

Neuronal Avalanches Are Diverse and Precise Activity Patterns That Are Stable for Many Hours in Cortical Slice Cultures

John M. Beggs and Dietmar Plenz

Unit of Neural Network Physiology, Laboratory of Systems Neuroscience, National Institute of Mental Health, Bethesda, Maryland 20892

A major goal of neuroscience is to elucidate mechanisms of cortical information processing and storage. Previous work from our laboratory (Beggs and Plenz, 2003) revealed that propagation of local field potentials (LFPs) in cortical circuits could be described by the same equations that govern avalanches. Whereas modeling studies suggested that these “neuronal avalanches” were optimal for information transmission, it was not clear what role they could play in information storage. Work from numerous other laboratories has shown that cortical structures can generate reproducible spatiotemporal patterns of activity that could be used as a substrate for memory. Here, we show that although neuronal avalanches lasted only a few milliseconds, their spatiotemporal patterns were also stable and significantly repeatable even many hours later. To investigate these issues, we cultured coronal slices of rat cortex for 4 weeks on 60-channel microelectrode arrays and recorded spontaneous extracellular LFPs continuously for 10 hr. Using correlation-based clustering and a global contrast function, we found that each cortical culture spontaneously produced 4736 ± 2769 (mean \pm SD) neuronal avalanches per hour that clustered into 30 ± 14 statistically significant families of spatiotemporal patterns. In 10 hr of recording, over 98% of the mutual information shared by these avalanche patterns were retained. Additionally, jittering analysis revealed that the correlations between avalanches were temporally precise to within ± 4 msec. The long-term stability, diversity, and temporal precision of these avalanches indicate that they fulfill many of the requirements expected of a substrate for memory and suggest that they play a central role in both information transmission and storage within cortical networks.

Key words: activity pattern; avalanche; information storage; memory; cortical network; organotypic culture; spontaneous activity; multi-electrode array

Introduction

Because cortex is essential for cognition, many studies have been directed at uncovering mechanisms of information processing and storage there. Previous work from our laboratory (Beggs and Plenz, 2003) revealed a new mode of activity in which propagation of local field potentials (LFPs) in neocortical slice cultures and in acute neocortical slices could be described by the same equations that govern avalanches. These “neuronal avalanches” were generated spontaneously when the cortical network was in a critical state. The critical state was characterized by branching dynamics in which one LFP, on average, led to only one other LFP in the next time step. Although modeling work suggested that neuronal avalanches produced in the critical state were optimal for information transmission (Beggs and Plenz, 2003), the role

that these avalanches could play in information storage has not been clear.

Other laboratories investigating information storage have demonstrated that cortical structures can produce spatiotemporal patterns of activity that are significantly repeatable. Recordings from the hippocampus and neocortex in rats and primates, as well as in the avian motor cortex, have shown that neural activity patterns produced during behavioral episodes are significantly recapitulated during subsequent, but not previous, periods of sleep (Nadasdy et al., 1999; Dave and Margoliash, 2000; Louie and Wilson, 2001; Hoffman and McNaughton, 2002). This body of work strongly suggests that repeatable neural activity patterns could serve as a substrate for information storage in the cortex.

Motivated by these studies, we sought to examine whether neuronal avalanches in critical networks occurred in random patterns, or whether they contained reproducible structure that could satisfy many of the requirements needed to store useful information. If such repeatable activity patterns were to serve as a substrate for memory, they would be expected to exhibit stability over several hours. In addition, they should be diverse to encode the varied experiences of an individual and should not merely reflect a common anatomical theme that would be present in all members of the same species (Weliky et al., 1995). Whereas

Received Feb. 16, 2004; revised April 29, 2004; accepted April 29, 2004.

We thank Veronica Karpiak for expert technical assistance with the preparation of cultures. We also thank Drs. U. Czubayko and J. N. D. Kerr for comments on a previous version of this manuscript.

