Neuron
Article

Stochastic Emergence of Repeating Cortical Motifs in Spontaneous Membrane Potential Fluctuations In Vivo

Alik Mokeichev,[1,2] Michael Okun,[1,3] Omri Barak,[1] Yonatan Katz,[1] Ohad Ben-Shahar,[2] and Ilan Lampl[1,*]
[1] Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel
[2] Department of Computer Science and The Zlotowski Center for Neuroscience, Ben Gurion University of the Negev, Beer-Sheva 84105, Israel
[3] Department of Computer Science, The Hebrew University of Jerusalem, Jerusalem 91904, Israel
*Correspondence: ilan.lampl@weizmann.ac.il
DOI 10.1016/j.neuron.2007.01.017

SUMMARY

It was recently discovered that subthreshold membrane potential fluctuations of cortical neurons can precisely repeat during spontaneous activity, seconds to minutes apart, both in brain slices and in anesthetized animals. These repeats, also called *cortical motifs*, were suggested to reflect a replay of sequential neuronal firing patterns. We searched for motifs in spontaneous activity, recorded from the rat barrel cortex and from the cat striate cortex of anesthetized animals, and found numerous repeating patterns of high similarity and repetition rates. To test their significance, various statistics were compared between physiological data and three different types of stochastic surrogate data that preserve dynamical characteristics of the recorded data. We found no evidence for the existence of deterministically generated cortical motifs. Rather, the stochastic properties of cortical motifs suggest that they appear by chance, as a result of the constraints imposed by the coarse dynamics of subthreshold ongoing activity.

INTRODUCTION

A long-lasting debate in neural coding revolves around the role of precise timing of action potentials. Different studies have shown that information in the central nervous system is well represented by the average firing rate of the neurons (Mazurek and Shadlen, 2002; Richmond and Optican, 1990; Richmond et al., 1987; Shadlen and Newsome, 1998). Precise spike times, in the order of a millisecond, may carry additional information (Dan et al., 1998; Mainen and Sejnowski, 1995; Victor and Purpura, 1998) that was suggested to play an important role in various perceptual processes in the cortex. Precise temporal firing in the form of coherent oscillatory activity among different neurons may link different features of a single object (Eckhorn, 1994; Engel et al., 1992; Gray et al., 1992), and precise oscillations were also suggested as a key feature of place cells (Burgess and O'Keefe, 1996; Huxter et al., 2003).

In the above studies, the precision of firing was suggested to be manifested in the form of brief spatiotemporal correlations and synchrony that emerge in the cortex and are strongly related to stimulation and behavior. Whether cortical mechanisms may control the precision of firing over a relatively long duration (up to several hundreds of milliseconds) and create complex spatiotemporal patterns of spike appearance is controversial. Clearly, a network-controlled timing of nonoscillatory spike patterns can encode information more efficiently than rate, count, or oscillations. Precise temporal patterns may reflect functional connectivity or task-dependent cohorts of cooperative neurons (Abeles et al., 1993). Precisely repeated temporal patterns of spikes, distributed across multiple neurons, have been considered as an indication of the existence of such functionally active groups of neurons (Dayhoff and Gerstein, 1983a, 1983b; Lestienne and Strehler, 1987; Prut et al., 1998). These claims are highly sensitive to the underlying statistical assumptions, and it was suggested that repeats of these types of patterns occur by chance (Baker and Lemon, 2000; Oram et al., 1999), therefore not necessarily implying the existence of controlled mechanisms for generation of repeating spike patterns (Richmond et al., 1999).

Intracellular recordings provide an additional method for searching for patterns of cortical firing. A single cortical neuron receives inputs from thousands of other neurons, so that precise repeating activation patterns of these neurons may lead to repeating patterns in the subthreshold fluctuations of the postsynaptic neuron, mediated by postsynaptic potentials. Repeated epochs of spontaneous synaptic potentials ("motifs") were indeed reported in several recent studies of cortical slices using intracellular and imaging techniques (Cossart et al., 2003; Ikegaya et al., 2004; MacLean et al., 2005; Mao et al., 2001), and in in vivo intracellular recordings (Ikegaya et al., 2004). (By way of full disclosure, I.L. is an author on Ikegaya et al. [2004].) A single motif may include several noticeable

waves of synchronized synaptic potentials separated by quiet hyperpolarized periods. Some motifs show extraordinary degree of precision in the reappearing activity, accurate down to a few milliseconds. The duration of a motif is on the order of one second, with a time span of seconds to minutes between its successive appearances (Ikegaya et al., 2004). Repeats which occur beyond the chance level provide a compelling evidence for the existence of mechanisms that enable precise regeneration of specific long-lasting patterns. Indeed, such a conclusion was argued by Ikegaya et al. (2004), who tested whether motifs occur more often than predicted by chance by applying various tests (see the Discussion below).

In this study, we reexamine the appearance of motifs in spontaneous membrane potential fluctuations of neurons in the rat barrel cortex. Using a search algorithm similar to the one described in Ikegaya et al. (2004), a large number of motifs were found in long continuous recordings, in some cases with a higher similarity than previously reported. We have applied three different methods to generate surrogate data, each corresponding to a different model of randomness. In the first method, the surrogate data were constructed by a time domain shuffling of the original traces. In the second method, the data were randomized in the frequency domain. In the third method, we optimized the parameters of a simulated passive cell receiving random synaptic inputs in order to elicit an activity with dynamics similar to that of the recorded cell. Motifs were found in all types of surrogate data. Moreover, their numbers, similarity scores, and basic emergence statistics closely resembled the recorded data. Based on these findings, we suggest that repeated patterns of synaptic activity appear at random as a result of the coarse dynamical properties of the subthreshold fluctuations in cortical cells, rather than by controlled network mechanisms.

RESULTS

Continuous intracellular recordings of spontaneous activity of 10–20 min in duration from five cells of the barrel cortex in anesthetized rats were included in the analysis. Each cell was recorded from a different animal, from the granular or the supragranular layers. Spontaneous firing rate was low, on average a cell fired 0.18 spikes/s. Examples of spontaneous subthreshold activity recorded from these five cells (Figure 1A) demonstrate different dynamics, which is also evident from power spectrum density functions and voltage histograms (see Figures 1B and 1C). The membrane potential distribution covers a large range of potentials and in most of the cells shows some bimodal tendency (see Figure 1B), cf. (Anderson et al., 2000; Lampl et al., 1999; Petersen et al., 2003; Sachdev et al., 2004). The recordings had a stable membrane potential activity for the entire period, evident from stable voltage distributions and power spectra along the recording (data not shown).

Motifs Appear in Recordings of Spontaneous Activity of Cortical Neurons

The search for motifs in spontaneous membrane potential fluctuations of cortical neurons was performed by employing a method similar to the one previously reported in Ikegaya et al. (2004). As in this previous study, we avoided analyzing cells with clear and robust oscillatory activity, when inspected by eye and by power spectral analysis. Many repeating pairs of a high temporal precision and of similar amplitude of voltage deflections were found in all cells. Motif detection was based on the high-resolution index (HRI) scores (see Experimental Procedures and the Supplemental Data available with this article online), which were successful in capturing remarkable examples of repeating activity. For example, despite being recorded 27 s apart, the two 1 s segments presented in Figure 1D (upper left pair of traces) show that every swing of membrane potential in the first segment is accompanied by the same pattern in the repeat. These selected examples represent only a small fraction of motifs found in each cell. The two examples of motifs in the right panels of Figure 1D were ranked as the 1000th most similar pair, yet they show a marked similarity.

Numerous motifs were also found when we searched for motifs of 2 s duration. The number of 2 s motifs was 25%–30% of the number of 1 s motifs. To compare the distributions of HRI scores of motifs of different durations, we sorted each set of scores and plotted it, normalized at the abscissa by its size. The normalized distribution of 1 s motifs is similar in its shape to the distribution of 2 s motifs and these two distributions show substantial difference among the cells (Figure 2). Since this suggests that the normalized distribution of motifs HRI scores is insensitive to motif duration, we limited our analysis to motifs of 1 s in duration, a choice which also follows the average motif length found in Ikegaya et al. (2004).

