

Discovering Spike Patterns in Neuronal Responses

Jean-Marc Fellous,¹ Paul H. E. Tiesinga,^{1,2} Peter J. Thomas,¹ and Terrence J. Sejnowski^{1,3}

¹Computational Neurobiology Laboratory, Howard Hughes Medical Institute, The Salk Institute for Biological Studies, La Jolla, California 92037,

²Department of Physics and Astronomy, University of North Carolina, Chapel Hill, North Carolina 27599-3255, and ³Division of Biological Sciences, University of California, San Diego, La Jolla, California 92093

When a cortical neuron is repeatedly injected with the same fluctuating current stimulus (frozen noise) the timing of the spikes is highly precise from trial to trial and the spike pattern appears to be unique. We show here that the same repeated stimulus can produce more than one reliable temporal pattern of spikes. A new method is introduced to find these patterns in raw multitrail data and is tested on surrogate data sets. Using it, multiple coexisting spike patterns were discovered in pyramidal cells recorded from rat prefrontal cortex *in vitro*, in data obtained *in vivo* from the middle temporal area of the monkey (Buracas et al., 1998) and from the cat lateral geniculate nucleus (Reinagel and Reid, 2002). The spike patterns lasted from a few tens of milliseconds *in vitro* to several seconds *in vivo*. We conclude that the prestimulus history of a neuron may influence the precise timing of the spikes in response to a stimulus over a wide range of time scales.

Key words: cluster; lateral geniculate (LGB, LGN); network; synapse; vision; MT

Introduction

Repeated injection of a fluctuating current into the soma of a cortical neuron *in vitro* yields similar spike patterns, precise to within a few milliseconds (Mainen and Sejnowski, 1995). Similar observations have been made *in vivo* using visual stimuli in the retina, lateral geniculate nucleus (LGN), and visual cortices (Berry et al., 1997; Gur et al., 1997; Buracas et al., 1998; Kara et al., 2000; Reinagel and Reid, 2000, 2002) and in invertebrate preparations (Bryant and Segundo, 1976; Wehr and Laurent, 1996; de Ruyter van Steveninck et al., 1997; Warzecha et al., 2000). The spike pattern depends on stimulus features such as frequency and amplitude (Hunter et al., 1998; Fellous et al., 2001; Hunter and Milton, 2003) and on the intrinsic properties of the neuron, such as its refractory period (Berry and Meister, 1998; Reich et al., 1998; Kara et al., 2000), intrinsic channel noise (White et al., 2000), and overall level of excitability (Schreiber et al., 2004). When various types of perturbations were made to the injected current in both *in vitro* neurons and computational models, the neural responses relaxed to the original spike pattern within a few tens of milliseconds (Tiesinga, 2002; Tiesinga et al., 2002).

Different spike patterns can occur *in vitro*, for example, when a neuron is reliably entrained in the 1:2 regimen by an injected current sine wave of sufficiently high frequency to generate one spike every two cycles. In the noiseless case, whether the spikes occur on the odd or even cycles of the sine wave is determined solely by the initial conditions (Tiesinga, 2002) and depends on the prestimulus history of the neuron. Finding different patterns

in real data has been difficult because neurons typically may jump from one pattern to another at times that cannot be predicted and because single patterns are themselves difficult to identify in the presence of neuronal noise. To circumvent these difficulties, spike patterns have been viewed as patterns of firing rates (Zipser et al., 1993), rather than patterns of spike times. Such analyses, however, are very sensitive to the averaging window (several hundreds of milliseconds), and do not capture fine temporal spiking patterns on the order of a few spikes, which, if they exist, would have significant consequences for our understanding of neural information processing. Neurons in a complex interconnected network can also produce several possible spike patterns, also called attractors (Schoner and Kelso, 1988; Zipser et al., 1993; Seidemann et al., 1996; Amit et al., 1997; Bodner et al., 1998; Rabinovich and Abarbanel, 1998; Diesmann et al., 1999; Freeman, 2000; Tsuda, 2001).

We introduce a new method for finding temporal spike patterns in multitrail recordings that are time-locked to the onset of a stimulus. This method is tested on surrogate data sets that include controlled amounts of noise. We apply it to *in vitro* data, for which the pattern structure is known by design; to *in vitro* data, for which the pattern structure is discovered; and to previously published data recorded *in vivo*, for which the patterns were previously undetected. The neurons both *in vitro* and *in vivo* generated several different but reliable spike patterns in response to repeated presentations of the same stimulus.

Materials and Methods

Neurophysiology

In vitro experiments. Coronal slices of rat prelimbic and infralimbic areas of prefrontal cortex were obtained from 2- to 4-week-old Sprague Dawley rats. Rats were anesthetized with isoflurane (Abbott Laboratories, Abbott Park, IL) and decapitated. Their brains were removed and cut into 350 μm thick slices using standard techniques. Patch-clamp was performed under visual control at 30–32°C. In most experiments Lucifer

Received Oct. 14, 2003; revised Jan. 23, 2004; accepted Jan. 25, 2004.

This work was funded by the Howard Hughes Medical Institute. We thank Drs. Pamela Reinagel, Clay Reid, Giedrus Buracas, and Tom Albright for helpful discussions and comments on this manuscript and for providing data.

Correspondence should be addressed to Dr. Jean-Marc Fellous, Computational Neurobiology Laboratory, Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037. E-mail: fellous@salk.edu.

DOI:10.1523/JNEUROSCI.4649-03.2004

Copyright © 2004 Society for Neuroscience 0270-6474/04/242989-13\$15.00/0

yellow (0.4%; Sigma/RBI, Poole, UK) was added to the internal solution. At the end of those experiments a fluorescent blue light was shined on the tissue and the cell morphology was visually inspected. Synaptic transmission was blocked by D-2-amino-5-phosphonovaleric acid (50 μ M), 6,7-dinitroquinoxaline-2,3-dione (10 μ M), and bicuculline methiodide (20 μ M). All drugs were obtained from Sigma/RBI, freshly prepared in artificial CSF and bath-applied. Whole-cell patch-clamp recordings were achieved using glass electrodes (4–10 M Ω) containing the following (in mM): 140 KmeSO₄, 10 HEPES, 4 NaCl, 0.1 EGTA, 4 Mg-ATP, 0.3 Mg-GTP, and 14 phosphocreatine. Data were acquired in current-clamp mode using an Axoclamp 2A amplifier (Axon Instruments, Foster City, CA). Most of the experiments presented here consisted of the repeated injection of the same current waveform. The dead time (0 current injection) between consecutive current injections was set to be at least twice the length of the current waveform. Only cells that had a stable (<0.5 mV SD) membrane potentials over the whole experiment were kept for study. We used regularly spiking layer 5 pyramidal cells.

Data were acquired using a Pentium-based computer. Programs were written using LabView 6.1 (National Instrument, Austin, TX), and data were acquired with a PCI16-E1 data acquisition board (National Instrument). The data acquisition rate was either 10 or 20 kHz. Stimuli were designed on-line, or off-line as text files. All experiments were performed in accordance with animal protocols approved by the National Institutes of Health. A total of 16 pyramidal cells were used in this study (*in vitro* experiments). Data were analyzed off-line using MATLAB (MathWorks, Natick, MA). Results are given as means \pm SD.

In vivo experiments: area MT. We analyzed published data obtained *in vivo* in awake monkeys by (Buracas et al., 1998). Rhesus monkeys (*Macaca mulatta*) were implanted with electrodes in area MT using standard techniques (Albright, 1984; Dobkins and Albright, 1994). The animals were alert, fixating, and passively viewing during stimulus presentation. Eye position was monitored. The data were obtained for stimuli based on Gabor patches subtending 4° of visual angle presented on a 21 inch computer monitor (refresh rate, 60 Hz). Motion was achieved by a spatial phase shift occurring in each video frame (12°/sec). The time-varying stimuli that resulted in the recordings used here consisted of rapid switching of Gabor patches from preferred to anti-preferred directions (~30–300 msec switch rate; stimulus length, 3 sec). The exact parameters of the switching were determined for each cell, to maximize its ability to discriminate between the two directions. Additional details can be found in Buracas et al. (1998).

In vivo experiments: LGN. We analyzed published data obtained *in vivo* in anesthetized cats by Reinagel and Reid (2002). Briefly, LGN X ON cells were recorded extracellularly in the anesthetized (sodium pentobarbital) cat while their receptive fields were stimulated by visual stimuli. The visual stimuli were uniform wide-field illuminations modulated in time and were repeated 128 times. The stimulus length was 8 sec; data were sampled at 1 kHz. The temporal modulation was set according to the distribution of luminance measured in a natural stimulus. Details of the surgical procedures and experimental designs have been published previously (Reinagel and Reid, 2000, 2002).

Clustering algorithm

Spike patterns can be identified on the basis of the rastergram by grouping together trials that are similar to each other. In the following we present an algorithm that achieves this grouping, we compare it on a surrogate data set for which the clusters are known, and we apply it to *in vitro* and *in vivo* data. The details of this algorithm are given in the Appendix.

Measure of trial similarity and reliability. The inputs (each input point is a trial) were preprocessed to focus on spike timing rather than on spike shape or subthreshold membrane fluctuations. After extracting spike times from the voltage traces using standard methods, we generated smoothed traces by convolving each spike train (a sequence of δ functions) with a Gaussian function of fixed width σ . This procedure replaced each spike by a standard waveform of known time span and increased the speed of the algorithm. The choice of σ will be discussed below.