Correspondence should be addressed to Dr. Dietmar Plenz, Unit of Neural Network Physiology, Laboratory of Systems Neuroscience, National Institute of Mental Health, Building 36, Room 2D-26, 9000 Rockville Pike, Bethesda, MD 20892. E-mail: plenzd@intr.nimh.nih.gov.

J. M. Beggs's present address: Department of Physics, Biocomplexity Institute, Indiana University, Swain Hall West 169, 727 East Third Street, Bloomington, IN 47405-7105.

DOI:10.1523/JNEUROSCI.0540-04.2004

Copyright © 2004 Society for Neuroscience 0270-6474/04/245216-14\$15.00/0

pathological conditions like epilepsy often produce highly repeatable network activity (Albowitz et al., 1990), such a state is not diverse and would not be capable of storing varied memories. Finally, activity patterns would be predicted to possess millisecond temporal precision to allow them to rapidly process information and to perform exact sequence control.

The present experiments were conducted to examine whether neuronal avalanches produced by networks in the critical state could satisfy the requirements of a memory substrate described above. To address these issues, we grew cortical slice cultures on 60-channel microelectrode arrays and recorded their spontaneous LFP activity continuously for 10 hr. Avalanches were extracted off-line, and statistically significant activity patterns were identified as families of avalanches that reoccurred more often than expected by chance. Some of these data have been published previously in abstract form (Beggs and Plenz, 2001).

Materials and Methods

Preparation of organotypic cultures on multi-electrode arrays

For the preparation of cortex–striatum–substantia nigra organotypic cultures (Plenz and Kitai, 1998), coronal sections from rat brains (Sprague Dawley; Taconic Farms, Germantown, NY) at postnatal days (P) 0–2 were cut on a vibroslicer (VT 1000 S; Leica Microsystems Inc., Allendale, NJ). Cortical slices were cut to a thickness of 350 μm and taken from the dorsolateral somatosensory cortex. For the striatum, slices were cut to a thickness of 500 μm . For the substantia nigra (including pars compacta and pars reticulata), ventrolateral sections from mesencephalic slices (500 μm thick) were taken, and medial tissue regions were avoided. After cutting, all tissue sections were submerged in 15 μl of chicken plasma (Sigma, St. Louis, MO) on a multi-electrode array (MEA), and 15 μl of bovine thrombin was added later (1000 NIH units/0.75 ml; Sigma). The tissue was arranged such that the white matter of the cortical slice was aligned with the bottom row of the 8×8 microelectrode array. After plasma coagulation, 800 μl of culture medium was added, consisting of 50% basal medium Eagle, 25% HBSS, 25% horse serum, 0.5% glucose, and 0.5 mM L-glutamine (all from Invitrogen, Grand Island, NY). MEAs were sealed inside a small culture chamber enclosing the electrode array. The array with its chamber was then rotated in a modified “rollertube” incubator (0.6 rpm; Heraeus GmbH, Göttingen, Germany) at 35.5°C in normal atmosphere. After 3 and 27 d *in vitro* (DIV), 10 μl of mitosis inhibitor was added for 24 hr (4.4 mM cytosine-5-b-arabino-furanosid, 4.4 mM uridine, and 4.4 mM 5-fluorodeoxyuridine; calculated to final concentration; all from Sigma). Medium was changed every 3–5 d.

Mature slice cultures prepared from rat cortex at P1 and grown for ~4 weeks have been shown to have anatomical features that closely match those found *in vivo* (Caesar and Schuz, 1989; Caesar et al., 1989; Gotz and Bolz, 1992; Plenz and Aertsen, 1996a,b; Plenz and Kitai, 1996) and develop precise efferent connections with target structures (for review, see Bolz, 1994). Because the current project is part of a larger study to analyze corticostriatal processing, cortex slices were cocultured with striatum and substantia nigra (Plenz and Kitai, 1998). In the mature triple cultures, no projections from these target structures back to the cortex exist, and no differences were found for spontaneous activity in single cortices and cortices cocultured with these target structures (Beggs and Plenz, 2003).