A possible source for repeating motifs could arise from oscillatory activity in the recorded data. Such activity may increase the tendency of epochs to repeat, provided that the oscillations are accurate and reproducible during different periods of ongoing activity. To find if oscillatory activity and HRI scores are correlated, we measured the oscillatory index (OSI), defined as the amplitude of the first satellite peak in the normalized autocorrelation of each of the 1 s motifs. A higher OSI index corresponds to a more oscillatory subthreshold activity. When motifs were ranked by their OSI, we found that in many cases the top 1000 oscillatory epochs demonstrated a peak between 3 to 20 Hz in their power spectrum. However, the correlation between HRI values and their corresponding OSI values was almost negligible, demonstrating correlation coefficients between −0.004 and 0.104 (n = 5 cells, the correlation was negative in one cell). In addition, for each recorded cell, we compared the OSI scores of the top 50 HRI-ranked motifs to those of the next 10,000 top-ranked motifs. A statistically significant difference was found only in one cell (cell 3, t test, $p < 0.01$). Even for this cell, the mean OSI of top 50 HRI-ranked motifs (mean

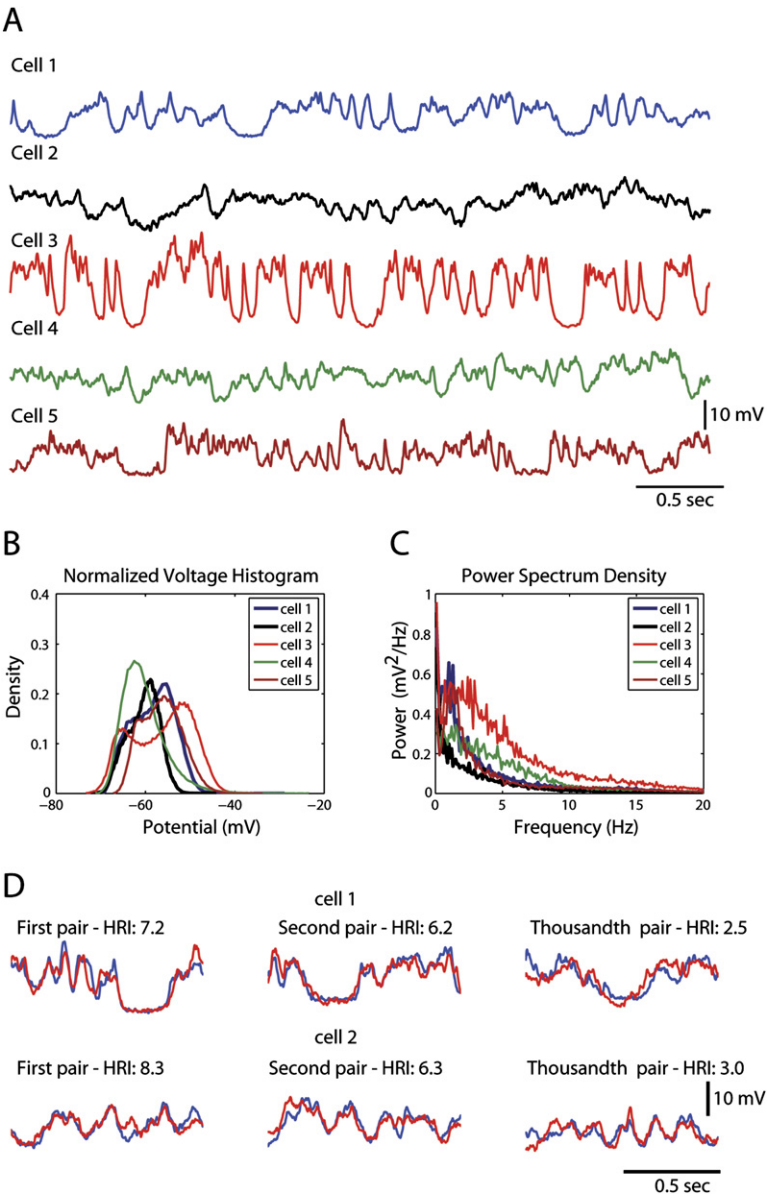

Figure 1. Emergence of Repeating Epochs of Spontaneous Subthreshold Activity

(A) Examples of spontaneous activity recorded from five cells. Note that rapid synaptic activity is observed throughout the entire 4 s epochs.

(B) Membrane potential distributions.

(C) Power spectrum of the recorded cells, obtained by averaging the power spectrum of a 10 s sliding window (without overlapping).

(D) Examples of repeating motifs from two different cells. Each row presents examples from the same cell, first the two motifs with the highest HRI values, then a motif which has the 1000th highest HRI rank (still showing a marked similarity). Such highly similar motifs were found in all the recorded cells in spite of the differences between cells in their statistical properties.

OSI = 0.27) was higher than the mean OSI of the next 10,000 rank ordered motifs only by 12%, while the mean OSI of the top 50 most oscillatory motifs was 0.54. The low correlation between HRI and OSI suggests that the oscillatory activity in the data rarely repeats with a precision that is high enough to result in high HRI. Therefore, the entire data were analyzed irrespectively of OSI scores.

Types of Surrogate Data

Our hypothesis is that motif statistics in spontaneous activity are determined by the stochastic properties of subthreshold voltage dynamics. One way to test this hypothesis is to compare the statistics of motifs in the recording to that of *random* surrogate signal which has dynamics similar to the recorded data. To this end, the most straightforward approach to produce surrogate data is by shuffling the recorded trace in the time domain. Our time-domain-shuffling method ("Interval shuffling," Figure 3A) closely preserves the voltage distribution and higher-order statistics of the recorded data. The resulting power spectrum is also very similar (see Figure S1). The shuffled surrogate data was obtained by fragmenting the recording into short segments and assembling them in a random order (see Experimental Procedures and the Supplemental Data). An example of a motif found in the time-domain-shuffled data is depicted in the right panel of Figure 4A. For comparison, a motif that was detected in the recorded data is shown in the left panel. Note that the dynamics of the surrogate data, as we mentioned above, seem to be unaffected by the shuffling. Since every

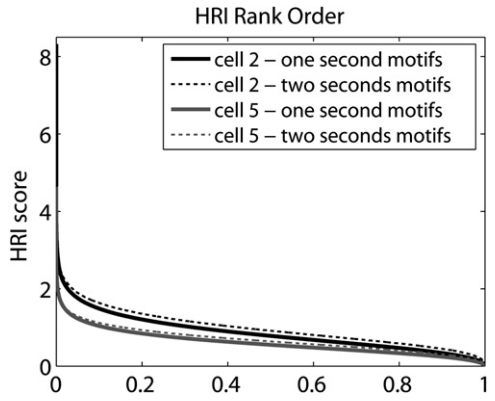

Figure 2. HRI Scores of Motifs of 1 s and 2 s Duration, Sorted from Highest to Lowest, for Two Different Cells

Abscissa is normalized by the total number of motifs. In each cell, distributions of motifs of 1 s and 2 s durations show remarkable similarity when compared to differences across the cells.

second of the shuffled data was randomly composed from at least three segments, each shorter than 500 ms, the depicted 1 s motif must have repeated by chance.

A different type of surrogate data was obtained by using the phase randomization method (see Figure 3B and Experimental Procedures) to generate a random signal whose power spectrum is identical to that of the original data. Phase randomization is equivalent to filtering a white noise signal with a filter whose modulation transfer function (MTF) is identical to the physiological power spectrum, and it is known to produce a voltage trace with normal distribution (Schreiber and Schmitz, 1996), which could differ from the distribution of our recorded data. An important consequence of this randomization method is that any particular sequence of events in the original data, e.g., distinct fast rising synaptic potentials, disappears in the phase randomized trace. Still, we found repeated motifs in this type of surrogate data as well, and one of them is illustrated in Figure 4B (right panel), next to an example from the recorded data (left panel).

Both of the methods described above acted upon the entire recorded data in order to produce a corresponding random surrogate data. We also developed a method to generate randomized data based on a small number of parameters extracted from the recorded data, by using a simulation of a cell that receives Poisson-distributed random synaptic inputs (see Figure 3C and Experimental Procedures). The aim of the simulation was to generate voltage fluctuations similar to those in the physiological data using a relatively simple and comprehensible random process, even if such a simulation does not fully reflect actual presynaptic properties and rates. Indeed, simulated voltage distributions were unimodal (see Supplemental Data, Figures S2A and S2C), as expected from similar stochastic models (Rudolph et al., 2004), whereas most of our recordings had a tendency toward bimodal voltage distribution. Although these simulations fed cells with completely stochastic synaptic inputs, the resultant traces exhibited motifs (Figure 4C), as was also found by the other two methods.

