest amount of sensitivity ($k=0.2$–0.4) is required, as is shown by the location of the peak of the surfaces along the k-axes. Finally, we note that the added information provided by the neuron of origin is independent of the added information provided by spike timing. That is, for each value of k, the amount of information by $D^{\text{spike}}[0,k]$ (only the neuron of origin is tracked, but not the spike time) is less than the information that can be extracted by $D^{\text{spike}}[q,k]$ for optimal values of q. The optimal sensitivity to timing for decoding multineural responses is indepen-

dent of k and appears similar to that for individual responses, as seen by the ridge of the surfaces running parallel to the k-axis, for q in the range 30 to 100.

This information loss may appear to be modest when only pairs of neurons are considered. However, a related analysis [47]) of similar datasets via extensions of the 'direct method' shows that the information lost by simple summation increases with the number of neurons jointly analyzed (up to 8 neurons). This approach, unlike the labeled metrics, allowed consideration of larger clusters of neurons, but only brief time intervals.

These complementary approaches show that the information available in the response of a local cluster of visual cortical neurons is substantially greater than what might be extracted by simple summation — even though it is firmly established [26,41] that neighboring neurons have many similar response properties. To extract the additional information that reflects more subtle differences in response properties [17], a postsynaptic neuron needs to be sensitive to the timing pattern of its individual inputs, as well as the neuron of origin of the individual spikes. Passive summation has the benefit of simplicity, while decoding processes sensitive to details of spike timing or neuron of origin doubtless entail a metabolic penalty [32]. It appears that neurons are willing to pay this price [1,10,39], and we propose that the reason for this is the additional information extracted.

References

[1] L.F. Abbott, J.A. Varela, K. Sen, S.B. Nelson, Synaptic depression and cortical gain control, Science 275 (1997) 220–224.

[2] M. Abeles, Role of the cortical neuron: integrator or coincidence detector?, Isr. J. Med. Sci. 18 (1982) 83–92.

[3] M. Abeles, Local Cortical Circuits, an Electrophysiological Study, Springer, Berlin, 1982.

[4] M. Abeles, G.L. Gerstein, Detecting spatiotemporal firing patterns among simultaneously recorded single neurons, J. Neurophysiol. 60 (1988) 909–924.

[5] N. Abramson, Information Theory and Coding, McGraw-Hill, New York, 1963.

[6] E.D. Adrian, The Basis of Sensation., W.W. Norton, New York, 1928.

[7] A. Aertsen, G.L. Gerstein, H.K. Habib, G. Palm, Dynamics of neuronal firing correlation: modulaton of effective connectivity, J. Neurophysiol. 61 (1989) 900–917.

[8] W. Bialek, F. Rieke, R.R. de Ruyter van Steveninck, D. Warland, Reading a neural code, Science 252 (1991) 1854–1857.

[9] T.V. Bliss, G.L. Collingridge, A synaptic model of memory: long-term potentiation in the hippocampus, Nature 361 (1993) 31–39.

[10] H.R. Bourne, R. Nicoll, Molecular machines integrate coincident synaptic signals, Cell 72/Neuron 10 Suppl. (1993), 65–85.

[11] S.E. Brodie, B.W. Knight, F. Ratliff, The spatiotemporal transfer function of the Limulus lateral eye, J. Gen. Physiol. 72 (1978) 167–202.

[12] I.H. Brivanlou, D.K. Warland, M. Meister, Mechanisms of concerted firing among retinal ganglion cells, Neuron 20 (1998) 527–539.

[13] T.H. Bullock, Signals and signs in the nervous system: the dynamic anatomy of electrical activity is probably information-rich, Proc. Natl. Acad. Sci. USA 94 (1997) 1–6.

[14] A.G. Carlton, On the bias of information estimates, Psychological Bulletin 71 (1969) 108–109.

[15] Y. Dan, J.M. Alonso, W.M. Usrey, R.C. Reid, Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus, Nature Neuroscience 1 (1998) 501–507.

[16] A.F. Dean, The variability of discharge of simple cells in the cat striate cortex, Exp. Brain Res. 44 (1981) 437–440.

[17] G.C. DeAngelis, G.M. Ghose, I. Ohzawa, R.D. Freeman, Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons, J Neurosci. 19 (10) (1999) 4046–4064.

[18] M.R. DeWeese, M. Meister, How to measure the information gained from one symbol, Network 10 (4) (1999) 325–340.

[19] M.S. Fee, P.P. Mitra, D. Kleinfeld, Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability, J. Neurosci. Methods 69 (1996) 175–188.

[20] R. FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, Biophys. J. 1 (1961) 445–466.