After convolving all trials with the Gaussian, the dot product between

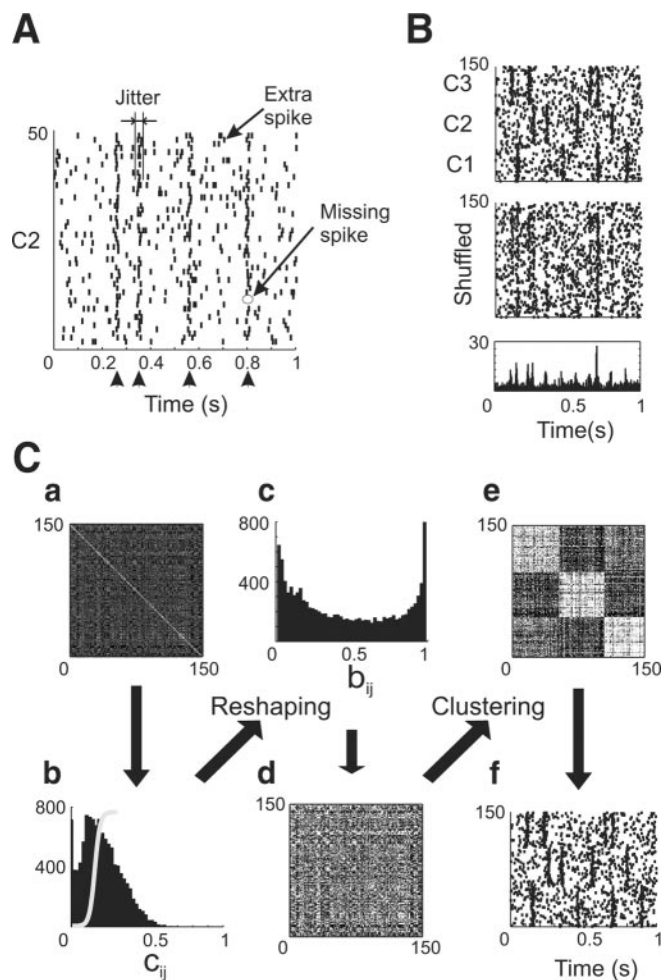


Figure 1. Clustering method. *A*, Example of a surrogate cluster of 50 trials generated around four randomly timed events (arrowheads below the time axis) with a controlled amount of jitter (5 msec), six extra spikes per trial, and 20% missing event spikes. *B*, Two other clusters (C1 and C3) were similarly generated around different events (top), and the three clusters were randomly shuffled (middle) to yield the input to the clustering algorithm. The original and shuffled rastergrams had the same spike time histogram (bottom). *C*, Clustering method. The similarity matrix was first computed on the basis of the shuffled set of trials (*a*); the distribution of similarity values was determined (*b*) and passed through a sigmoid function centered on the peak of the distribution yielding the flattest possible histogram (*c*; slope = 0.14). The sigmoid was then applied to the similarity matrix (*d*), which was used as inputs to the clustering algorithm. The result of the clustering was a reordering of the columns of the matrix such that similarities within clusters were maximized and similarities between clusters were minimized (*e*; dark pixels indicate low similarity values). The same reordering was applied to the input trials (*f*; performance was 97% correct; fuzziness factor, 2).

all pairs of trials $\langle i, j \rangle$ was taken and the symmetric similarity matrix S (s_{ij}) was formed (Schreiber et al., 2003). For N trials,

$$s_{ij} = \frac{\tilde{g}_i \cdot \tilde{g}_j}{\|\tilde{g}_i\| \cdot \|\tilde{g}_j\|}$$

where $\tilde{g}_i \in \mathcal{R}^P$ represented the convolution of the spike train i with the Gaussian kernel, with each trial sampled with a total of P points. The normalized dot product s_{ij} was the cosine of the angle between vectors \tilde{g}_i and \tilde{g}_j . Because all the coordinates of these two vectors were positive, s_{ij} was a number between 0 and 1 and served as a measure of the similarity (correlation) between the input vectors \tilde{g}_i and \tilde{g}_j . Its value approached 1 as the inputs became equal and 0 as the inputs became orthogonal. The distribution of s_{ij} was often approximately Gaussian, centered at values significantly smaller than 0.5 (Fig. 1*Cb*). The dynamic range of similarity values was small, and most points were confined to a small region of the

range of similarity values. To augment the dynamic range, we rescaled the similarity matrix by applying a sigmoid function to all similarity values (see Appendix).

Other distance measures could in principle be used in place of the similarity measure used here. Such measures include Minkowski-generalized L_r measures, Pearson correlation measures (Everitt et al., 2001), or other “binless” measures specifically designed for spike-train analysis (Victor and Purpura, 1996). These measures were not studied here.

Based on this similarity measure, a new measure of reliability was introduced as the average of all similarity values (excluding the diagonal):

$$R = \frac{2}{N(N-1)} \sum_{i=1}^N \sum_{j=i+1}^N s_{ij}$$

where R is bounded between 0 (low reliability; all trials were dissimilar) and 1 (high reliability; all trials were identical). A comparison between this measure of reliability and the histogram-based measure used in Mainen and Sejnowski (1995) showed that this measure degraded in a smoother manner with the addition of noise, with erasure of event spikes, and with increased spike jitter, and was more robust than the histogram measure when fewer trials were considered (Schreiber et al., 2003).

Clustering. If a spike train is cut into sufficiently brief time bins, each spike train is a vector, the coordinates of which are 0 if no spike occurred and 1 if a spike occurred. This vector would have as many coordinates as sampling points, and clustering methods such as the K -means algorithm would typically require a very large number of samples. We reduced the dimensionality of the space in which to perform spike train clustering by representing each of the N spike trains j as an N -dimensional vector $\vec{p}_j \in \mathfrak{R}^N$ (j th column of S) of similarities between trials j and all other trials $i = 1, \dots, N$. The result of this “recoding” can be represented by the symmetric $N \times N$ matrix S showing these vectors in columns (Fig. 1). In this representation, trials that belong to the same pattern have vectors that are close in the Euclidian sense. Three different clustering algorithms were tested. Compared with the basic K -means and the extended K -means algorithms, the fuzzy K -means (FKM) technique yielded the most stable results with respect to initial conditions and the best performance on surrogate data sets (see below and Appendix). Other clustering algorithms have not been tested (Hartigan, 1975; Jain, 1988; Everitt et al., 2001; Theodoridis and Koutroumbas, 1998).

Surrogate data sets

To test the clustering method, we designed sets of surrogate data containing a known number of clusters. Each cluster consisted of a sequence of E randomly chosen spike-time events T_E (E between 4 and 6) (Fig. 1A, arrowheads below the time axis). Each cluster was expressed as a rastergram and was parameterized with M (percentage of missing event spikes across all trials), J (average SD in msec of the jitter across all events), and X (number of extra non-event-related spikes per trial). Each cluster contained I trials (Fig. 1A,B, $I = 50$). Each trial i of cluster c was a set of spike times R_i^c :

$$R_i^c = \{t_{p,e}^i\} \cup \{t_x^i\}$$

where p is a uniformly distributed random variable in the interval $(0, 1)$, such that if $p < 1 - M$, then $t_{p,e}^i = T/e + \varepsilon$ otherwise the spike time was not included, and where ε was a normally distributed random variable with 0 mean and SD J . The t_x^i were X random times uniformly distributed in the time interval (Fig. 1A,B, between 0 and 1000 msec). N clusters were then joined in a combined rastergram, and the trials were shuffled (Fig. 1B). Note that the histogram of spike times was identical in the non-shuffled and shuffled cases, and did not reflect the temporal structure of any of the two clusters (Fig. 1B). In this study, 20 rastergrams were generated for each triplet (M, J, X) , and results are presented as the average and SD across these 20 instances. Note that the actual jitter, and amount of missing spikes were recalculated from each of the formed rastergrams.

The data set was first clustered with a small number of clusters (two to five), the strongest cluster was extracted (using the clustering strength D

in the Appendix), and a new round of clustering was performed on the remaining data, until no data points are left, or until the difference in strength between the new clusters became smaller than a given criterion. In practice for *in vitro* and *in vivo* data, one pass was sufficient to yield a good clustering of the data. To test this method, we generated three families of surrogate sets with two, three, and five clusters respectively, leaving for additional studies the case of more than five clusters. The clustering programs and other data analysis routines were written in MATLAB (MathWorks) and executed on Pentium-based computers running windows 98SE, Windows 2000 (Microsoft, Seattle, WA) or FreeBSD. The surrogate data sets are available from <http://www.cnl.salk.edu/~fellous/data/JN2004data/data.html>, to facilitate comparisons with other clustering algorithms.

Clustering procedure

The same clustering procedure was used for surrogate data sets and for *in vitro* and *in vivo* data. Briefly, with N trials available (Fig. 1B, shuffled), (1) The $N \times N$ matrix S of pairwise trial similarity was computed (Fig. 1Ca); (2) the histogram of all similarity values was computed, its mean was determined, and a sigmoid function centered at the mean was calculated by iteratively varying the slope of the sigmoid until the histogram flattened (see Appendix, Fig. 1Cb,Cc); (3) the sigmoid function was then applied to the similarity matrix (Fig. 1Cd); (4) the FKM algorithm was run on the reshaped matrix, considering each column as an input point. The inputs were labeled and sorted according to their cluster membership, and the matrix was reordered accordingly (Fig. 1Ce). The cluster strength D was computed for each cluster (see Appendix); and (5) the same reordering was applied to the original rastergram (Fig. 1Cf).

For surrogate data sets, when the clusters are known (Fig. 1B, top), the original and extracted clusters were compared using all permutations, because cluster order is not guaranteed, and performance was measured as the fraction of correctly grouped trials for the best permutation.