Statistics of Motifs in Physiological and in Surrogate Data Are Similar

The appearance of high-precision motifs in randomized data might be considered surprising. In itself, however, it is insufficient for drawing any conclusions. Hence, we proceeded with a detailed comparison between various statistical properties of motifs found in the physiological and the randomized data. First, we compared the entire distribution of motifs, when they were rank ordered by their HRI scores. The distributions of the HRI scores of all the motifs found in the physiological data and in the corresponding surrogate data are presented for two of the recorded cells in Figures 5A and 5B. The data is presented in a semi-log scale to emphasize higher-ranked motifs. The HRI distributions of phase randomization and Poisson simulations were obtained from single randomized surrogate, whereas for the time domain interval-shuffling method the green solid curve was obtained by averaging HRI-ranked distributions of 40 independent shuffles of the original recorded trace. Dashed green curves are the 99% confidence limits of this average distribution (see "Statistical Analysis" in the Experimental Procedures section). For both cells, the distribution of the top recorded motifs is within the confidence limits range of data obtained using time domain shuffling. At lower HRIs, the physiological curve slightly departs from the upper confidence limit of the time-domain-shuffled data.

To compare the number of motif repetitions and their mean HRI scores across cells, rather than just between each physiological data and its corresponding random data, we limited the search for motifs to the duration of the shortest recording, i.e., to the first 10 min of continuous recording from each cell. In some cells, more motifs were found in the physiological data, while in others it was the surrogate data that exhibited more motifs (Figure 5C). However, the differences in the number of motifs were more significant across the cells ($p < 0.04$, two-way ANOVA) than within cells ($p > 0.1$) for different data sets. For example, the number of motifs in the phase randomized data of cell 1 is by 25% higher than the number of motifs found in the phase randomized data of cell 3, yet the number of motifs in these randomized traces is smaller only by 5% than the number of motifs found in the corresponding original data.

To compare the mean HRI scores of motifs found in recorded data and in the three types of surrogate data, for each data set we calculated the mean HRI value of three subpopulations—the mean HRI score of the entire distribution (Figure 5D), the mean HRI of the top 1000 HRI-ranked motifs (Figure 5E), and the mean HRI of the top 10 motifs (Figure 5F). Across cells, the mean HRI score of the recorded data was slightly higher than the value measured from the time domain shuffles (2.0%, 3.6%, and 8.0% difference in Figures 5D–5F, respectively). At the upper range of HRI scores (Figures 5E and F),

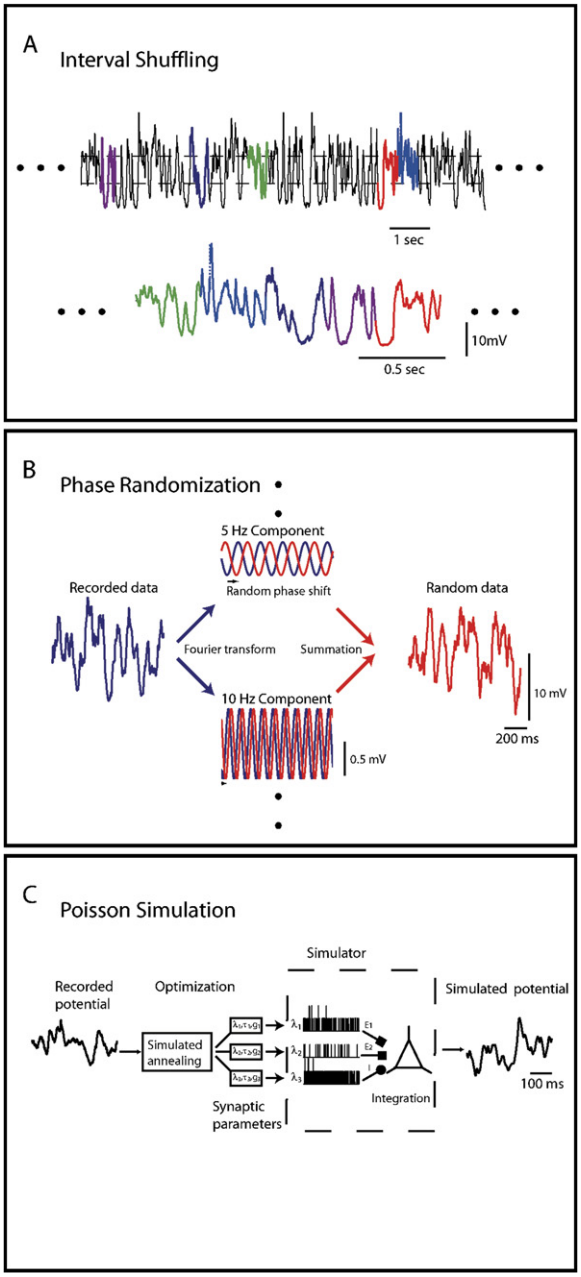

Figure 3. Stochastic Surrogate Data Were Generated by Three Different Methods

(A) Time domain interval shuffling. Using two levels of potentials (determined from 1/3 and 2/3 of the density distribution), the data were fragmented into short segments. Each segment starts and ends at one of the two levels, and its duration was the longest possible below 500 ms. Five different segments are marked by colors. The fragments were then randomly assembled to generate a new continuous voltage trace.

(B) Frequency domain phase randomization. The recorded data were decomposed into its frequency components using FFT. For demonstration, two components, 5 Hz and 10 Hz, are shown (blue). Each component was shifted by a random phase (red), and new random data were constructed by summing up all phase randomized components.

however, confidence limits cover a larger range, and in some cells they indicate that the differences were nonsignificant. Moreover, we found a strong correlation between the physiological HRI scores and those of the time-domain-shuffling surrogate ($r^2 > 0.98$, $p < 0.002$ for each of the three ranges) and also between the physiological HRI scores and simulated surrogate ($r^2 > 0.93$, $p < 0.002$ for each of the three ranges). For the phase-randomized surrogate, the correlation was nonsignificant due to the lower HRI scores obtained for cell 3. We conclude that the three types of surrogate data capture the primary statistical properties of the recorded HRI distributions.

One issue of concern with any method that produces surrogate data from recorded data is its effect on genuine motifs, if they do exist. In the Supplemental Data, we show that even if one epoch (selected randomly) is artificially implanted in random places of a 20 min continuous data, approximately once per minute, it is easily "detected" by comparing the HRI scores distribution of this synthetic data to its shuffles. Thus, the small differences between the distributions of the physiological and the randomized data that were found suggest that the vast majority of motifs are a result of the coarse characteristics of the recorded trace, such as its power spectrum and voltage distribution. Nevertheless, we proceeded with examining additional statistical properties which further support this hypothesis.

In Ikegaya et al. (2004), it was reported that the mean duration between the repeats of a motif can vary substantially, from less than a second to minutes. To understand the source of this variability, we compared the distribution of the time interval between the repeats of each motif in 10 min of recorded data to that in the random data (Figure 6). Panel A of Figure 6 shows that shorter intervals between repeats were more abundant and that the density function declines linearly with the increasing interval length. The same behavior is observed for the top 1000 motifs (inset in Figure 6A), both from the physiological data and from the random data. The observed density function coincides with the distribution of the distance between two points that are randomly and independently selected in a unit interval (each point with a uniform probability). In this case the expectation of the distance is 1/3, which closely matches the mean time interval (normalized by the recording duration) between a motif and its repeat. This result was obtained for all cells and all randomization methods (Figures 6D–6F) and independently of the recording duration (Figures 6B and 6C). This result seems to match the measures of (Ikegaya et al., 2004), after normalizing the reported mean time interval by the reported mean duration of the recordings. These observations indicate that motifs' appearance times are random.

(C) Simulation of a single compartment neuron with Poisson-distributed synaptic inputs. The simulation parameters (average rates of excitatory and inhibitory synaptic inputs, their maximal conductance, and time constants) were obtained using a simulated annealing algorithm by minimizing the differences in power spectra and voltage distribution between the simulated and the recorded data.