[21] L.J. Frishman, A.W. Freeman, J.B. Troy, D.E. Schweitzer-Tong, C. Enroth-Cugell, Spatiotemporal frequency responses of cat retinal ganglion cells, J. Physiol. 341 (1987) 279–308.

[22] T.J. Gawne, T. W Kjaer, B.J. Richmond, Latency: another potential code for feature binding in striate cortex, J. Neurophysiol. 76 (1996) 1356–1360.

[23] C.M. Gray, P.E. Maldonado, M. Wilson, B. McNaughton, Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recording in cat striate cortex, J. Neurosci. Methods 63 (1-2) (1995) 43–54.

[24] G.R. Holt, W.R. Softky, C. Koch, R.J. Douglas, Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons, J. Neurophysiol. 75 (1996) 1806–1814.

[25] J.J. Hopfield, Pattern recognition computation using action potential timing for stimulus representation, Nature 376 (1995) 33–36.

[26] D.H. Hubel, T.N. Wiesel, Ferrier Lecture. Functional architecture of macaque monkey visual cortex, Proc. Roy. Soc. Lond. B. 198 (1977) 1–59.

[27] H. Jahnsen, R. Llinás, Electrophysiological properties of guinea-pig thalamic neurons: An in vitro study, J. Physiol. (Lond.) 349 (1984) 205–226.

[28] H. Jahnsen, R. Llinás, Ionic basis for the electroresponsiveness and oscillatory properties of guinea-pig thalamic neurons in vitro, J. Physiol. (Lond.) 349 (1984) 227–247.

[29] B.W. Knight, Dynamics of encoding in a population of neurons, J. Gen. Physiol. 59 (1972) 734–766.

[30] B.W. Knight, The relationship between the firing rate of a single neuron and the level of activity in a population of neurons, J. Gen. Physiol. 59 (1972) 767–778.

[31] J. Larson, D. Wong, G. Lynch, Patterned stimulation at the theta frequency is optimal for the induction of hippocampal long-term potentiation, Brain Res. 368 (1986) 347–350.

[32] S.B. Laughlin, R. R deRuyter van Stevenink, J.C. Anderson, The metabolic cost of neural information, Nature Neurosci. 1 (1998) 36–41.

[33] J.E. Lisman, Bursts as a unit of neural information: making unreliable synapses reliable, Trends Neurosci. 20 (1997) 38–43.

[34] Z. Mainen, T.J. Sejnowski, Reliability of spike timing in neocortical neurons, Science 268 (1995) 1503–1506.

[35] H. Markram, J. Lubke, M. Frotscher, B. Sakmann, Regulation of synaptic efficacy by coincidence of postsynaptic AP's and EPSP's, Science 275 (1997) 213–215.

[36] D.A. McCormick, J.R. Huguenard, A model of the electrophysiological properties of thalamocortical relay neurons, J. Neurophysiol. 68 (1992) 1384–1400.

[37] F. Mechler, J.D. Victor, K.P. Purpura, R. Shapley, Robust temporal coding of contrast by V1 neurons for transient but not steady-state stimuli, J. Neurosci. 18 (1998) 6583–6598.

[38] M. Meister, Multineuronal codes in retinal signaling, Proc. Natl. Acad. Sci. USA 93 (1996) 609–614.

[39] B.W. Mel, Synaptic integration in an excitable dendritic tree, J. Neurophysiol. 70 (1993) 1086–1101.

[40] G.A. Miller, Note on the bias on information estimates. Information Theory in Psychology; Problems and Methods II-B (1955) 95–100.

[41] V.B. Mountcastle, Modality and topographic properties of single cortical neurons of cat's somatic sensory cortex, J. Neurophysiol. 20 (1957) 408–434.

[42] M.W. Oram, M.C. Wiener, R. Lestienne, B.J. Richmond, Stochastic nature of precisely timed spike patterns in visual system neuronal responses, J. Neurophysiol. 81 (1999) 3021–3033.

[43] S. Panzeri, A. Treves, Analytical estimates of limited sampling biases in different information measures, Network 7 (1996) 87–107.

[44] K.P. Purpura, N.D. Schiff, The thalamic intralaminar nuclei: a role in visual awareness, The Neuroscientist 3 (1997) 314–321.

[45] D.S. Reich, F. Mechler, K.P. Purpura, B.W. Knight, J.D. Victor, Multiple timescales are involved in coding of contrast in V1 (Abstract), Soc. Neurosci. Abstr. 24 (1998) 1257.

[46] D.S. Reich, F. Mechler, K.P. Purpura, J.D. Victor, Interspike intervals, receptive fields, and information encoding in primary visual cortex, J. Neurosci. 20 (2000) 1964–1974.