Results

Intrinsic noise can switch a neuron from one stable pattern to another

When a cell was injected with a sine wave current, its spike timing was most reliable when the frequency matched the frequency of the intrinsic subthreshold oscillations obtained by depolarizing the cell to the same average DC membrane potential (Hunter et al., 1998; Fellous et al., 2001). In this regimen, the cells were entrained 1:1. Higher sine wave frequencies yielded cycle skipping and loss of reliability. Figure 2A shows the response of a pyramidal cell to a 7 sec long sine wave of 6 Hz that was twice its most reliable frequency (1:1 regimen was reliably obtained at 3 Hz), whereas synaptic transmission was blocked pharmacologically. The top two traces show typical voltage responses. After a transient of ~ 3 sec the cell spiked every other cycle and maintained this 2:1 regimen for the duration of the current injection (4 sec). The initial voltage before current injection was stable (-70 ± 0.3 mV; $n = 100$ trials), but slight variations, presumably caused by intrinsic noise, eventually yielded spiking on either the even (top trace) or the odd cycles (middle trace) of the input current.

Figure 2B shows the rastergram obtained by injecting the same sine wave 100 times. Considering all the trials together, the reliability of this cell was 0.49 (computed after the transient, between 3 and 7 sec; 0.48 ± 0.02 ; $n = 8$). Figure 2C shows the same rastergram but the trials have been reordered depending on whether the cell mainly fired on the even or odd cycles of the sine wave injection during the last 4 sec of the recording. The two patterns of spiking are clearly apparent, and when considered individually yielded a reliability of 0.89 (0.86 ± 0.03 ; $n = 8$) and 0.92 (0.87 ± 0.0 ; $n = 8$). Each pattern contained the same average number of spikes (i.e., had the same average firing rate), but differed only in spike timing.

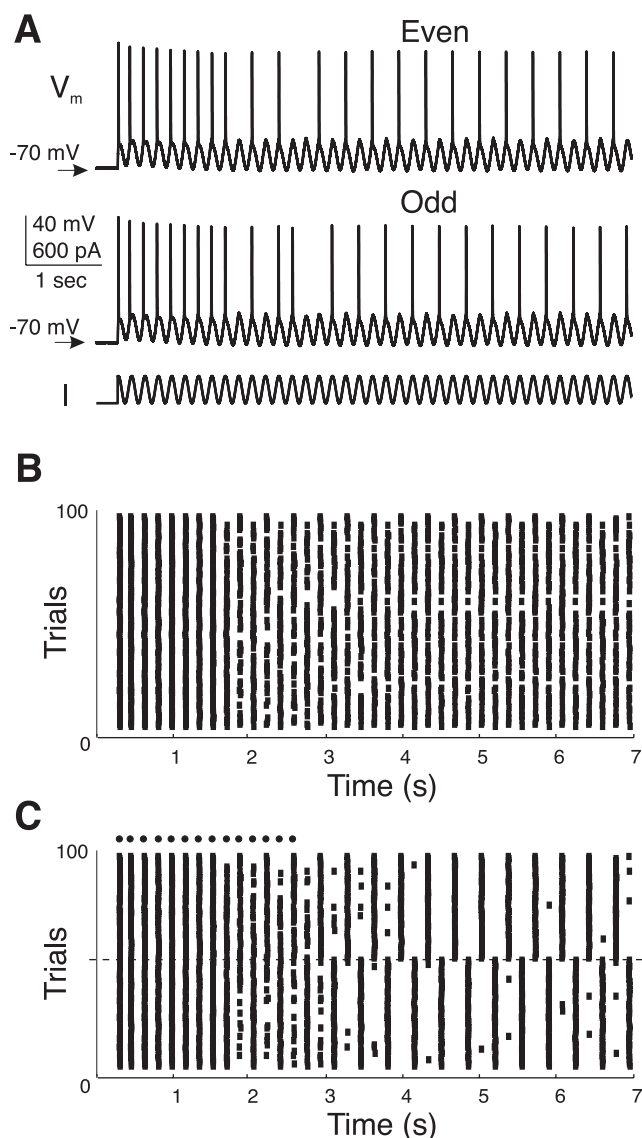


Figure 2. Creation of spike patterns *in vitro*. *A*, A prefrontal pyramidal neuron was somatically injected with a sine wave current waveform (bottom trace). The frequency was twice the preferred frequency of the cell for this sine wave amplitude and DC level (Fellous et al., 2001). After a transient, the cell spiked once every two cycles, either on the even (top) or odd (middle) cycles. The overall reliability of the cell was 0.49. *B*, Rastergram showing 100 trials for the same cell as in *A*. *C*, Same rastergram as in *A*, with trials reordered according to whether the cell settled on the even or odd cycles. Two separate reliable sets of trials are apparent for the even (reliability, 0.89) and odd (reliability, 0.92) cycles. The dashed line separates the two sets. Dots indicate spike times that are common to the two sets.

Note that although the experiment typically lasted for several hours, the resting membrane potential of the cell remained stable and there were no correlations between the initial voltage before current injection and the spike pattern (odd or even cycles) converged on. Similar observations were made for seven other cells, although the transient could vary significantly depending on the cell and on the amplitude and frequency of the stimulus (range, 2–4 sec; $n = 8$). These transients were likely to depend on the kinetics and relative conductances of the various membrane currents present in each cell. Because synaptic transmission was blocked in these experiments, the variability in the length of the transient for a given cell was caused by intrinsic noise (Manwani and Koch, 1999). A quantitative analysis of these transients was beyond the scope of this study.

These results indicated that cells in some conditions can produce a small number of distinct and reliable firing patterns in response to the same stimulus. After a brief transient, these patterns could last for several seconds. Note that for smaller drive amplitudes the cell could jump from one spike pattern to the other.

Clustering

The goal was to identify and extract temporal spike patterns from multitrail rastergrams such as shown in Figure 2*B*. We first applied the method to *in vitro* data, for which the stimulus is known, and next to *in vivo* data, for which the input stimulus to the cell is unknown. We used the fuzzy K -means clustering method with an initial fuzziness factor equal to 2 and ran the algorithm only once (see Appendix). In the experimental data used here, it is not known whether clusters exist. Therefore, the clustering method was assessed on a surrogate data set for which the cluster composition is known, and a confidence measure of the clustering was established to differentiate weak (almost random) from significant clustering results.

To assess the performance of the clustering method we used surrogate data sets (see Materials and Methods) and chose the value of the Gaussian width σ used to compute trial-to-trial similarities to match the worst cases observed *in vitro* in cortex (Mainen and Sejnowski, 1995; Fellous et al., 2001) and *in vivo* in LGN (Reinagel and Reid, 2002). The performance of the clustering was defined as the percentage number of trials correctly attributed to their original clusters, before the shuffling took place (Fig. 1*A,B*). Because across multiple trials, the peristimulus histograms (PSTHs) of these experiments typically had strong peak widths of 10 msec or less, we chose $\sigma = 5$ msec. Figure 3 shows two example surrogate sets for two (*A*, three extra spikes/trials; 15% missing spikes; 10 msec jitter) and five (*B*) clusters (fuzzy K -means with initial fuzziness factor $f = 2$). Panels labeled *a* show the original surrogate rastergrams; *b*, the corresponding similarity matrix; *c*, the clustering achieved by the fuzzy cluster algorithm; and *d*, the rastergram reordered accordingly. In these two examples the final fuzziness factor was $f = 2$ and 1.65, and the performance was 100% and 93.1% for two and five clusters, respectively. In both cases the slope of the reshaping sigmoid was 0.14 (see Appendix).

Figure 4 shows the overall performance results for two, three, and five clusters in panels *A–C*, respectively. Seven different numbers of extra spikes (0–30 extra spikes/trial) and seven different levels of event-spike jitter (0–40 msec) were used. The fraction of missing event spikes was kept constant at 15%. Left panels show the performance as a function of the amount of jitter present in the data, for various numbers of extra spikes per trials (different curves). The value of σ was set to the average width of the PSTH peaks, for each of the rastergrams considered (Fig. 4). The initial fuzziness factor was set to 2. In control experiments we ran the clustering algorithm on random rastergrams containing varying amounts of extra spikes per trials, as in the surrogate set, but not containing any event. The left panels show the classification result on this random data set (arrowheads) for 2 (55%), 3 (37%), and 5 (27%) clusters. Each point represents an average of 20 random rastergrams. These results are slightly different from what would be expected from an infinite data set (50.0, 33.3, and 20.0%, respectively). The data of the left panels are replotted in the right panels as a function of the number of extra spikes per trial, for various amounts of jitter (different curves). The dotted curve represents chance levels computed from random rastergrams containing no events. The cluster strengths D of these




Figure 3. Sample clustering results on surrogate data sets containing two clusters (*A*, 4 events per cluster; 10 msec jitter; three extra spikes/trial; 15% missing spikes; 35 trials per cluster) and five clusters (*B*, 4–5 events per cluster; 10 msec jitter; three extra spikes/trial; 15% missing spikes; 35 trials per cluster). For each panel: *a*, the input to the clustering algorithm; *b*, the corresponding distance matrix; *c*, the result of the clustering on the distance matrix; and *d*, the corresponding reordering of the rastergram are shown.

eventless rastergrams were all smaller than 1.5. However, in rastergrams in which events were present and for which 90% of the trials were correctly clustered, *D* was always much greater than 2.0 (i.e., for each cluster, the average distance of all the cluster members to their center was at most half the average distance of all cluster members to all the other centers; see Appendix). This criterion served as a confidence measure; a clustering was considered valid when all clusters had $D > 2$.