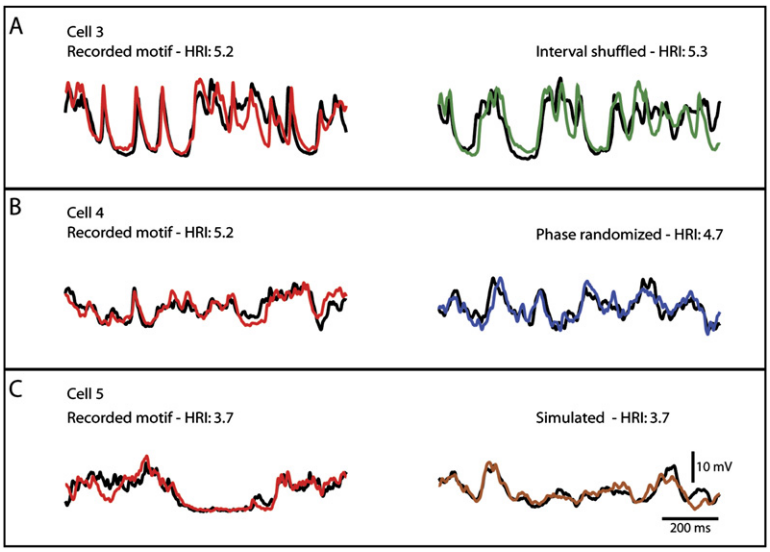

Figure 4. Motifs from the Original Data and the Corresponding Surrogate Data, from Three Different Cells and from Three Different Surrogate Data Types, Chosen from the 50 Top-Ranked Motif Pairs

(A) On the left side, an example of a motif pair found in recorded data, and on the right, a motif found in its interval-shuffled surrogate.

(B) Left side, a motif from a different cell, and on the right, a motif from its phase-randomized surrogate.

(C) An example of a motif from a third cell and on its right a motif from a simulated neuron with random Poisson-distributed inputs.

An additional characteristic that may distinguish between deterministic and stochastic processes of motif generation is the *number* of times a motif is repeated. A network-controlled mechanism of repeating motifs is expected to repeat motifs more times than dictated by chance. To test this possibility, we searched for multiple repeats and compared them across different data sets. In this analysis, we limited our search in each cell to the top 1000 motifs as ranked by their HRI scores. Triplets, quadruplets, quintuplets, and multiple repeats of higher-order appeared in all the cells and all the data sets that were analyzed (Figure 7). The frequency of the number of multiple repeats decreases almost equally across different cells and data sets. Note that the differences between the physiological data and the time-domain-shuffled data are not significant, as is indicated by the confidence limits bars for this surrogate data.

DISCUSSION

Summary of the Findings

In this study of cortical neurons in the anesthetized rat, we find that 1 s epochs of spontaneous subthreshold activity can faithfully repeat seconds to minutes apart. To test the expected occurrence of motifs, we generated randomized data that preserve the dynamics of the recorded data. The statistical properties of the motifs found in the recorded and the surrogate data, such as their number, HRI distribution, the time interval between motif repeats, and the number of motifs with more than two appearances, were very similar, in particular when the most similar repeating patterns are considered.

Any of the randomization methods inevitably altered some of the characteristic of the recorded data. This is particularly true for the phase randomization and simulation of Poisson-distributed synaptic inputs. Therefore, we could expect, a priori, some statistical differences between physiological and surrogate data and we would be greatly surprised if they were completely identical in all statistical properties. The fact that the differences in motif statistics between recorded and random data were found to be very small (at the order of few percents at the most), together with the results of our test of synthetic data with planted motifs, strongly support our argument that 1 s repeating motifs emerge stochastically rather than by a network-controlled mechanism.

Comparison to Previous Studies

According to several measures (Friedberg et al., 1999), our recordings were obtained from lightly anesthetized rats. We support this by the fact that membrane potential fluctuations in our study somewhat resemble the membrane potential activity recorded from awake animals (Crochet and Petersen, 2006; DeWeese and Zador, 2006). Cells recorded in the present study show less bimodality of voltage distribution when compared to several previous studies describing UP and DOWN states of membrane potential fluctuations (Anderson et al., 2000; Compte et al., 2003; Leger et al., 2005; Sachdev et al., 2004; Wilson and Kawaguchi, 1996). At the same time, the membrane potential fluctuations that we find in our study span a range of potentials similar to that reported in these previous studies. Analysis of the transition times between UP and DOWN states suggests that these transitions are nonperiodic, uncorrelated, and their generation by stochastic or chaotic oscillators could not be ruled out (Stern et al., 1997). Indeed, recent models of UP and DOWN states attribute this activity to stochastic mechanisms (Compte et al., 2003; Holcman and Tsodyks, 2006).

The stochastic nature of spontaneous activity was challenged in several papers that demonstrated repeating temporal structures. A study which was performed in cortical slice cultures, revealed brief repeating patterns (few milliseconds long) that have been hypothesized to

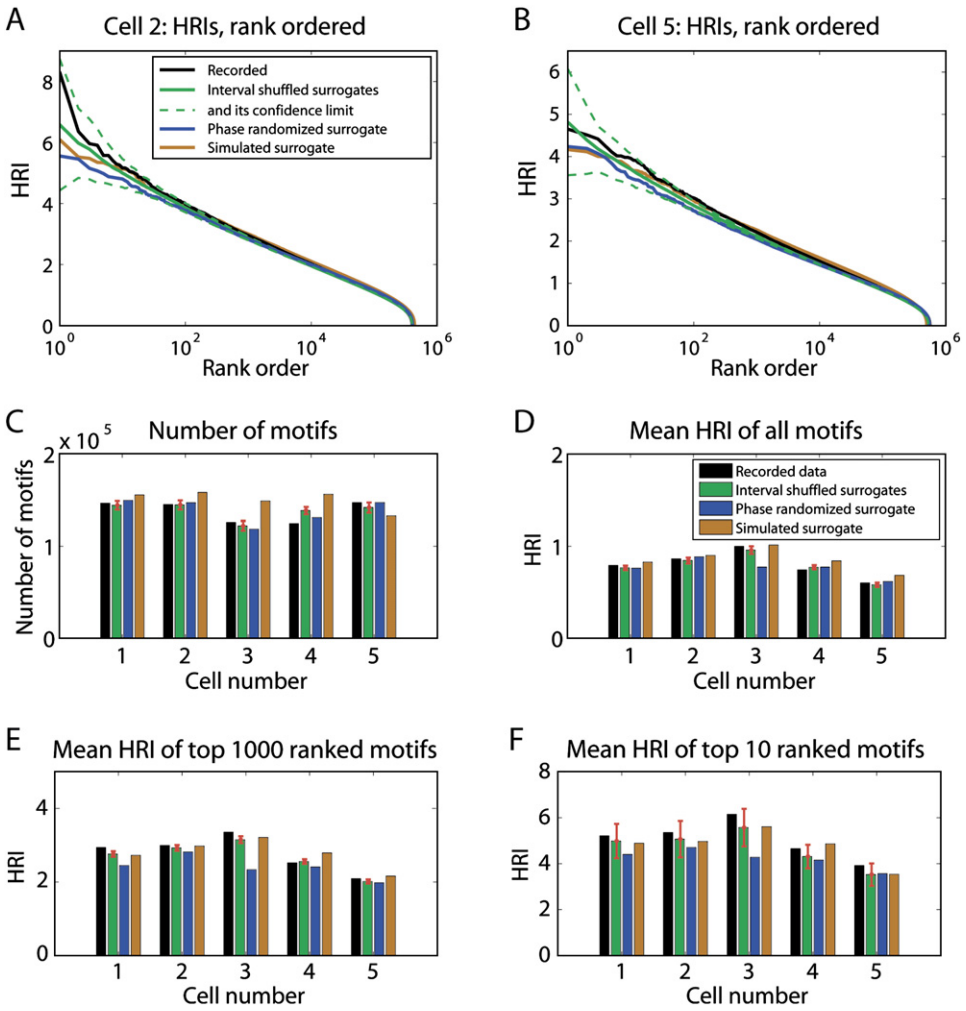

Figure 5. Distributions of Number of Motifs and HRI Scores in Recorded Cells and Their Corresponding Surrogate Data

(A and B) HRI scores for all the motifs found in two different cells, sorted in decreasing order. Recorded data and the three types of surrogate data have very similar distributions of HRI scores (see also Figure S3). The time-domain-shuffled curves (green) are the averages of 40 independent shuffles of the original recording, and the dashed green curves are their confidence intervals ($p = 0.99$).

(C) The total number of motifs found in the first 10 min of recording of each cell and in the corresponding surrogates. Observe that the number of motifs found in the recorded data and in the surrogates produced by time domain shuffling and by phase randomization closely match. Error bars at the time domain interval shuffling bars correspond to upper and lower 99% confidence limits. A two-way ANOVA test shows no significant difference between recorded data and different surrogates and at the same time a significant difference across the cells ($p < 0.05$).