Preparation of MEAs

The organotypic cultures were assembled and grown on MEAs (Multichannelsystems, Reutlingen, Germany). Each MEA consisted of a square (5×5 cm) plate of glass with a circular glass well attached to the center that functioned as a culture chamber (Karpiak and Plenz, 2002). The center of the well contained a square array of 60 microelectrodes made of titanium nitride (8×8 grids with the corners missing). Electrodes were flat and disk-shaped with a diameter of 20 μm , had an inter-electrode spacing of 200 μm , and were attached to gold leads that ended at the edge of the glass square. These leads served as contacts for the

spring-loaded probes of the amplifier headstage (Multichannelsystems). A silver chloride electrode was used for ground reference.

Data acquisition, filtering, and down-sampling of field potentials

Raw data sampled at 1 kHz was low-pass filtered by convolving it with a Gaussian kernel (corner frequency of bode plot, ~50 Hz). This filtering removed high-frequency artifacts and single-unit activity but preserved LFP waveforms. After filtering, the data were usually down-sampled by a factor of 4, resulting in a temporal bin width of 4 msec (i.e., every fourth point of the original data were kept). For comparison, some data sets were also down-sampled by factors of 1, 2, and 8, resulting in temporal bin widths of 1, 2, and 8 msec, respectively. Previous work indicated that binning at 4 msec was most appropriate for cortical culture data collected on these MEAs (Beggs and Plenz, 2003). This is because the spacing between electrodes on the array was 200 μm , and the propagation velocity of activity across the arrays was ~50 $\mu\text{m}/\text{msec}$, giving an average temporal interval between electrodes of 4 msec. We reasoned that binning the data at 8 msec would tend to clump distinct avalanches of activity together, producing an overestimate of true avalanche duration. In a similar manner, binning at 1 msec would tend to separate continuous avalanches, producing an underestimate of true avalanche duration. Accordingly, unless explicitly stated otherwise, all of the results reported here were obtained after the data had been down-sampled by a factor of 4. While down-sampling by factors other than 4 also produced statistically significant activity patterns (see below), these patterns were less likely to reflect the intrinsic timing of the cultures on these MEAs. To obtain a threshold for each channel, the SD was calculated for a 10 min data stream after filtering and down-sampling. A threshold line was chosen at -3 times the SD, well beyond the range of noise. As will be shown later, the results did not differ significantly with the choice of threshold. Previous work (Beggs and Plenz, 2003) used receiver operating characteristic (ROC) curves to calculate the threshold that would minimize both type I and type II errors. The threshold determined by ROC on a similar data set was -2.86 times the SD, suggesting that the threshold of 3 SDs used here was reasonable. LFPs that were more negative than the threshold were then searched for the point of their maximum excursion within 20 msec after threshold crossing. The time of this maximum was recorded as the time of occurrence for the field potential (see Fig. 1). A refractory period of 20 msec was imposed after each maximum to prevent large excursions from being counted more than once. The imposed refractory period was well within the naturally observed refractory period for negative LFP deflections (20 msec). For each suprathreshold LFP, its time point and maximal amplitude were stored (for details, see Beggs and Plenz, 2003).

LFP events were analyzed in binary or amplitude form. In binary form, suprathreshold time bins were represented by ones, whereas subthreshold time bins were represented by zeros. In amplitude form, each suprathreshold event was represented by the maximum amplitude (in microvolts) at that time, whereas subthreshold time bins were represented by zeros. Average event amplitudes divided by the SD of the analog signal were used for the calculation of signal-to-noise ratio based on a 10 sec recording for each electrode.