[47] D.S. Reich, F. Mechler, J.D. Victor, Neighboring V1 neurons convey mostly independent messages. (Abstract) Soc. Neurosci. (2000) accepted.

[48] D.S. Reich, F. Mechler, J. D. Victor, Temporal coding of contrast in primary visual cortex: when, what and why? J. Neurophysiol. (2000) submitted.

[49] R.C. Reid, J.D. Victor, R.M. Shapley, The use of m-sequence in the analysis of visual neurons: linear receptive field properties, Visual Neurosci. 14 (1997) 1015–1027.

[50] F. Rieke, D. Warland, R.R. de Ruyter van Steveninck, W. Bialek, in: Spikes: exploring the neural code, Vol. ix+, MIT Press, Cambridge, 1997, p. 394.

[51] G.M. Rose, T.V. Dunwiddie, Induction of hippocampal long-term potentiation using physiologically patterned stimulation, Neurosci. Lett. 69 (1986) 244–248.

[52] N.D. Schiff, K.P. Purpura, J.D. Victor, Gating of local network signals appears as stimulus-dependent activity envelopes in striate cortex, J. Neurophysiol. 82 (1999) 2182–2196.

[53] P.H. Sellers, On the theory and computation of evolutionary distances, SIAM J. Appl. Math. 26 (1974) 787–793.

[54] K. Sen, J.C. Jorge-Rivera, E. Marder, L.F. Abbott, Decoding synapses, J. Neurosci. 16 (1996) 6307–6318.

[55] M.N. Shadlen, W.T. Newsome, The variable discharge of cortical neurons: implications for connectivity, computation, and information coding, J. Neurosci. 18 (1998) 3870–3896.

[56] C.E. Shannon, W. Weaver, The Mathematical Theory of Communication, University of Illinois Press, Urbana, IL, 1949.

[57] W.R. Softky, C. Koch, The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSP's, J. Neurosci. 13 (1993) 334–350.

[58] S. Song, J.A. Varela, L.F. Abbott, G.G. Turrigiano, S.B. Nelson, A quantitative description of synaptic depression at monosynaptic inhibitory inputs to visual cortical pyramidal neurons, Neurosci. Abstr. 23 (1997) 2362.

[59] R.H. Stein, The information capacity of nerve cells using a frequency code, Biophys. J. 7 (1969) 797–826.

[60] S.P. Strong, R. Koberle, R.R. de Ruyter van Steveninck, W. Bialek, Entropy and information in neural spike trains, Phys. Rev. Lett. 80 (1998) 197–201.

[61] E. Sutter, A practical nonstochastic approach to nonlinear time-domain analysis, in: V. Marmarelis (Ed.), Advanced Methods of Physiological System Modeling, Biomedical Simulations Resource, Los Angeles, 1987, pp. 303–315.

[62] E. Sutter, The fast m-transform: A fast computation of cross-correlations with binary m-sequences, SIAM J. Comput. 20 (1991) 686–694.

[63] E. Sutter, A deterministic approach to nonlinear systems analysis, in: R. Pinter, B. Nabet (Eds.), Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks, CRC Press, Cleveland, 1992, pp. 171–220.

[64] W.M. Usrey, J.B. Reppas, R.C. Reid, Paired-pulse interactions between retina and thalamus, Neurosci. Abstr. 23 (1997) 170.

[65] J.D. Victor, Temporal aspects of neural coding in the retina and lateral geniculate: a review, Network 10 (1999) R1–R66.

[66] J.D. Victor, K. Purpura, Nature and precision of temporal coding in visual cortex: a metric-space analysis, J. Neurophysiol. 76 (1996) 1310–1326.

[67] J.D. Victor, D.S. Reich, F. Mechler, D. Aronov, Multineuronal signaling of spatial phase in V1. (Abstract). Soc. Neurosci. (2000) accepted.

[68] J.D. Victor, K.P. Purpura, Metric-space analysis of spike trains: theory, algorithms, and application, Network 8 (1997) 127–164.

[69] J.D. Victor, K.P. Purpura, Spatial phase and the temporal structure of the response to gratings in V1, J. Neurophysiol. 80 (1998) 554–571.

[70] P. Viviani, Eye movements in visual search:cognitive, perceptual, and motor aspects, in: E. Kowler (Ed.), Eye Movements and their Role in Visual and Cognitive Processes, Elsevier, Amsterdam, 1990, pp. 353–393.

[71] S. Wuerger, L.T. Maloney, J. Krauskopf, Proximity judgments in color space: tests of a Euclidean color geometry, Vision Res. 35 (1995) 827–835.