(D–F) Mean HRI score for each cell's original trace and for the corresponding surrogates calculated for the entire distributions, top 1000 ranked motifs, and top 10 ranked motifs, correspondingly. Error bars denote the 99% confidence interval for time domain interval shuffling.

serve as building blocks for information representation (Beggs and Plenz, 2003, 2004). The existence of short patterns in in vivo recordings was not addressed in our study. The most striking evidence for repeating temporal patterns is provided by several studies which demonstrated long repeating patterns in spontaneous firing and voltage fluctuations of cortical neurons (Cossart et al., 2003; Ikegaya et al., 2004; MacLean et al., 2005; Mao et al., 2001). In our study, the number of motifs was higher than previously reported for cells recorded in vivo (Ikegaya et al., 2004), while the mean HRI score was lower. Yet, in both cases the maximal HRI scores were similar. We conjecture that the high number of motifs in our study is a consequence of longer recordings (since the number of candidates grows quadratically with recording duration) and of the differences in the dynamics of the activity.

In previous studies (Ikegaya et al., 2004; MacLean et al., 2005), two statistical methods were used to test whether motif occurrence exceeds chance level. First, recordings from *different* cells were searched for matching patterns (Ikegaya et al., 2004) and the low number of cross-cell motifs was regarded as an indication for the significance of motif appearance in any given cell. However, if the activity dynamics of the two cells is different, as we find

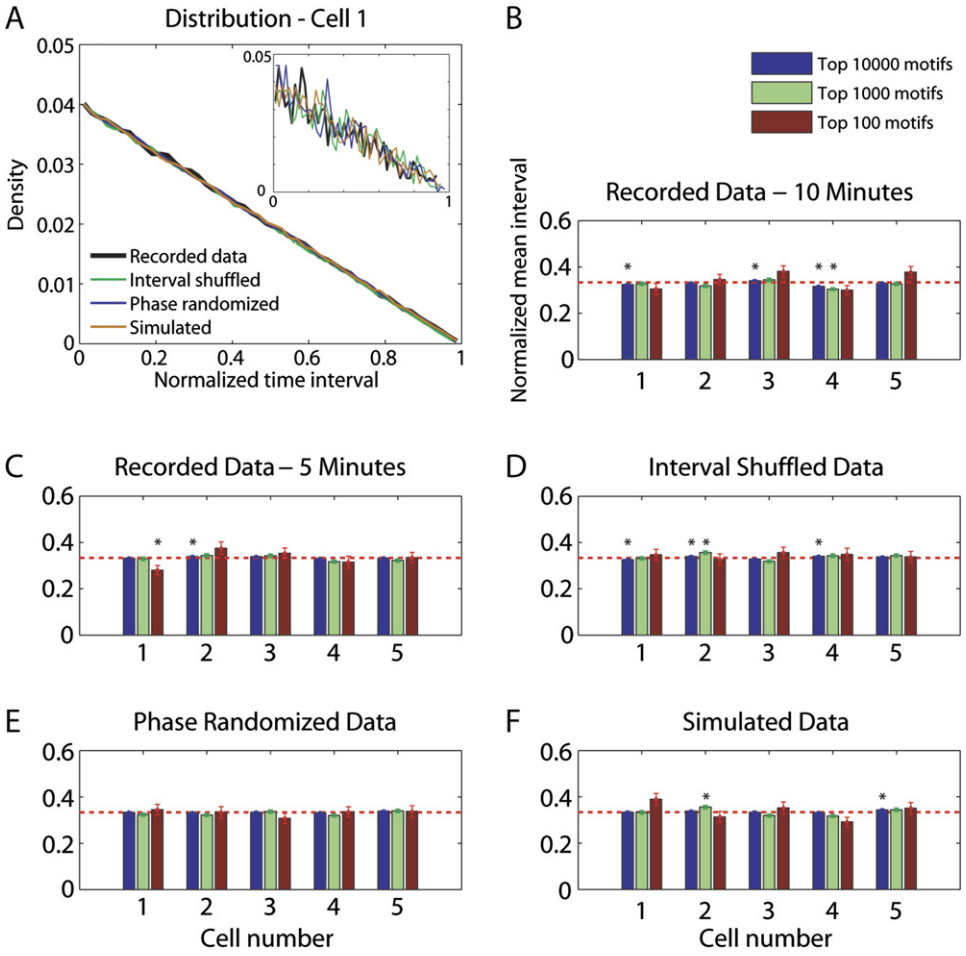

Figure 6. Appearance Times of Motifs Are Random

(A) Density function of the normalized time interval between the repeats of a motif, for all motifs found in the first 10 min of recorded data and three types of surrogate data. Inset shows the density function for the top 1000 ranked motifs. Total recording duration was normalized to a unit length and divided into 50 bins. Frequency of occurrences at each bin was normalized by the total number of motifs to obtain the density function for each data type. The observed density function matches the assumption of randomly distributed motif appearance times.

(B–F) Average time interval between a motif and its repeat. The interval length is normalized by the length of the data. Blue, green, and brown bars are the results for the top 10^4, 10^3, and 10^2 motifs, respectively. The red line denotes the expectation of the distance between two points selected independently and uniformly from a [0,1] interval, which is 1/3. Error bars are SEM and asterisks indicate significant results ($p < 0.01$). Significant but small deviations from this chance level occur even in surrogate data.

in our sampled population, the number of pattern matches across cells is indeed likely to be lower than within cells, whether they occur by chance or not. To support their conclusion, Ikegaya et al. (2004) also compared motif statistics from recorded data against surrogate data. A threshold was set to identify large PSPs, while the voltages below the threshold were set to the threshold value. Then, surrogate data were generated by shuffling the inter-PSP periods. This method transforms the continuous recording into distinct events by a user-defined subjective threshold, thus ignoring a large portion of the recording, such as negative deflections and depolarizing events below the threshold. Using this method, Ikegaya et al. (2004) found that the number of motifs in the surrogate data was considerably lower than in the recorded data. We suspect that this type of shuffling affects local correlations that may exist in the original data, e.g., the typical time between a large PSP and the following PSP, thus reducing the number of motifs in the shuffled data. We note that the above shuffling method was not applicable to our data, since the activity was rarely quiet for a sufficiently long period in order to be substituted by segments of constant voltage.

Appearance of Motifs in Recorded and Randomized Data

In the time-domain-shuffling method used in the present study, the data were fragmented into segments shorter than 500 ms and then randomly reassembled, resulting in a continuous subthreshold activity with voltage distribution identical to the original one. Because fragmented

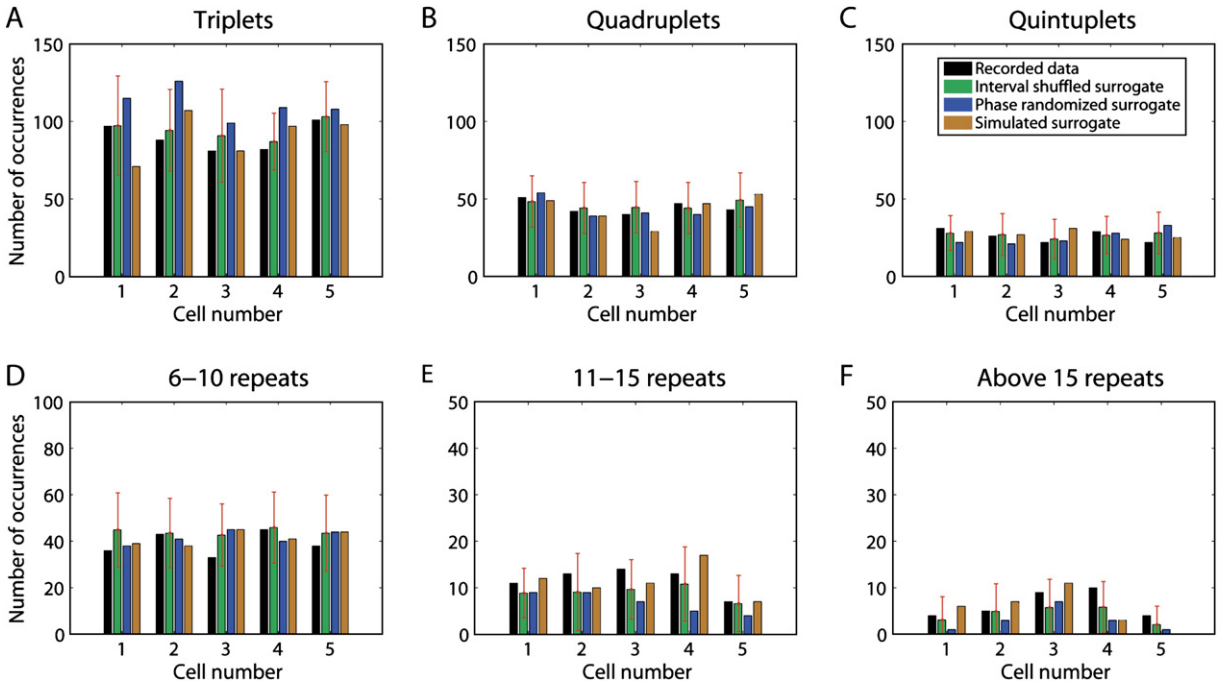

Figure 7. Number of Motifs that Repeat Multiple Times among the Top 1000 HRI-Ranked Pairs, in Recorded Data and Three Types of Surrogate Data