Spike patterns *in vitro* can be created by narrowband or broadband oscillatory inputs

The clustering algorithm was run on the data in Figure 2 ($n = 8$ cells; $f = 2$; $\sigma = 5$ msec; $K = 2$) using the whole record (7 sec long, including the transient). In all cases the two clusters found were balanced (equal number of trials ± 2) and corresponded to the even and odd cycles, as expected. The mean cluster strength *D* was 15.1 ± 1.4 ($n = 8$). Note that the patterns initially shared “common events,” for which an event is defined as a peak in the PSTH that contains at least 40% of the trials. Fourteen events common to the two clusters are shown in Figure 2*C* (dots). After these events, the two patterns differed and did not overlap (no common events) for the remainder of the recordings. These experiments showed that, given an appropriate stimulus, it was possible to create two stable patterns that differed from each




Figure 4. Overall performance of the fuzzy *K*-means algorithm. The fuzzy clustering algorithm was used to recover two, three, and five clusters (*A–C*, respectively) in the surrogate data sets. The results are plotted as a function of jitter for different numbers of extra spikes (left), and replotted as a function of the number of extra spikes for different amounts of jitter (right). Chance levels of performance are shown on the left (arrowheads), and on the right (dashed curves) for different numbers of extra spikes. Each point is the average across 20 random rastergrams.

other for an arbitrarily long time. The data presented here can be extended to generate an arbitrary number of patterns by increasing the frequency of the sine wave and entraining the neuron in a 1:*q* regimen (*q* patterns could then be created; Tiesinga, 2002). For sufficiently high amplitude these spike patterns were stable and nonoverlapping (data not shown). In the following we show that hidden spike patterns are present for more general stimulus conditions.

In a second set of experiments, pyramidal cells *in vitro* were recorded intracellularly and were repeatedly injected with a fixed current waveform that consisted of a 3 sec linear mixture of two broadband stimuli: 75% of one with a bandpass in the range 2–8 Hz and 25% of another with a bandpass in the range 30–50 Hz. Such waveforms qualitatively reproduced the membrane potential fluctuations observed during spontaneous theta oscillatory episodes in the rat *in vivo* (Kamondi et al., 1998), but were not meant to address this body of data. An example of this stimulus and the response of a sample cell are shown in Figure 5, *B* and *A*, respectively. The response of the cell to these injections exhibited reliable events that were apparent on the rastergram of 150 trials,




Figure 5. Evidence for spike patterns *in vitro*. *A*, A layer 5 prefrontal pyramidal cell was repeatedly injected with the same current waveform; 150 trials are shown. The histogram is shown below the rastergram. Labels 1–4 indicate the sample times when histogram peaks are clearly detected (arrows). *B*, The stimulus was a linear superposition of two broadband stimuli. *C*, *E*, Expanded view of two sample sections depicted by the left and right boxes in *A*, respectively. *D*, *F*, show the reordering of the trials in *C* and *E*, using the fuzzy *K*-mean clustering algorithm. The algorithm was run separately on *C* and *E*. The clustering algorithm reveals the presence of temporal structure within trials. In the two cases, the trials can be separated in two groups that show distinct firing patterns (*a*–*d*, arrows). These patterns were not apparent from the raw rastergrams (*C*, *E*). The stimulus is reproduced under each rastergram.

and as peaks in the PSTH (Fig. 5*A*, top and bottom traces, respectively). After a transient of ~ 2 sec, four of these events were selected for clustering. Figure 5*C* shows events 1–3 before clustering, and 5*D* shows the result of the reordering obtained by running the FKM algorithm on the rastergram ($f = 2$; $\sigma = 5$ msec; $K = 2$).

Two separate clusters were clearly distinguishable. One cluster corresponded to events 2 and 3 (cluster *a*). The second cluster consisted in what appeared in *C* as interevent extra spikes, but was in reality a separate spiking pattern (*b*) that overlapped very little with *a* (no spikes in event 3; rare spikes in event 2). Note that the reason the cell was not firing between events 2 and 3 in cluster *a* was not because of its refractory period, because the cell could fire with short interspike intervals in trials *b*. If the cell fired in event 2, it would fire only on event 3 (with two exceptions). If the cell fired just before event 2 (time, 2.1 sec), it would not fire on event 2 or 3. Figure 5*E* shows an expansion of event 4. The rastergram seems to show two subevents. The FKM algorithm was run, and the reordering is shown in Figure 5*F* ($f = 2$; $\sigma = 5$ msec; $K = 2$). What appeared to be two subevents in Figure 5*E* resolved as two distinct and mutually exclusive spike times. When the cell spiked at one time, it never spiked at the other. The individual

reliability of each of clusters *c* and *d* was ~ 1.0 , but would have led to a reliability of ~ 0.5 if all trials were taken together. A careful observation of the stimulus (reproduced in Fig. 5*D,F*, bottom traces) revealed that locally, the temporal structure of the pattern was a result of oscillatory dynamics (cycle skipping) similar to that shown in Figure 2; each event in each pattern occurred near the peak of a local oscillatory bout of stimulation.

Grouping events 1, 2, 3, and 4 did not reveal any common clusters, indicating that in this particular experiment, clusters *a* and *b* on the one hand and clusters *c* and *d* on the other were independent and local. Qualitatively similar results were obtained with nine other cells presented with similar stimuli. In all cells tested, no global patterns could be found: running the clustering algorithm on the whole spike train (3 sec long) did not show any strong clusters ($D < 1.25$, when two, three, and five clusters were assumed). Patterns could be detected for only short sequences of events spanning several tens of milliseconds, such as the ones presented in Figure 5, and depended on the exact nature of the stimulus. Thus, for these *in vitro* experiments, and in more general stimulation conditions such as α -filtered noise (Mainen and Sejnowski, 1995), spike patterns are expected to be transient. Similar results were obtained for eight cells, stimulated by three different input waveforms.

Evidence for transient nonoverlapping spike patterns *in vivo*: monkey area MT

We used 10 cells from the study of Buracas et al. (1998), for which at least 60 trials were collected (range, 60–68). Figure 6*A* shows the response of one of these cells to 64 presentations of the same pattern of motion. As observed previously in their study, spikes align with a high degree of precision (low jitter) at specific times during the stimulus presentation (Fig. 6*A*, arrows). These times correspond to peaks of the peristimulus spike histogram across all trials, and can be obtained in a manner similar to event detection in Mainen and Sejnowski (1995). Peaks were considered events if they contained at least 40% of the trials. In this cell, 19 events were detected (20.5 ± 4.8 ; $n = 10$), and the jitter of spike times within each event was on average 5.1 msec (7.1 ± 3 msec; $n = 10$). To find spike patterns, we first windowed the data to include a chosen number of events *E* (between one and five), and then looked for *K* clusters (between two and five) in each window. There were too few trials available to detect a higher number of clusters. The FKM clustering algorithm was run on each of these windows ($f = 2$) with σ set to twice the actual average precision of the spikes in all the events for each cell (here, 5.1 msec). We adopted a stringent criterion for cluster validity: clusters with $D > 3$ (i.e., the distances between clusters were on average at least three times greater than the distances within clusters), and at least six trials per clusters were considered valid. The cell shown in Figure 6 presented 25 valid clusters (34.7 ± 16 ; $n = 10$; range, 10–59). Across all conditions, the average cluster strength was 3.6 (3.9 ± 0.7 ; $n = 10$), the average number of events in the windows yielding valid clustering configurations was 1.8 (E , 1.4 ± 0.5 ; $n = 10$), and the average number of clusters participating in valid configurations was 3 (K , 2.8 ± 0.8 ; $n = 10$). The average window length for valid cluster configurations was 235 ± 106 msec for this cell (189 ± 9 msec; $n = 10$). Figure 6 shows two examples of valid cluster configurations for $E = 2$, $K = 3$ (left dashed box in *A*, reproduced in *B* and clustered in *C*) and for $E = 4$, $K = 5$ (right dashed box in Fig. 6*A*, reproduced in *D* and clustered in *E*). Note that the two reorderings applied to the two boxes shown were different. There were no clustering overlaps in




Figure 6. Evidence for spike patterns in monkey area MT *in vivo* (Buracas et al., 1998). *A*, Rastergram response of MT cells during passive viewing of an alternating sequence of moving Gabor patches in the preferred and anti-preferred directions. Arrows mark spike times when a reliable event was detected. Two sections of this rastergram are indicated by dashed boxes and are reproduced in *B* (left box) and *D* (right box). *C*, *E*, Corresponding reorderings of these two sections according to the fuzzy clustering algorithm (see Materials and Methods) that was run for $K = 2, f = 2$ (*C*; average $D = 3.5$) and $K = 5, f = 2$ (*E*; average $D = 3.0$). Horizontal dashed lines in *C* and *E* separate the different clusters.

all cases tested (i.e., the clustering in one window did not yield clustering in any other window that did not overlap with the window analyzed).