Data analysis

Extraction of neuronal avalanches. After binning, the data typically showed long periods of quiescence punctuated by brief periods of activity (see Fig. 1B). Frames and avalanches were used to characterize activity periods. A frame was defined as the pattern of suprathreshold activity on the 60-channel electrode grid during one time bin (Jimbo et al., 2000). An avalanche of length N was defined as N consecutively active frames, preceded by a blank frame and followed by a blank frame (see Figs. 1D, 2A). Previous work showed that all cultures self-organized into a critical state, which simulations suggested was optimum for information transmission. This critical state was best observed when the data were binned at 4 msec (Beggs and Plenz, 2003). In an effort to continue to examine the dynamics of these networks at the critical state, we sought to extract avalanches primarily at this temporal resolution (4 msec).

Similarity matrices. Before calculating similarity between avalanches, each two-dimensional frame from an avalanche had to be converted into

a one-dimensional vector. To do this, each 8×8 frame was “unfolded” into a 1×64 vector (the four corner electrodes were always blank, leaving the possibility of only 60 active channels). An avalanche of length N was then formed by concatenating N linear vectors of 1×64 into a single vector of $1 \times N64$. These vectors could then be compared for similarity using three measures: Boolean similarity, binary correlation, and amplitude correlation. “Boolean similarity” between avalanche A and avalanche B was defined as the intersection of active electrodes in avalanches A and B divided by the union of active electrodes in avalanches A and B (Jimbo et al., 2000). Boolean similarity ranged between 0 (least similar) and 1 (perfectly similar):

$$\text{Sim}(A, B) = \frac{(A \cap B)}{(A \cup B)} \quad (1)$$

The correlation between avalanches A and B was given by:

$$\text{Corr}(A, B) = \frac{E(AB) - E(A)E(B)}{\sigma(A)\sigma(B)} \quad (2)$$

where $E(\cdot)$ is the expected value, $\sigma(\cdot)$ is the SD, and AB is the vector dot product.

Correlation values could range from -1 (perfectly antisimilar) to $+1$ (perfectly similar). A “binary correlation” was computed when the correlation formula was given binary data, whereas an “amplitude correlation” was computed when the correlation formula was given amplitude data. When all possible similarity values between a set of avalanches of a given length were calculated, they could be assembled into a matrix (see Fig. 3). Each element (A, B) in the matrix contained the similarity or correlation value between avalanche A and avalanche B . Note that $\text{Similarity}(A, B) = \text{Similarity}(B, A)$, and $\text{Similarity}(A, A) = 1$, leading to a symmetric matrix with ones along the diagonal.

To verify that the results were robust with respect to changes in temporal alignment, some similarity matrices were constructed using a shift procedure. In this shift procedure, the similarity of avalanche A with another avalanche B was calculated under three conditions: a positive shift in time ($+1$), no shift (0), and a negative shift in time (-1). Under a $(+1)$ shift, each frame (i) in avalanche A was compared with frame $(i - 1)$ in avalanche B . In the case where $i = 1$, the first frame of avalanche A was compared with a blank frame. Under a (0) shift, each frame (i) in avalanche A was compared with frame (i) in avalanche B . Under a (-1) shift, each frame (i) in avalanche A was compared with frame $(i + 1)$ in avalanche B . In the case where $i = N$ (where N is the length of the avalanche), the last frame of avalanche A was compared with a blank frame. To form a shifted similarity matrix, only the maximum similarity value from the three shift conditions was placed in the matrix. This shifted similarity matrix ensured that the distinct families of avalanches that were extracted from the data (see below) were not merely the result of a single type of avalanche appearing in three different time positions, as could have been caused by temporal misalignment.