For the interval-shuffled surrogates (green), the mean count in 40 shuffles is presented and the error bars are the confidence interval (p = 0.99). For other types of data, the presented values were obtained from a single randomization. The number of motifs with multiple repeats were within the confidence intervals of the interval-shuffled surrogates in all cells and all data types, except for two cases (number of triplets in phase randomized surrogate of cell 2 and cell 4 [A]).

epochs were relatively long, this shuffling procedure maximally preserved the local correlations in the data, yet enabled determining the significance of motifs of 1 s length. If the top-ranking instances of motifs occur beyond chance level, then cutting these motifs into segments and shuffling them among a much greater number of nonmotif segments is expected to greatly reduce the number of highly similar motifs. Our results show that this is not the case (Figures 5A and 5B). In fact, confidence limits for the time domain shuffling (obtained using multiple independent shuffles), indicate that the top similar motifs are expected by chance. As the similarity index drops, small deviations above the upper confidence limit were found. These small differences are unlikely to reflect the existence of (even a small number of) repeating motifs above the chance level: when we tested our detection method, using an artificially implanted motif (inserted 20 times at random locations), we found that the distribution of HRI scores was profoundly reduced after shuffling (Supplemental Data, Figure S4).

The observed small differences between the physiological and the time-domain-shuffled distributions may arise from several sources of bias introduced into the HRI statistics. The time domain shuffling is limited in its ability to produce dynamics identical to the original data, for several reasons. First, segments longer than 500 ms were excluded from reassembly (1%–2% of the data). Omitting these segments might bias the number of motifs toward lower values. A second type of bias is introduced by (slope) discontinuities at the end points of the fragments, which may produce distorted synaptic events at the reassembly points. The presence of distorted short events is likely to reduce, to some extent, the correlations between epochs. A third source of bias is introduced by a reduction of local correlations that exist in the physiological data and extend beyond 500 ms, which may result from synaptic mechanisms such as short-term depression or slow voltage-dependent currents. Eventually, despite our requirement for a highly stable activity, nonvarying power spectrum and similar membrane potential distribution along the entire recording period, slow and small changes in the dynamics of the physiological activity may exist. Such variations are also expected to be eliminated by the shuffling procedure and therefore bias the results. Despite these factors, the number of motifs and the distribution of similarity index were found to be very close to those of the recorded data.

The two additional randomization methods (phase randomization and the simulation of a point neuron with Poisson-distributed synaptic inputs) provide a better understanding of how motifs emerge out of a stochastic process. The presence of a large number of highly similar motifs in the surrogate data produced by the phase-randomization method is particularly surprising. This finding

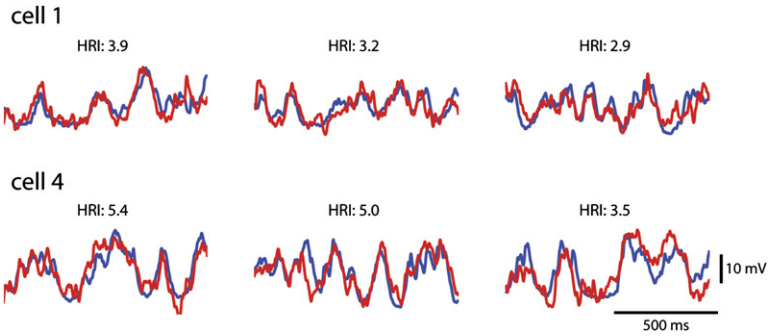

Figure 8. Examples of Pattern Matches between Physiological Traces (Blue) and Simulated Poisson Process Surrogates (Red), from Two Cells

indicates that repeating epochs appear by chance in a noise signal, provided it is long enough and its spectrum is constrained to the one observed in the subthreshold activity. We do note, however, that the resemblance of motif statistics in the original and the phase-randomized data depends on the type of subthreshold activity. For example, phase randomization in cell 3 shows lower mean HRI when compared to the recorded data and other types of randomized data (Figures 5D–5F). This cell had fast and smooth transitions between two voltage states, whereas phase randomization produces noisy fluctuations at all voltage levels. Yet, for most cells, both the time-domain-shuffling and the phase-randomization methods produce a miniscule difference in the number of motifs and in their HRI scores.

In the third method, random data were generated using simulations of a neuron receiving Poisson-distributed synaptic inputs. The Poisson distribution is frequently used to describe the firing of cortical cells (Abeles et al., 1993; Aertsen et al., 1989; Lestienne and Tuckwell, 1998; Riehle et al., 1997). With this method, we have directly tested the hypothesis that repeating motifs may arise from random synaptic inputs, in the absence of any network mechanisms. In most cases, both the number of motifs found in simulated data and their mean HRI scores were higher than those measured in the physiological data. These results should be cautiously interpreted to avoid concluding that the observed spontaneous activity in cortical neurons is a result of Poisson-distributed synaptic inputs. Rather, the aim of the method was to produce a random activity of maximal similarity to the recorded data. Indeed, voltage fluctuations that were produced by the simulation resemble those that were recorded in vivo, and matches of similar patterns across these two data types could be found (Figure 8). The obtained set of simulation parameters was not necessarily within the physiological range of cortical synapses. For example, the time constants of excitatory and inhibitory inputs (on average 5.1 ms and 13.4 ms, respectively) are several times larger than experimentally measured for synapses in cortical neurons (Beierlein et al., 2003; Cowan and Stricker, 2004). Thus, each synaptic event in our simulations may actually represent a transient conductance rise, generated by a synchronized volley of synaptic inputs that lasts up to tens of milliseconds. Nevertheless, regardless of the actual underlying synaptic mechanisms that generate spontaneous activity, the method that was used to simulate the data captured the major dynamics of the activity and indicates that motif statistics can be recapitulated by random synaptic inputs.

Several statistical properties further suggest that small differences between the physiological and surrogate data in the number of motifs and their mean score do not reflect any controlled mechanisms for motif generation. The distribution of the time interval between two repeats of a motif in the recorded data and in the three types of randomized data matches that of a random process (see Figure 6) and seems to be consistent with previous reports (Ikegaya et al., 2004). In addition, the number of triplets, quadruplets, quintuplets, and multiple repeats of higher order is similar to that found in the randomized data (Figure 7). Moreover, HRI statistics differences between the cells were significantly higher than the differences between the data recorded from a cell and its corresponding types of surrogate data (Figures 5C and 5D). This observation implies that the random models that were used to test for motif significance captured the properties of the activity that affect motif statistics. The close similarity between the statistical properties of motif appearance in the physiological data and those found in the surrogate data suggests that coarse factors, such as the dynamical characteristics of subthreshold activity and in particular its limited power spectrum, can closely predict motif emergence statistics. Coarse temporal statistics were also proposed to account for repeating patterns of spikes (Oram et al., 1999).