To determine whether the clusters could have occurred by chance, we ran the same analysis on a trial-shuffled version of all 10 rastergrams used. This shuffling conserved the average spike density of the data (histogram across trials) but eliminated all within-trial temporal correlations (spike patterns). Therefore this procedure assessed the clustering that would occur by chance alone, given the particular firing rate of each cell. For the cell in Figure 6, the number of valid cluster configurations [(number of event, number of clusters) pairs] found decreased significantly, from 25 (34.7; $n = 10$) to 4 (9; $n = 10$). Each configuration consisted of a single event (1.1; $n = 10$) that was split into two or three clusters (2.6; $n = 10$). The average cluster length for such valid configurations was 228 msec (136 msec; $n = 10$). These results indicate that ~75% of the spike patterns (clusters) ob-

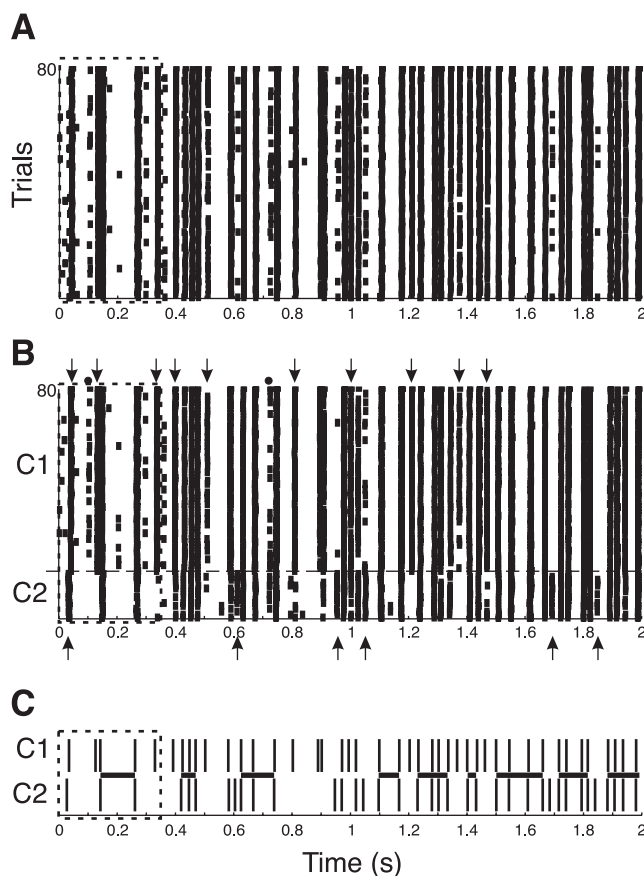


Figure 7. Evidence for spike patterns in cat LGN *in vivo*. *A*, Rastergram of an LGN ON X cell recorded *in vivo* in the anesthetized cat. The receptive field of the cell was presented with 128 repetitions of the same 8 sec long visual stimulus (for clarity, only 80 of 128 trials for the first 2 sec are shown). The clustering algorithm was applied to this data, assuming the presence of two clusters on the first five events. *B*, Reordered rastergram showing two strong clusters. Arrows point to events that were present in one cluster but not the other. *C*, Plot of the reliable events in each cluster (vertical bars). The dashed boxes indicate the region that yielded maximal cluster strength, and the horizontal lines show the regions where the two patterns overlap (within a 5 msec window).

served in the data were caused by temporal correlations rather than by chance alone.

Evidence for long-lasting and overlapping patterns *in vivo*: cat LGN

We used six cells from the Reinagel and Reid study (2002). Figure 7 shows the rastergram of an ON X cell presented with 128 repetitions of the same 8 sec long stimulus [cell 6 in Reinagel and Reid (2002), only the first 80 repetitions and the first 2 sec of the stimulus are shown for clarity]. The rastergram shows spike alignments indicative of reliable events, and the histogram across multiple trials showed peaks that were considered events if they contained at least 40% of the trials. These peaks were fitted by a Gaussian function and a mean width of 2.1 ± 0.9 msec [Reinagel and Reid (2002), their Fig. 2]. The similarity matrix of the spikes was computed with a Gaussian kernel of 3 msec (set according to the histogram peak width measurements mentioned above). The fuzzy clustering algorithm was applied with an initial fuzziness factor of 2, and a desired number of clusters $K = 2$. For this cell, the clustering strength D was simultaneously maximal for the two clusters when only the first 11 events were included in the analysis. In this case, the smaller cluster contained 16 trials (The cluster

strength for the first cluster was $D_1 = 5.86$) and the second, larger cluster contained 112 trials ($D_2 = 8.2$). However, reducing the number of events considered in the analysis did not change the cluster composition until only five events were included ($D_1 = 4.77$ for the 16 trials cluster and $D_2 = 6.95$ for the 112 trials cluster). Including more events did not change the cluster composition (up to 20 events), but reduced cluster strengths.

Figure 7B shows the reordering of the initial histogram according to the clustering performed on the first five events (0–370 msec) (Fig. 7A, dashed box). The histogram across multiple trials was then computed for cluster 1 (C_1) and cluster 2 (C_2) separately (1 msec bins) and peaks were identified. The arrows in Figure 7B indicate the events that were present in one cluster, but not the other. In the first 2 sec, as shown in Figure 7, some spikes that did not belong to any event (defined as above) did show some timing specificity (e.g., dots) and cluster specificity (e.g., left dot). Figure 7C summarizes these two clusters. Vertical marks indicate times (events) when the cell spiked on >40% of the total number of trials in each cluster. The horizontal lines show the regions with common events ($n = 20$ over 8 sec; mean length, 255 msec). The two clusters did not differ in their average firing rates ($C_1 = 27.4 \pm 1.2$ Hz; $C_2 = 26.8 \pm 1.7$ Hz; computed every second).

When the clustering algorithm was run on the same data sets with $K = 3, 4,$ and 5 the smaller cluster obtained with $K = 2$ always appeared without changing the other clusters. The same clustering analysis was run on a control data set obtained with the same cell and consisting of 32 spike trains, each of which was obtained in response to a different stimulus (Reinagel and Reid, 2002). The fuzzy clustering algorithm returned two weak clusters (C_1 , 15 trials, $D_1 = 1.13$; C_2 , 17 trials, $D_2 = 1.3$; data not shown) but no events (i.e., no peaks containing >40% of the trials in each cluster).

These results show that the cell was responding to the same visual stimulus by generating two different (but overlapping) spike patterns that were identifiable from the first 370 msec (five events) of the trials and persisted for several seconds.

The same analysis was run on five other cells recorded from different animals, but stimulated by the same visual stimulus as the cell above (Reinagel and Reid, 2002). Two of these cells showed very similar results as the one presented in Figure 7. One of these [Reinagel and Reid (2002), their Fig. 1B,D, cell 4] had a clustering that was maximal when the first three events were considered (two clusters: $D_1 = 5.95$; $D_2 = 8.31$). The two clusters had significant areas of overlap [$n = 13$ (8 sec), mean length, 469 msec]. The control experiment on 32 different stimuli presented once returned two weak clusters (C_1 , 15 trials, $D_1 = 1.14$; C_2 , 17 trials, $D_2 = 1.24$; data not shown) without any event. A second cell had a clustering that was maximal when the first four events were considered (two clusters: $D_1 = 4.13$, $D_2 = 6.84$). The two clusters also had significant areas of overlap ($n = 20$ over 8 sec; mean length, 253 msec). The control experiment on 32 stimuli presented once returned two weak clusters (C_1 , 14 trials, $D_1 = 1.2$; C_2 , 18 trials, $D_2 = 1.2$; data not shown) without any events. For each of these three cells, the average firing rates within each of the two clusters were not significantly different from each other (ANOVA; $p > 0.2$). The clusters differed only by the precise timing of the events. Note that in these three cells, in spite of common events between the two clusters, the two patterns are significantly different well beyond the first few events that were used in our analysis, as late as 7 sec after stimulus onset (data not shown).

Another cell had two weak clusters, but only when the total length of the recording was considered (8 sec; C_1 , 44 trials, $D_1 = 2.05$; C_2 , 84 trials, $D_2 = 3.48$). Unlike the three cells discussed above, the two clusters obtained for this cell differed significantly in their average firing rates (C_1 , 4.5 ± 1.2 Hz; C_2 , 10.6 ± 2.7 Hz) and only very few and shorter overlaps could be noted [$n = 6$ (8 sec); mean length, 173 msec]. Two other cells presented clear events across all trials, but no clusters with a strength $D > 2$ were found.

Discussion

The stable spike patterns lasting several seconds observed *in vivo* in the cat LGN neurons are unlikely to be caused solely by their intrinsic membrane properties, because long stable overlapping patterns were not observed *in vitro* (Figs. 2, 5). Intrinsic currents such as Ca-activated K currents, or slow afterdepolarization-potential currents (Caeser et al., 1993; Magee and Carruth, 1999) that were responsible for the “history” after a spike lasted ~100 msec *in vitro*, and were effectively “reset” (or “refreshed”) by the spike. If this reset occurred at the same time in two patterns, the two trajectories would converge and become identical. However, neurons *in vivo* are spontaneously active and in a high conductance state (Paré et al., 1998; Destexhe and Paré, 1999; Destexhe et al., 2003), and their adaptation currents are significantly reduced by neuromodulators such as acetylcholine or norepinephrine (Hasselmo, 1995). Under these conditions the history after a stimulus-driven spike should be reduced to a few tens of milliseconds. In contrast, the spike patterns observed in area MT of the macaque lasted for several hundreds of milliseconds, making it unlikely that the spike timing reliability observed in MT depended exclusively on the intrinsic properties of the individual cells. A more likely explanation is that this precision of spike timing is an emergent property of the network with which the cell interacts. This explanation is further supported by our analysis of the LGN data. In this preparation, spike patterns lasted several seconds, in spite of long sequences of common events that *in vitro* would undoubtedly reset the cell. In the LGN, *in vivo* spike patterns are therefore likely to be the result of network dynamics that involve retinal ganglion cells inputs, intrinsic thalamic connections, cortical feedback and possibly neuromodulatory inputs from the brain stem as well as intrinsic neural properties.