Paired clustering algorithm. A paired clustering algorithm (Matlab 6.0; “dendrogram”) was used to sort the correlation matrix. This procedure grouped avalanches with high similarity values into families. Briefly, this algorithm grouped pairs of avalanches together in order of least distance in similarity space. Thus, if avalanche A and avalanche B were the most similar pair in a set of M avalanches, then avalanche A and avalanche B would be grouped together first. Once grouped together, avalanches A and B would thereafter be represented by a single composite avalanche AB , the location of which in similarity space would be (some function of) the average location of avalanche A and avalanche B . After forming composite avalanche AB , the algorithm would proceed to find the next most similar pair in the remaining set of $(M-1)$ avalanches. The next most similar pair could be formed by two new avalanches, say avalanche C and avalanche D , or it could be formed by a new avalanche with a composite avalanche, say avalanche C and avalanche AB . Thus, for a set of M avalanches, the algorithm performed $(M-1)$ pairings until one large composite family containing all avalanches was left.

For comparison, some avalanches were grouped into families using a simulated annealing algorithm. The advantage of simulated annealing is that it is not a “greedy” algorithm like paired clustering, which could

often get trapped in local minima that represented good, but not best, solutions to the clustering problem. Rather, the simulated annealing algorithm always sought, but did not necessarily find, a global minimum that represented the best solution. In general, the results of this algorithm were similar to those produced by paired clustering, but the simulated annealing process took much longer. For this reason, the results reported here were all derived from paired clustering.

Dendrogram. The sequence in which avalanches were grouped into families by the pairing algorithm was graphically represented as a hierarchy in a dendrogram. At the bottom level of the dendrogram, each avalanche was in a family by itself; at levels above that, several avalanches could be grouped together in one family, and several families could exist; at the top level of the dendrogram, all avalanches were grouped together into one large family (see Fig. 5B). For a set of M avalanches, the dendrogram had $(M-1)$ levels, and each level represented a unique grouping of families.

To create a correspondence between the similarity matrix and the dendrogram, the rows and columns of the similarity matrix were arranged in the order found along the bottom of the dendrogram. This had the effect of concentrating the highest similarity values along the diagonal of the similarity matrix so that families of avalanches could be more easily visualized (see Fig. 5). The square regions of dark pixels along the diagonal formed families of avalanches, and the similarity values found within such squares would, in general, be higher than the similarity values found in regions away from the diagonal that were not in such squares.

Contrast function. For a given level of the dendrogram, a contrast number could be calculated as a measure of the contrast between similarity values found within all families to similarity values found outside of all families:

$$C = \frac{(D_{\text{in}} - D_{\text{out}})}{(D_{\text{in}} + D_{\text{out}})} \quad (3)$$

where D_{in} = (the sum of similarity values in all families)/(the number of similarity values in all families) and D_{out} = (the sum of similarity values outside all families)/(the number of similarity values outside all families).

Because a contrast number could be obtained for every level of the dendrogram, a contrast function could be constructed to indicate the quality of family groupings formed by the clustering algorithm at every level (see Fig. 5C). The peak of this contrast function indicated the level of the dendrogram at which the most distinct grouping of families occurred. To prevent the contrast function from always peaking at the bottom level of the dendrogram where every avalanche was merely paired with itself (a trivial and uninformative result), the diagonal elements of the similarity matrix were changed from all ones to all zeros during contrast calculations.

Shuffling methods. To measure the statistical significance of any putative structure in the data, it is necessary to compare the actual data to what would be expected by chance. Shuffled data sets can provide a way to generate chance expectations without having to make assumptions about how the data should be distributed theoretically. Four methods of shuffling were used in the results reported here: frame, electrode, combined, and matched (see Fig. 6A). “Frame shuffling” randomly permuted the sequence of active frames within a 1 hr recording session and was used to obtain a chance estimate of temporal structure while preserving the original spatial structure. “Electrode shuffling” randomly permuted the location of active electrodes (inactive electrode locations were not permuted) within the 60-electrode array for each avalanche and was used to obtain a chance estimate of spatial structure while preserving the original temporal structure. “Combined shuffling” was equivalent to a frame shuffle, followed by an electrode shuffle, and randomly permuted both the sequence of frames and the active electrodes within each avalanche. Combined shuffling was used to obtain a chance estimate of both spatial and temporal structure. “Matched shuffling” was also used to obtain a chance estimate of spatial and temporal structure but did not alter the firing rates on each electrode (which were altered by electrode shuffling and combined shuffling). Matched shuffling consisted of randomly permuting the activity on each electrode within the time slots of the active