The stochastic emergence of motifs is not limited to the barrel cortex. We used the time domain interval-shuffling method to test the statistical significance of motifs found in spontaneous activity of three cells recorded from the primary visual cortex of the cat (Ikegaya et al., 2004; Lampl et al., 1999). We found no significant differences between the distributions of motifs in the physiological data and their interval-shuffled surrogates (Figure 9). Thus, we believe that our results may be generalized to other cortical areas and other species. Nevertheless, we cannot exclude the possible existence of deterministically repeating motifs in the subthreshold activity of cortical neurons in brain slices or during sensory stimulation. Sensory stimulation, in particular, may drastically change the dynamics of cortical activity (Anderson et al., 2000). Yet, several studies

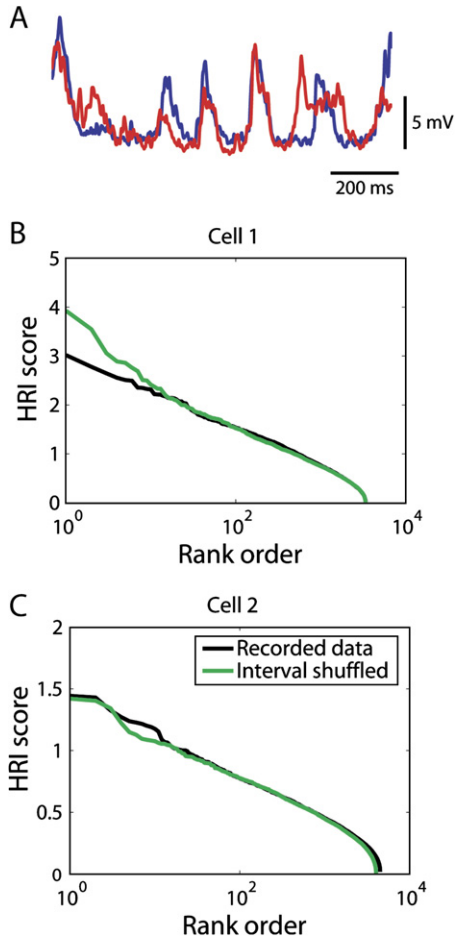

Figure 9. Analysis of Repeating 1 s Patterns in Two Cells Recorded In Vivo from the Cat Primary Visual Cortex (93 and 190 s Recording Durations, Respectively)

We selected for analysis three recordings that met similar criteria as were used for the rat barrel cortex recordings: significant subthreshold activity with amplitude of at least 5 mV, stability of recordings, absence of significant voltage drifts, and absence of significant oscillatory epochs.

(A) An example of high-similarity motif found in recorded data of cell 1 (HRI = 2.8).

(B) HRI scores for all the motifs, sorted in decreasing order. There are no significant differences between HRI distributions of the recorded data and the time domain interval-shuffled surrogate ($p > 0.51$, Kolmogorov-Smirnov) or between the corresponding distributions of the top 100 most similar motifs ($p > 0.27$).

(C) HRI distribution of motifs found in recording from another cell. There are no significant differences between HRI distributions of the recorded data and the time domain interval-shuffled surrogate ($p > 0.07$, Kolmogorov-Smirnov) or between the distributions of the top 100 most similar motifs ($p > 0.9$).

show strong link between evoked and spontaneous activity (Arieli et al., 1996; DeWeese and Zador, 2006; Lampl et al., 1999; Petersen et al., 2003; Sachdev et al., 2004), implying that a large portion of the sensory evoked response is predicted from spontaneous activity and that both may share similar dynamical properties. However, the possibility that nonstochastic, network generated motifs may be found during sensory processing remains open. The methods presented here could be useful in future studies of motifs under different experimental conditions.

EXPERIMENTAL PROCEDURES

Electrophysiology

Intracellular membrane potential recordings of neurons in the barrel cortex of anesthetized rats were performed as described previously (Katz et al., 2006) using the standard patch technique. Briefly, after induction of anesthesia with ketamine (100 mg/kg) and xylazine (1 mg/kg) mixed with saline (50%), tracheotomy was made and the animal was placed in a standard stereotaxic device. Body temperature was kept at 37.0°C ± 0.1°C by means of a heating blanket and rectal thermometer (TC-1000, CWE). Maintenance and control of the anesthesia were achieved by a mixture of air and halothane (<0.8%), by means of artificial respiration at a rate of 100–115 bpm, while monitoring the levels of end-tidal CO_2 and heart rate. Anesthesia level was monitored by heartbeat rate (300–450/m), eyelid reflex, vibrissa, and pinch withdrawal movements. We assess the anesthesia level that was used in our recordings to be between stages III-3 and III-2. Craniotomy of 1 mm in diameter above the barrel cortex was performed, and the dura was removed. All surgical and experimental procedures were performed in accordance with the Weizmann Institute Animal Care and Use Committee.

Standard patch electrodes were filled with an intracellular solution containing 136 mM K-gluconate, 10 mM KCl, 5 mM NaCl, 10 mM HEPES, 1 mM MgATP, 0.3 mM NaGTP, 10 mM phosphocreatine, 310 mM mOsm. Patch electrode resistance was 5–8 MΩ. Series resistance during the recordings was 40–100 MΩ. Compensation for junction potential was not applied. After the electrode was advanced a few microns inside the cortex, the brain was protected from drying by covering the craniotomy with warm agar (3%). Recordings were obtained in the absence of whisker stimulation and current injection. Signals were amplified using Axoclamp-2B (Axon Instruments) and low-passed at 3 kHz before digitization at 15 kHz or 20 kHz.

Data Analysis

Below we provide a short description of the methods used for detection and statistical analysis of motifs in subthreshold recorded activity. For a detailed description of the data analysis methods see the Supplemental Data.

Motif Search Methodology

Repeated segments of subthreshold spontaneous activity of similar shape (cortical motifs) were detected and quantified employing the algorithms described in Ikegaya et al. (2004). Briefly, spikes in long voltage traces (10–20 min) were removed from the digitized data for subsequent analysis by interpolating the membrane potential trajectory from 1 ms before spike threshold to 3 ms after the threshold using a fifth-order polynomial fitting based on three samples on each side of this window. These traces were digitally low-pass filtered at 500 Hz and resampled at 1 ms resolution. We searched the data for motifs of 1 or 2 s length by comparing segments of same length against each other with steps of 250 ms. First, the two segments were aligned to maximize the correlation value. A pair of segments with coefficient of correlation above the threshold value of 0.45 was considered a motif and its HRI similarity score was calculated (see Supplemental Data of the present paper and the supporting online material of Ikegaya et al. [2004]).

High similarity of two segments yields a high HRI score. Other important properties of the HRI score are amplitude invariance and a relative lack of sensitivity for short subregions of dissimilarity. A pair which is similar almost entirely, except for a few short subregions, would still get a high HRI score.

The search algorithm was implemented in C++ (VC7 compiler, Microsoft Windows and GNU C++ compiler, version 4.0.3, Linux).

Surrogate Data

For each recorded cell, we produced three different types of surrogate data to be used in the statistical analysis of motif appearance. The first type of surrogate data is created via time domain interval shuffling (see a schematic illustration in Figure 3A). Briefly, the data are fragmented into intervals shorter than 500 ms, thus breaking each 1 s motif in the original data into three pieces at the least. Then the intervals are randomly shuffled and reassembled. Time domain shuffling produces data which has the same voltage distribution as the physiological data and with almost similar power spectrum density (see Figure S1).

The second type of surrogate data was generated by shuffling the original data in the frequency domain (for a schematic illustration of the method see Figure 3B). We used fast Fourier transform (FFT) to decompose the original data into its frequency components, each characterized by its amplitude and relative phase. While keeping the amplitude of each component, the original phases were replaced by randomly chosen phases, and the surrogate data was created by summing all the phase-randomized frequency components. This procedure generates a trace of identical power spectrum to the original trace, but eliminates any specific temporal structures of the original recording.

The third type of surrogate data was based on a simulation of a passive, single compartment neuron receiving inputs modeled by two excitatory and one inhibitory Poisson process, each representing a different population of synapses (Figure 3C). The rate, strength, and time constants of the inputs were optimized, using a simulated annealing algorithm, to produce subthreshold fluctuations with power spectrum, voltage distribution, and voltage variance that are as similar as possible to the recorded data.

All the computations described above were implemented in Matlab7 (The Mathworks, Inc).

Statistical Analysis

Several statistics were used to compare between recorded data and random surrogates: the number of motifs, their HRI scores, the time between motif appearances, and the number of multiple repetitions.

Confidence intervals for the statistics of motifs found in the time-domain-shuffled surrogates were estimated for the number of motifs, means of HRI scores, and numbers of multiple repeats, from the distributions of 40 random shuffles (for each cell). These confidence intervals (99%), presented in Figures 5 and 7, were calculated based on the t distributions after the null hypothesis of normal distribution could not be rejected (Kolmogorov-Smirnov test, $\alpha = 0.05$). Assuming that the corresponding statistics of other data sets in each cell is also normally distributed, the confidence limits could be used for statistical comparisons.