The high reliability of neuronal firing *in vivo* is puzzling because of the notorious unreliability of synaptic transmission at individual synapses (Allen and Stevens, 1994; Zador, 1998; Dobrunz and Stevens, 1999). Even if the same presynaptic spike train arrived at the dendritic arbor of a given neuron, the current seen by the action potential generating zone is expected to be different from trial to trial, because of synaptic transmission failures and complex synaptic short-term dynamics (Maass and Zador, 1999). In addition, neurons *in vivo* are constantly bombarded by background synaptic inputs that are likely to interfere with the ability of the neurons to detect specific presynaptic signals (Paré et al., 1998; Destexhe and Paré, 1999). These sources of synaptic unreliability and synaptic noise need to be overcome to yield the millisecond spiking precision observed in LGN (Reinagel and Reid, 2000, 2002), MT (Buracas et al., 1998), and other structures. The means by which such reliability can be achieved are still unclear, but may involve bursting (Lisman, 1997; Reinagel et al., 1999) and temporally or spatially correlated inputs (Salinas and Sejnowski, 2001). The stimuli used in the two studies used here were changing on fast time scales of order 10–50 msec

for full-field flicker stimuli in the cat LGN (Reinagel and Reid, 2002) and 30–300 msec for moving Gabor patches in the monkey area MT (Buracas et al., 1998). It may be argued that these time scales are significantly faster than would be expected in natural viewing conditions, and that similar clusters might not appear for “natural” stimuli. Eye movements, however, induce fast-changing retinal visual stimuli in otherwise static visual scenes on the time scale of those used experimentally in MT (Buracas et al., 1998). Furthermore, when presented with slower luminance changes closer to natural conditions, two LGN neurons showed reliabilities across trials that were qualitatively similar to those obtained with faster-changing stimuli (Reinagel and Reid, 2002), although their responses were less precise. More experimental recordings are needed to explore this issue. Another possibility is that the different spike patterns are influenced by intrinsic physiological processes such as the respiratory cycle or sleep rhythms. Although this was not the case for the data analyzed here, this possibility should be considered when analyzing other data sets.

Does the fact that individual spike times and spike patterns are reproducible across multiple trials mean that the neuron actually uses a temporal encoding scheme? For rapidly varying signals such as the ones used here, Theunissen and Miller (1995) suggest that “a temporal encoding scheme is one in which there is significant additional correlation between a frequency component of a dynamic stimulus signal and a higher-frequency component of the corresponding elicited spike pattern.” When a model neuron is driven by a rapidly varying periodic stimulus waveform and the noise strength is sufficiently small, it will become entrained and will produce a reliable and precise spike pattern. The occurrence of multiple spike patterns across trials will depend on the noise strength (Tiesinga et al., 2002). Spike patterns will indeed add higher-frequency components to spike trains compared with the stimulus frequency. However, we did not perform a systematic study to determine whether these higher-frequency components were correlated with a stimulus parameter. Hence, although spike patterns allow for the possibility of temporal encoding, more experiments are required to determine whether the nervous system uses temporal encoding as defined by Theunissen and Miller (1995).

Our study showed that several different spike patterns (“messages”) might be generated in response to the same stimulus. Which pattern appears at any particular time may depend partly on the history of spiking before stimulus onset, and on the neuromodulatory state of the network. This result has at least two possible consequences. First, because the input was the same, the meaning of the different spike patterns could also be the same; that is, they form synonymous messages and could be interpreted (“decoded”) similarly by downstream neurons. Second, downstream neurons would respond differently depending on which pattern was generated. By biasing which pattern the neuron produces, through neuromodulators or other means, sensory information flow can be modulated. In that case the patterns are not synonymous, but inherently carry additional contextual information that may be used to “route” information flow. It remains for future investigations to determine which is valid under which conditions.

Unsupervised learning techniques have been applied to neural data for the purpose of spike sorting (for review, see Lewicki, 1998), but have not been used previously for discovering spike patterns in neural spike trains across multiple presentations of the same stimulus. Recently, J. D. Hunter, J. Wu, and J. G. Milton

(personal communication) have applied an unsupervised learning technique based on adaptive resonance theory (Carpenter et al., 1991) to order neurons in an ensemble according to the input that each neuron received. This method could also be useful for discovering spike patterns.

The clustering method we propose here is applicable to cases in which the neuronal responses are time-locked to the onset of an event (a stimulus for example). The knowledge of a “time 0” allows for an easy computation of the similarity between spike trains. If such a time is not available, possible time shifts between different recordings have to be considered (Nawrot et al., 2003). The problem then becomes one of detecting repeating spike patterns of given length (or given number of spikes) in running bouts of recordings (Frostig et al., 1990; Lestienne and Tuckwell, 1998; Prut et al., 1998; Oram et al., 1999; Abeles and Gat, 2001; Mao et al., 2001). Note that although our technique does not assume stationarity in spike pattern temporal structure (i.e., absence of temporal modulation), the performance of the algorithm is expected to decrease if the analysis is performed on spike trains that are obtained from neurons that jump from one spike pattern to another (Fig. 5). In such cases, a windowing procedure can be used and the clustering algorithm should be systematically applied to fixed or variable-width windows that are moved along the records, from one subset of events to the next (as in Fig. 6). We also note that this method can be used on simultaneously recorded data from multiple neurons. In that case, all neurons share the same time reference, and the similarity between them (instead of between trials) can be compared. The analyses of such data could go beyond the detection of coincident occurrences of spiking (Riehle et al., 1997; Grun et al., 1999), and units sharing common spike patterns could then be identified.

Appendix

The fuzzy K -means algorithm used in this study has been derived from a standard K -means algorithm that operated on a reshaped similarity matrix. We present in this appendix the steps taken to derive this algorithm, show how to determine the various parameters of the algorithms and show that the fuzzy cluster algorithm yields the most reliable and best clustering across the clustering methods tested here. Other clustering methods, such as projective clustering or hierarchical clustering were not tested. The surrogate data sets are available from <http://www.cnl.salk.edu/~fellous/data/JN2004data/data.html>, to facilitate the comparison with other clustering techniques.

Basic K -means algorithm

Under the assumption that there may be K clusters hidden in the data set, the standard K -means clustering algorithm (Ball and Hall, 1967) identifies cluster centers (patterns) and determines for each trial its distance (cluster membership) to each of the centers. Briefly, the algorithm includes the following steps (details can be found in Everitt et al., 2001): (1) Randomly choose K centers in the N -dimensional space described above (uniform distribution). Randomly assign all points (trials) to one of K equally sized clusters. (2) Compute the distances of all data points (columns of S) to the K centers. (3) Assign each point to the closest cluster. (4) Let the new K cluster centers be the center of mass (hence, K -means) of all the data points that were assigned to them.

Repeat steps 2–4 until no data points are reassigned.
This algorithm is referred to as the basic K -means algorithm.

Extended K -means algorithm

Input reshaping

The distribution of similarities (see Materials and Methods) between trials was often approximately Gaussian, centered at values significantly smaller than 0.5 (Fig. 1*Cb*). The dynamic range of similarity values was small, and most points were confined to a small region of the similarity space. To augment the dynamic range, we reshaped the similarity matrix by applying a sigmoid function to all similarity values.

$$b_{ij} = \frac{1}{1 + e^{-\frac{(s_{ij} - s_{in})}{\tau}}}$$

The sigmoid was centered at the mean of the similarity values, and its slope τ was determined empirically by systematic variation between 0.01 and 0.3 in steps of 0.005. For each slope value, the histogram of similarity values was computed (50 bins, between 0 and 1). The slope value yielding the minimum SD among all bins (i.e., the “flattest histogram”) was chosen (Fig. 1*C*). The search for the flattest histogram stopped if the values in the smallest bin vanished: the slope was then too shallow and no dissimilar values existed. Values of chosen τ typically ranged between 0.05 and 0.15 (see below). This type of reshaping is the goal of Infomax Independent Component Analysis, which maximizes the entropy of the outputs $\{b_{ij}\}$ (Bell and Sejnowski, 1995).

Initial conditions selection

The clustering obtained by the basic K -means algorithms may depend on the choice of the initial conditions (step 1, above) and may converge to a local minimum of the objective function O (defined below). To study the sensitivity of the basic K -means algorithm to the choice of initial conditions, we repeated the clustering 150 times for different initial conditions, and for each cluster computed the cluster strength D_k , defined as the ratio between the average of the intercluster distances of the k th cluster to the average of the intracluster distances. An overall measure of the clustering was obtained as D , the average of all D_k (see below).

The histogram of the distribution of all 150 D s was then plotted (50 bins) and one of the clustering results belonging to the highest bin was randomly selected. Note that this choice is not guaranteed to yield the best possible clustering. Instead, choosing a member of the highest bin is equivalent to choosing a clustering scheme that is reached from most initial conditions (i.e., forming the largest basin of attraction).

The basic K -means algorithm with reshaping and initial condition selection is referred to below as the extended K -means algorithm.

Fuzzy K -means algorithm

In the standard K -means algorithm, data points are assigned to one and only one of the clusters. This may lead to a large number of local minima when only a few data points (trials) are available, hence, a far-from-optimal clustering. A variant of the standard clustering method was introduced that relaxed the strict membership condition and replaced it by probability-like membership. The fuzzy K -means clustering algorithm (FKM) (Bezdek, 1981; Hoppner et al., 1999; Dumitrescu et al., 2000) assigns to a

data point a probability of belonging to each one of the K clusters, instead of a unique membership label.