frames. Because each electrode would not, in general, be shuffled the same way as its neighbors, this resulted in a spatial uncoupling of the activity patterns that had been together in the frames of the original data. Because matched shuffling preserved more of the statistical structure of the data than any other form of shuffling, it allowed the most stringent test of significance (Oram et al., 1999).

All four of the above methods of shuffling only operated on sections of the data at which activity was present (i.e., within an avalanche). Other methods of shuffling that operate over the entire length of the data set could have taken activity within avalanches and transferred it to times outside of avalanches. Because all of the putative structure reported here was found within avalanches, these other shuffling methods would have caused a net loss of activity within the avalanches of the shuffled data set. Because these shuffling methods would have artificially biased comparisons between the actual data and shuffled data in favor of finding structure, they were not used to obtain any of the results that pertained to structure within avalanches.

Tests for statistical significance. To measure statistical significance, the actual data were compared with 50 sets of matched shuffled data. While all other methods of shuffling listed above were tried, they were found to be less stringent than matched shuffling (see Results). Shuffled data produced similarity matrices of exactly the same size as actual data, allowing for straightforward comparison.

Significant avalanches. After all 50 similarity matrices from shuffled data were assembled, a maximum matrix could be constructed. The maximum of element (i,j) (representing the correlation between avalanches i and j) from the 50 shuffled matrices was used to create the value of element (i,j) in this maximum matrix. A correlation between avalanches i and j was considered significant at the $p < 0.02$ level if element (i,j) in the correlation matrix from the actual data were greater than element (i,j) in the maximum matrix.

Significant families. Shuffled data could also be used to determine statistically significant families. To do this, shuffled data were subjected to the clustering algorithm, and families were extracted at the level at which the contrast function was at a maximum (see contrast function above). The average correlation within a family was just the sum of the correlations within the family divided by the number of avalanches in the family. The average correlations from all families produced by the shuffled data were then compared with the average correlations from all families produced by the actual data. A family was considered significant at the $p < 0.02$ level if it had a higher average correlation than all families produced by the 50 sets of shuffled data.

Correlations shared over time. Correlations were used to quantify the stability of avalanches between two different time periods (t_1, t_2). To perform this measurement, two activity periods of at least 1 hr each were first recorded. Avalanches were then extracted from the activity of each period, and all avalanches of the same length were assembled into a merged correlation matrix. This merged matrix had three regions, as illustrated in Figure 7A. Black lines divide the merged matrix into correlations within t_1 (square in bottom left), correlations within t_2 (square in top right), and an interaction area that contains correlations between avalanches from t_1 and t_2 (two rectangular regions on either side of the matrix diagonal). The correlations in the interaction area indicate the extent to which the avalanches produced in period t_2 were similar to the avalanches produced in t_1 . The correlations in the interaction area produced by actual data were compared with the correlations in the interaction area produced by 50 sets of matched shuffled data. The fraction of elements in the interaction area that exceeded significance was then plotted as a function of time between the recording periods ($t_2 - t_1$). This fraction gave the probability that an avalanche observed in time t_1 would be repeated in time t_2 in a manner that was more similar than expected by chance. Although this fraction indicates how often avalanches are being repeated in a statistically significant manner, it tells nothing about which avalanches are being repeated. For example, it might be possible that 5% of the avalanches recorded in the second hour are significantly similar to the avalanches produced in the first hour. In the third hour, 5% of the avalanches might also be significantly similar to the avalanches produced in the first hour. However, the avalanches produced in the second and third hours might share no significant similar-

ities. To examine this possibility, another measure of stability over time was also used.