Two-way ANOVA test was used only in Figure 5C. In this case, the statistic of the number of motifs in the multiple time domain shuffles passed Bartlett's test of equal variances across cells ($p > 0.33$). We also assumed equal variance across cells for the other data sets, where only a single sample was obtained.

Supplemental Data

The Supplemental Data for this article can be found online at http://www.neuron.org/cgi/content/full/53/3/413/DC1/.

ACKNOWLEDGMENTS

We thank David Ferster for helpful discussions and for allowing us to use the data presented in Figure 9 which were recorded in his lab by I.L. and were part of the data analyzed in Ikegaya et al. (2004). We would like to thank Ronny Aloni, Israel Nelken, Yosef Yarom, and Mayer Goldberg for their critical comments on our study and all the members of Lampl's lab for their helpful contribution to this work. We thank Gilad Jacobson for his insightful comments during the preparation of this manuscript. We thank the MOSIX group for providing computational resources on the MOSIX Grid at the Hebrew University. I.L. is an incumbent of the Carl and Frances Korn Career Development Chair in the Life Sciences. This work was supported by grants from The Israel Science Foundation (1037/03), the National Institute for Psychobiology in Israel, by the Henry S. and Anne Reich Research Fund for Mental Health, the Asher and Jeanette Alhadeff Research Award (I.L.), Sir Charles Clore fellowship (M.O.), Toman and Frankel funds of Ben Gurion University of the Negev and the Paul Ivanier Center for Robotics Research (O.B.-S.).

Received: September 25, 2006
Revised: December 11, 2006
Accepted: January 10, 2007
Published: January 31, 2007

REFERENCES

Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. J. Neurophysiol. *70*, 1629–1638.

Aertsen, A.M., Gerstein, G.L., Habib, M.K., and Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity". J. Neurophysiol. *61*, 900–917.

Anderson, J., Lampl, I., Reichova, I., Carandini, M., and Ferster, D. (2000). Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. Nat. Neurosci. *3*, 617–621.

Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science *273*, 1868–1871.

Baker, S.N., and Lemon, R.N. (2000). Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. J. Neurophysiol. *84*, 1770–1780.

Beggs, J.M., and Plenz, D. (2003). Neuronal avalanches in neocortical circuits. J. Neurosci. *23*, 11167–11177.

Beggs, J.M., and Plenz, D. (2004). Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. J. Neurosci. *24*, 5216–5229.

Beierlein, M., Gibson, J.R., and Connors, B.W. (2003). Two dynamically distinct inhibitory networks in layer 4 of the neocortex. J. Neurophysiol. *90*, 2987–3000.

Burgess, N., and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. Hippocampus *6*, 749–762.

Compte, A., Sanchez-Vives, M.V., McCormick, D.A., and Wang, X.J. (2003). Cellular and network mechanisms of slow oscillatory activity (<1 Hz) and wave propagations in a cortical network model. J. Neurophysiol. *89*, 2707–2725.

Cossart, R., Aronov, D., and Yuste, R. (2003). Attractor dynamics of network UP states in the neocortex. Nature *423*, 283–288.

Cowan, A.I., and Stricker, C. (2004). Functional connectivity in layer IV local excitatory circuits of rat somatosensory cortex. J. Neurophysiol. *92*, 2137–2150.

Crochet, S., and Petersen, C.C. (2006). Correlating whisker behavior with membrane potential in barrel cortex of awake mice. Nat. Neurosci. *9*, 608–610.

Dan, Y., Alonso, J.M., Usrey, W.M., and Reid, R.C. (1998). Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. Nat. Neurosci. *1*, 501–507.

Dayhoff, J.E., and Gerstein, G.L. (1983a). Favored patterns in spike trains. I. Detection. J. Neurophysiol. *49*, 1334–1348.

Dayhoff, J.E., and Gerstein, G.L. (1983b). Favored patterns in spike trains. II. Application. J. Neurophysiol. *49*, 1349–1363.

DeWeese, M.R., and Zador, A.M. (2006). Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. J. Neurosci. *26*, 12206–12218.

Eckhorn, R. (1994). Oscillatory and non-oscillatory synchronizations in the visual cortex and their possible roles in associations of visual features. Prog. Brain Res. *102*, 405–426.

Engel, A.K., Konig, P., Kreiter, A.K., Schillen, T.B., and Singer, W. (1992). Temporal coding in the visual cortex: new vistas on integration in the nervous system. Trends Neurosci. *15*, 218–226.

Friedberg, M.H., Lee, S.M., and Ebner, F.F. (1999). Modulation of receptive field properties of thalamic somatosensory neurons by the depth of anesthesia. J. Neurophysiol. *81*, 2243–2252.

Gray, C.M., Engel, A.K., Konig, P., and Singer, W. (1992). Synchronization of oscillatory neuronal responses in cat striate cortex: temporal properties. Vis. Neurosci. *8*, 337–347.

Holcman, D., and Tsodyks, M. (2006). The emergence of up and down states in cortical networks. PLoS Comput. Biol. *2*, e23. 10.1371/journal.pcbi.0020023.

Huxter, J., Burgess, N., and O'Keefe, J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. Nature *425*, 828–832.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., and Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. Science *304*, 559–564.

Katz, Y., Heiss, J.E., and Lampl, I. (2006). Cross-whisker adaptation of neurons in the rat barrel cortex. J. Neurosci. *26*, 13363–13372.

Lampl, I., Reichova, I., and Ferster, D. (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. Neuron *22*, 361–374.

Leger, J.F., Stern, E.A., Aertsen, A., and Heck, D. (2005). Synaptic integration in rat frontal cortex shaped by network activity. J. Neurophysiol. *93*, 281–293.

Lestienne, R., and Strehler, B.L. (1987). Time structure and stimulus dependence of precisely replicating patterns present in monkey cortical neuronal spike trains. Brain Res. *437*, 214–238.

Lestienne, R., and Tuckwell, H.C. (1998). The significance of precisely replicating patterns in mammalian CNS spike trains. Neuroscience *82*, 315–336.

MacLean, J.N., Watson, B.O., Aaron, G.B., and Yuste, R. (2005). Internal dynamics determine the cortical response to thalamic stimulation. Neuron *48*, 811–823.

Mainen, Z.F., and Sejnowski, T.J. (1995). Reliability of spike timing in neocortical neurons. Science *268*, 1503–1506.

Mao, B.Q., Hamzei-Sichani, F., Aronov, D., Froemke, R.C., and Yuste, R. (2001). Dynamics of spontaneous activity in neocortical slices. Neuron *32*, 883–898.

Mazurek, M.E., and Shadlen, M.N. (2002). Limits to the temporal fidelity of cortical spike rate signals. Nat. Neurosci. *5*, 463–471.

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

Petersen, C.C., Hahn, T.T., Mehta, M., Grinvald, A., and Sakmann, B. (2003). Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. Proc. Natl. Acad. Sci. USA *100*, 13638–13643.

Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., and Abeles, M. (1998). Spatiotemporal structure of cortical activity: properties and behavioral relevance. J. Neurophysiol. *79*, 2857–2874.

Richmond, B.J., and Optican, L.M. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. J. Neurophysiol. *64*, 370–380.

Richmond, B.J., Optican, L.M., Podell, M., and Spitzer, H. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. J. Neurophysiol. *57*, 132–146.

Richmond, B.J., Oram, M.W., and Wiener, M.C. (1999). Response features determining spike times. Neural Plast. *6*, 133–145.

Riehle, A., Grun, S., Diesmann, M., and Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. Science *278*, 1950–1953.

Rudolph, M., Piwkowska, Z., Badoual, M., Bal, T., and Destexhe, A. (2004). A method to estimate synaptic conductances from membrane potential fluctuations. J. Neurophysiol. *91*, 2884–2896.

Sachdev, R.N., Ebner, F.F., and Wilson, C.J. (2004). Effect of subthreshold up and down states on the whisker-evoked response in somatosensory cortex. J. Neurophysiol. *92*, 3511–3521.

Schreiber, T., and Schmitz, A. (1996). Improved surrogate data for nonlinearity tests. Phys. Rev. Lett. *77*, 635–638.

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

Stern, E.A., Kincaid, A.E., and Wilson, C.J. (1997). Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. J. Neurophysiol. *77*, 1697–1715.

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

Wilson, C.J., and Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. J. Neurosci. *16*, 2397–2410.