FKM was based on the minimization of the objective function:

$$O_f(U, C) = \sum_{i=1}^N \sum_{j=1}^K u_{ij}^f \|\tilde{p}_i - \tilde{c}_j\|^2$$

where U is a fuzzy partition of the N trials \tilde{p}_i , having K cluster centers $\tilde{c}_j \in \mathcal{R}^N$, $C = \{\tilde{c}_j\}$. The numbers u_{ij} are the degree of membership of trial i to cluster j (the fuzzy partition) and f is the fuzziness factor such that $f > 1$. The fuzzy partition and the cluster centers were computed iteratively starting from an initial-




Figure 8. Choice of the parameters of the clustering algorithms. *A*, Sensitivity to initial conditions. *a–c* show the performance results for the K -means, extended K -means, and fuzzy K -means algorithms, respectively, on three surrogate rastergrams (diamonds, triangles, and circles). The small symbols show the performance (percentage correct) of 150 individual runs and are plotted against average cluster strength D (see Materials and Methods). The larger symbols in *b* and *c* are the runs that were chosen by the extended algorithms (region of densest D). The insets in *b* and *c* show this choice for 10 surrogate rastergrams based on the same statistics and include the three examples shown in the expanded graphs. *d* shows the average performance error for the three algorithms. The difference between extended K -means (EK) and K -means (K) was statistically significant ($p < 0.01$). The fuzzy K -means (FK) algorithm was the least sensitive to the choice of initial conditions. *B*, Choice of initial fuzziness factor. *a* (circles) shows the average performance of the fuzzy K -means algorithms for different values of the initial fuzziness factor f , computed on the data set used in *A*. Triangle and star symbols show the average performance for the K -means and extended K -means algorithms, respectively. *b* (circles) shows the mean difference in performance relative to $f = 2$. *C*, Actual fuzziness factor used as a function of the amount of jitter (*a*) and extra spikes (*b*) present in the surrogate data set (initial fuzziness factor was three in all cases). *c* shows the distribution of all actual fuzziness factors (bins = 0.05). For clarity, the peak at 3 was truncated.

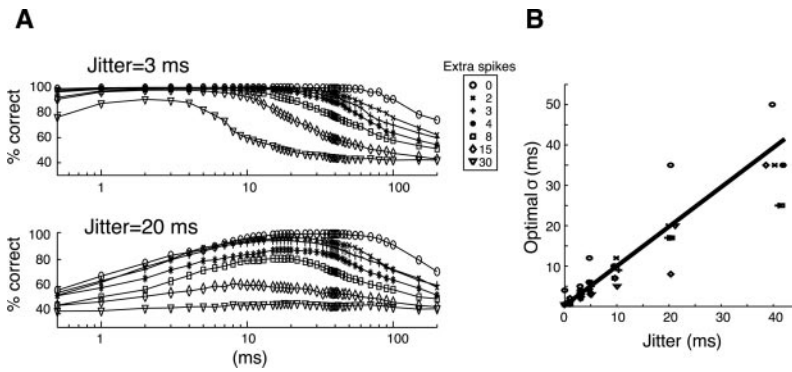


Figure 9. Determination of the optimal σ for the fuzzy K -means clustering algorithm. *A*, Clustering was performed on the three cluster surrogate data set processed with different values of σ . The performance of the clustering is plotted as a function of σ for two values of the jitter. Each curve represents a different number of extra spikes per trials. All curves show the maximum performance for a preferred σ . *B*, Preferred σ as a function of the amount of jitter in the rastergram. Each point is an average of 20 rastergrams. The distribution was fitted by a line of slope 1.0.

random partition. At each iteration, degree of membership and cluster centers were recomputed according to:

$$u_{ij} = \frac{1}{\sum_{k=1}^K \left(\frac{d_{ij}}{d_{ik}} \right)^{\frac{2}{f-1}}}$$

and

$$\tilde{c}_j = \frac{\sum_{i=1}^N u_{ij}^f \tilde{p}_i}{\sum_{i=1}^N u_{ij}^f}$$

with

$$d_{ij} = \|\tilde{p}_i - \tilde{c}_j\|$$

The iteration stopped when the maximal difference between two consecutive u_{ij} was smaller than a termination criterion ϵ set in advance (typically 10^{-12}). Subsequently, we computed the Euclidian distance between all cluster centers. If two cluster centers were identical (distance, $<10^{-6}$), the algorithm was repeated with a smaller fuzziness factor (decremented linearly by steps of 0.05). Note that in the limit, as f becomes 1, FKM becomes equivalent to the extended K -means algorithm described above (Dumitrescu et al., 2000).

Cluster strength

Typical clustering algorithms often converge to clustering solutions, even if no clusters are actually present in the data. To assess the quality of a cluster we computed the clustering strength D :

$$D = \frac{\sum_k D_k}{K}$$

with

$$D_k = \frac{\sum_{\tilde{p} \in Q_k} \|\tilde{p} - \tilde{c}_k\|}{\sum_{\tilde{p} \in Q_k} \|\tilde{p} - \tilde{c}_k\|} \times \frac{N_k}{N - N_k}$$

assuming K clusters Q_j ($j = 1 \dots K$) with centers C_j containing N_j points. D is a measure of the quality of clustering result: a high D value indicates a strong clustering with small within-cluster distances and large intercluster distances.

The fuzzy K-means algorithm is the least sensitive to the initial order of trials

In general it is not possible to predict the sensitivity of K -means algorithms to the initial order of the trials. Such an assessment needs to be done empirically. We compared each of the three K -means algorithms on the surrogate data set with 150 different initial conditions. Figure 8*A* shows an example for three surrogate rastergrams (symbols: diamond, triangle, and oval) that were built with 15 extra spikes, 15% missing event spikes, 10 msec jitter, four random events, and three clusters.

The smaller symbols represent the clustering performance (fraction of correctly clustered trials) as a function of the clustering strength D for each of the individual runs. The larger symbol in Figure 8, *Ab* and *Ac*, shows the position of the initial condition selected by the extended K -means and fuzzy K -means algorithms, respectively. Note that this choice is not necessarily optimal (diamond), although in most cases it results in good performance (triangle and oval). The fuzzy clustering method yields far less variation in performance when several initial conditions are considered (most individual symbols are hidden behind the larger symbols representing the densest D area). The insets in Figure 8, *Ab* and *Ac*, show the relative performance of the extended K -means and fuzzy K -means on 10 rastergrams with the same statistics. Both algorithms yield qualitatively the same relative performance (e.g., star is always the worst classification), although some differences were observed (e.g., square). Figure 8*Ad* shows the mean performance error for 150 initial conditions for the three algorithms studied here computed on all surrogate data sets for three clusters [920 rastergrams exhibiting 49 combinations of jitter (0–40 msec) and extra spikes (0–30 extra spikes per trials) with a fixed amount of missing spikes (15%)]. The fuzzy K -means algorithm yielded $<0.2\%$ error across all initial conditions tested. These results were similar for two and five clusters (data not shown) and indicate that of the three algorithms, the fuzzy K -means algorithm is the most insensitive to initial conditions. However, for all 150 initial conditions tested, the fuzzy K -means algorithm was significantly slower than the extended K -means algorithm by a factor of 4–7, depending on the number of clusters and their overall similarity (measured on a Pentium-based computer, using a MATLAB implementation). Therefore, the data analysis involves a tradeoff between accuracy (fuzzy K -means) and speed (extended K -means).

Choice of the initial fuzziness factor

The fuzziness factor f in the fuzzy clustering algorithm was adaptively set to avoid convergence on identical cluster centers. From a given starting value, f was linearly decreased (but with $f > 1$) until the K clusters were different. Figure 8*Ba* shows the overall performance of the fuzzy K -means clustering algorithm as a function of the initial fuzziness factor for the three cluster surrogate set. Note that the SD of the clustering performance increased with smaller initial fuzziness factors. For comparison, the performance using the basic and extended K -means algorithm is plot-

ted to the left of the curve (triangle and star, respectively). Maximal average performance and minimal variation in performance was obtained for the fuzzy K -means algorithm and an initial fuzziness factor >1.7 . Figure 8*Bb* shows the average difference of performance when the initial fuzziness factor was varied, compared with the case in which $f = 2$. This average is the smallest for fuzziness factors >1.7 (not accounting for $f = 2$ when it is 0 by construction) with no statistical differences for different values of f . Figure 8*C* shows the fuzziness factor used as a function of jitter (*a*) and amount of extra spikes per trials (*b*). In general, difficult cluster separation (high jitter, or high noise) required smaller fuzziness factors. Figure 8*Cc* shows the aggregate histogram of all the fuzziness factors used across all conditions, with a starting fuzziness factor equal to 3. The lower peak of the distribution occurred at 1.22 (not accounting for the peak at 3 which included 587 samples and which was truncated for clarity). For two and five clusters this secondary peak was located at 1.31 and 1.14, respectively. As the number of clusters increased, the effective f used became closer to 1 and the fuzzy K -means became equivalent to the extended K -means.

The fuzzy K -mean algorithm yielded the highest and least variable overall performance for initial fuzziness factors >1.7 . The clustering of the three-cluster surrogate set was repeated with an initial fuzziness factor equal to 2 instead of 3, with no significant difference in performance.

The clustering performance is optimal for width of the Gaussian kernel (σ) corresponding to the amount of jitter in the data

The temporal precision considered by the analysis depends on the width of the Gaussian kernel used to convolve the spike trains before the computation of the similarity matrix. When the jitter is low, a small σ correctly produces high similarities between trials, irrespective of the density of extra spikes present within the cluster. When the jitter is high, a large (but not a small) σ correctly produces high similarity values provided the density of extra spikes is low. In both cases, setting σ too high confounds trials that belong to different clusters. Choosing a small σ emphasizes the fine temporal structure, whereas a large σ will emphasize the overall number of spikes (firing rate), regardless of their precise temporal pattern.