Mutual information. Mutual information was used to evaluate whether the avalanches repeated in later time periods were, in fact, the same subset of avalanches hour after hour. To estimate the mutual information shared between two different time periods (t_1, t_2), a distribution of significant correlations was extracted from all avalanches of a given length from times t_1 and t_2 . A distribution of significant correlations was also extracted from the correlation matrix produced by the merger, or union, of t_1 and t_2 (as described above). The entropy (H), in bits, of each distribution is defined as follows (Shannon and Weaver, 1949):

$$H(t_1) = - \sum p_{t_1} \cdot \log_2(p_{t_1}) \quad (5)$$

$$H(t_2) = - \sum p_{t_2} \cdot \log_2(p_{t_2}) \quad (6)$$

$$H(t_1 \cup t_2) = - \sum p_{t_1 \cup t_2} \cdot \log_2(p_{t_1 \cup t_2}) \quad (7)$$

The mutual information, in bits, between two time periods is then given by:

$$MI(t_1, t_2) = H(t_1) + H(t_2) - H(t_1 \cup t_2) \quad (8)$$

Intuitively, the mutual information is a measure of the degree of overlap between the distributions produced at t_1 and t_2 . If these distributions are exactly the same (complete overlap), then the merged distribution will be exactly the same also, and mutual information will be at a maximum, equal to the entropy of one of the distributions [$H(t_1) = H(t_2) = H(t_1 \cup t_2) = MI(t_1, t_2)$]. If these distributions are completely different (no overlap), then the merged distribution, which is the union of the distributions produced at t_1 and t_2 , will have entropy that is equal to the sum of the entropies associated with t_1 and t_2 , and mutual information will be at a minimum [$H(t_1) + H(t_2) = H(t_1 \cup t_2) = MI(t_1, t_2) = 0$]. Mutual information was plotted as a function of time between the recording periods ($t_2 - t_1$). Note that all random contributions to mutual information are removed with this method, because only statistically significant correlations are used to compose the distributions. Mutual information provided a way to measure whether the avalanches that shared correlations were, in fact, the same avalanches hour after hour.

Jittering. A jittering procedure was used to evaluate the temporal precision of avalanches. The time points of suprathreshold activity, taken from original raster data, were each displaced (positively or negatively) by a small amount of time randomly drawn from a Gaussian distribution with a known SD. The SD of the Gaussian distribution was varied between 2 and 20 msec, in steps of 2 msec. A correlation matrix composed of all avalanches of length 5 was assembled for each culture, and the average correlation value within each matrix was extracted. Avalanche length 5 was chosen as an arbitrary representative of patterns produced by the cultures; avalanches of length 4 and 6 were also tested, and there were no significant differences in results between avalanche lengths. The average of all correlation values was then plotted against the standard deviation of Gaussian jitter. For this procedure, the original data were binned at 1 msec to allow maximum resolution of temporal precision.

Measures of robustness. Some remarks should be made about the choice of analysis parameters used to obtain the present results. There were six different analysis parameters, and each of these parameters could have been chosen in more than one way: threshold (-2 SD, -3 SD, -4 SD); bin width (1, 2, 4, 8, and 16 msec); data type (binary, amplitude); similarity measure (correlation, Boolean); shuffling type (frame, electrode, combined, matched); shift check (no shift, shift). The total number of possible combinations of these analysis parameters is 480, a number that would take prohibitive computing time to fully explore. The solution adopted here to this problem of multiplicity was to choose one standard set of plausible parameters (justified in Results). The robustness of this standard set was then explored by systematically varying one parameter at a time through all of its possible choices, while keeping the other parameters fixed at their standard values. The fraction of statistically significant correlations produced in 1 hr for avalanches of length 5 was used as the output variable of interest. The results produced by system-