Figure 9*A* shows examples of clustering performance for low (3 msec) and high (20 msec) jitter, in the case of three clusters (fuzzy K -means clustering; initial $f = 2$ with 15% missing spikes) as a function of the width of the Gaussian kernel. The curves had an inverted U shape, and performance was maximal for a limited range of σ . Note that this range increased as the jitter was decreased. Figure 9*B* plots the values of σ that yielded the best performance as a function of the jitter present in the data for all surrogate sets with three clusters. The relationship between optimal σ and input jitter was well fitted by a line of slope 1.0 and 0 offset (Nelder–Mead fitting method in MATLAB). The outliers arise from the flatness of some curves (e.g., circles) near their peaks. In these cases, the variability caused by jitter and extra spikes make it difficult to determine the position of the maximum of the curve accurately. Similar linear fits were obtained for two and five clusters. These results indicate that σ should be approximately equal to the average amount of jitter in the data. In some cases, this jitter can be estimated as the average peak width in the peristimulus time histogram (“precision” in Mainen and Sejnowski, 1995). However, in cases in which there are a few trials with a large number of clusters, no clear peaks can be detected in the PSTH, and σ should be systematically varied.

Overall test results using the extended K -means algorithm on

the surrogate data set (Fig. 4) showed that the fuzzy K -means algorithm performed, on average, 8.1, 9.1, and 15.5% better across all conditions (jitter and extra spikes) for two, three, and five clusters, respectively. Compared with the basic K -means algorithm, the mean increase in performance achieved by using the fuzzy K -means algorithm was 10.1, 14.3, and 22.5% for two, three, and five clusters, respectively.

References

- Abeles M, Gat I (2001) Detecting precise firing sequences in experimental data. *J Neurosci Methods* 107:141–154.
- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130.
- Allen C, Stevens CF (1994) An evaluation of causes for unreliability of synaptic transmission. *Proc Natl Acad Sci USA* 91:10380–10383.
- Amit DJ, Fusi S, Yakovlev V (1997) Paradigmatic working memory (attractor) cell in IT cortex. *Neural Comput* 9:1071–1092.
- Ball GH, Hall DJ (1967) A clustering technique for summarizing multivariate data. *Behav Sci* 12:153–155.
- Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput* 7:1129–1159.
- Berry MJ, Warland DK, Meister M (1997) The structure and precision of retinal spike trains. *Proc Natl Acad Sci USA* 94:5411–5416.
- Berry II MJ, Meister M (1998) Refractoriness and neural precision. *J Neurosci* 18:2200–2211.
- Bezdek JC (1981) Pattern recognition with fuzzy objective function algorithms. New York: Plenum.
- Bodner M, Zhou YD, Fuster JM (1998) High-frequency transitions in cortical spike trains related to short-term memory. *Neuroscience* 86:1083–1087.
- Bryant HL, Segundo JP (1976) Spike initiation by transmembrane current: a white-noise analysis. *J Physiol (Lond)* 260:279–314.
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969.
- Caeser M, Brown DA, Gahwiler BH, Knopfel T (1993) Characterization of a calcium-dependent current generating a slow afterdepolarization of CA3 pyramidal cells in rat hippocampal slice cultures. *Eur J Neurosci* 5:560–569.
- Carpenter GA, Grossberg S, Rosen DB (1991) Fuzzy art: fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Netw* 4:759–771.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808.
- Destexhe A, Paré D (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol* 81:1531–1547.
- Destexhe A, Rudolph M, Pare D (2003) The high-conductance state of neocortical neurons in vivo. *Nat Rev Neurosci* 4:739–751.
- Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Dobkins KR, Albright TD (1994) What happens if it changes color when it moves? The nature of chromatic input to macaque visual area MT. *J Neurosci* 14:4854–4870.
- Dobrunz LE, Stevens CF (1999) Response of hippocampal synapses to natural stimulation patterns. *Neuron* 22:157–166.
- Dumitrescu D, Lazzareni B, Jain LC (2000) Fuzzy sets and their application to clustering and training. New York: CRC.
- Everitt BS, Landau S, Leese M (2001) Cluster analysis, ed 4. London: Arnold.
- Fellous JM, Houweling AR, Modi RH, Rao RP, Tiesinga PH, Sejnowski TJ (2001) Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *J Neurophysiol* 85:1782–1787.
- Freeman WJ (2000) Mesoscopic neurodynamics: from neuron to brain. *J Physiol (Paris)* 94:303–322.
- Frostig RD, Frostig Z, Harper RM (1990) Recurring discharge patterns in multiple spike trains. I. Detection. *Biol Cybern* 62:487–493.
- Grun S, Diesmann M, Grammont F, Riehle A, Aertsen A (1999) Detecting unitary events without discretization of time. *J Neurosci Methods* 94:67–79.
- Gur M, Beylin A, Snodderly DM (1997) Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J Neurosci* 17:2914–2920.

- Hartigan JA (1975) Cluster algorithms. New York: Wiley.
- Hasselmo ME (1995) Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behav Brain Res* 67:1–27.
- Hoppner F, Kruse R, Klawonn F, Runkler T (1999) Fuzzy cluster analysis: methods for classification, data analysis and image recognition. New York: Wiley.
- Hunter JD, Milton JG (2003) Amplitude and frequency dependence of spike timing: implications for dynamic regulation. *J Neurophysiol* 90:387–394.
- Hunter JD, Milton JG, Thomas PJ, Cowan JD (1998) Resonance effect for neural spike time reliability. *J Neurophysiol* 80:1427–1438.
- Jain AK (1988) Algorithms for clustering data. Upper Saddle River, NJ: Prentice Hall.
- Kamondi A, Acsady L, Wang XJ, Buzsaki G (1998) Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. *Hippocampus* 8:244–261.
- Kara P, Reinagel P, Reid RC (2000) Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27:635–646.
- Lestienne R, Tuckwell HC (1998) The significance of precisely replicating patterns in mammalian CNS spike trains. *Neuroscience* 82:315–336.
- Lewicki MS (1998) A review of methods for spike sorting: the detection and classification of neural action potentials. *Netw Comput Neural Syst* 9:R53–R78.
- Lisman JE (1997) Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci* 20:38–43.
- Maass W, Zador AM (1999) Dynamic stochastic synapses as computational units. *Neural Comput* 11:903–917.
- Magee JC, Carruth M (1999) Dendritic voltage-gated ion channels regulate the action potential firing mode of hippocampal CA1 pyramidal neurons. *J Neurophysiol* 82:1895–1901.
- Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506.
- Manwani A, Koch C (1999) Detecting and estimating signals in noisy cable structure. I. Neuronal noise sources. *Neural Comput* 11:1797–1829.
- Mao BQ, Hamzei-Sichani F, Aronov D, Froemke RC, Yuste R (2001) Dynamics of spontaneous activity in neocortical slices. *Neuron* 32:883–898.
- Nawrot MP, Aertsen A, Rotter S (2003) Elimination of response latency variability in neuronal spike trains. *Biol Cybern* 88:321–334.
- Oram MW, Wiener MC, Lestienne R, Richmond BJ (1999) Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol* 81:3021–3033.
- Paré D, Shink E, Gaudreau H, Destexhe A, Lang EJ (1998) Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *J Neurophysiol* 79:1450–1460.
- Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, Abeles M (1998) Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J Neurophysiol* 79:2857–2874.
- Rabinovich MI, Abarbanel HD (1998) The role of chaos in neural systems. *Neuroscience* 87:5–14.
- Reich DS, Victor JD, Knight BW (1998) The power ratio and the interval map: spiking models and extracellular recordings. *J Neurosci* 18:10090–10104.
- Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. *J Neurosci* 20:5392–5400.
- Reinagel P, Reid RC (2002) Precise firing events are conserved across neurons. *J Neurosci* 22:6837–6841.
- Reinagel P, Godwin D, Sherman SM, Koch C (1999) Encoding of visual information by LGN bursts. *J Neurophysiol* 81:2558–2569.
- Riehle A, Grun S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–1953.
- Salinas E, Sejnowski TJ (2001) Correlated neuronal activity and the flow of neural information. *Nat Rev Neurosci* 2:539–550.
- Schoner G, Kelso JA (1988) Dynamic pattern generation in behavioral and neural systems. *Science* 239:1513–1520.
- Schreiber S, Fellous J-M, Tiesinga PH, Sejnowski TJ (2003) A new correlation-based measure of spike timing reliability. *Neurocomputing* 52–54:925–931.
- Schreiber S, Fellous JM, Tiesinga P, Sejnowski TJ (2004) Influence of ionic conductances on spike timing reliability of cortical neurons for supra-threshold rhythmic inputs. *J Neurophysiol* 91:194–205.
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16:752–768.
- Theodoridis S, Koutroumbas K (1998) Pattern recognition. San Diego: Academic.
- Theunissen F, Miller JP (1995) Temporal encoding in nervous systems: a rigorous definition. *J Comput Neurosci* 2:149–162.
- Tiesinga PH (2002) Precision and reliability of periodically and quasiperiodically driven integrate-and-fire neurons. *Phys Rev E Stat Nonlin Soft Matter Phys* 65:041913.
- Tiesinga PH, Fellous JM, Sejnowski TJ (2002) Attractor reliability reveals deterministic structure in neuronal spike trains. *Neural Comput* 14:1629–1650.
- Tsuda I (2001) Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behav Brain Sci* 24:793–848.
- Victor JD, Purpura KP (1996) Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol* 76:1310–1326.
- Warzecha AK, Kretzberg J, Egelhaaf M (2000) Reliability of a fly motion-sensitive neuron depends on stimulus parameters. *J Neurosci* 20:8886–8896.
- Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384:162–166.
- White JA, Rubinstein JT, Kay AR (2000) Channel noise in neurons. *Trends Neurosci* 23:131–137.
- Zador A (1998) Impact of synaptic unreliability on the information transmitted by spiking neurons. *J Neurophysiol* 79:1219–1229.
- Zipser D, Kehoe B, Littlewort G, Fuster J (1993) A spiking network model of short-term active memory. *J Neurosci* 13:3406–3420.