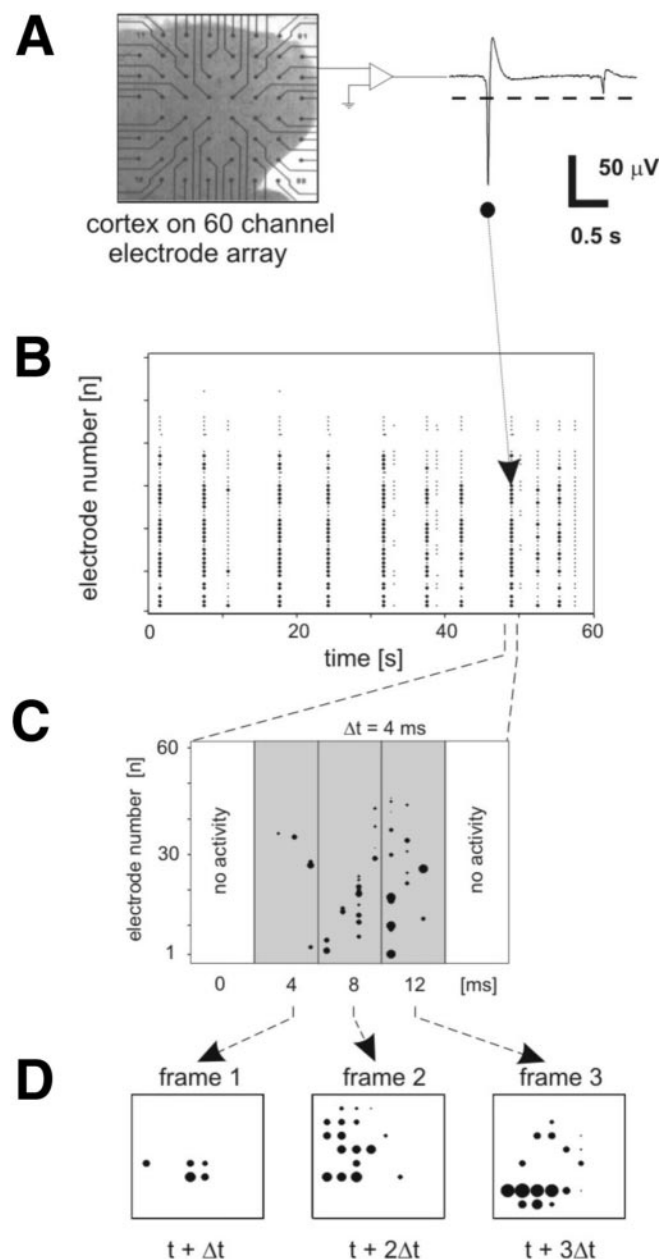


Figure 1. Multi-electrode recording of spontaneous activity in cortical slice culture. *A*, Left, Photo of culture grown on an MEA dish taken at 2 DIV. The interelectrode distance is 200 μm . Right, Raw activity recorded from one electrode, showing a dashed threshold line at -3 SDs. *B*, Raster plot shows periods of nearly synchronous activity separated by quiescent intervals of several seconds. *C*, Period of suprathreshold activity near the 50 sec mark in raster is binned at 4 msec, showing that activity is not exactly synchronous. Avalanches are defined as successive frames of activity with at least one electrode active. The avalanche shown spans three time bins and is preceded and terminated by a frame with no activity. *D*, One-dimensional arrangement of electrodes in a raster plot can be displayed as a two-dimensional frame matching the MEA grid. The activity shown in *C* is presented as a spatiotemporal sequence of activity on the MEA grid. In this case, an avalanche of three frames is shown. Figure 1 is reproduced from Beggs and Plenz (2003) (©2003 Society for Neuroscience).

atic variation were then compared with those generated by the standard set of analysis parameters.

Results

The results are composed of six main sections. First, we present general statistics, describing the cultured cortical networks and their spontaneous activity. Second, we establish the existence of

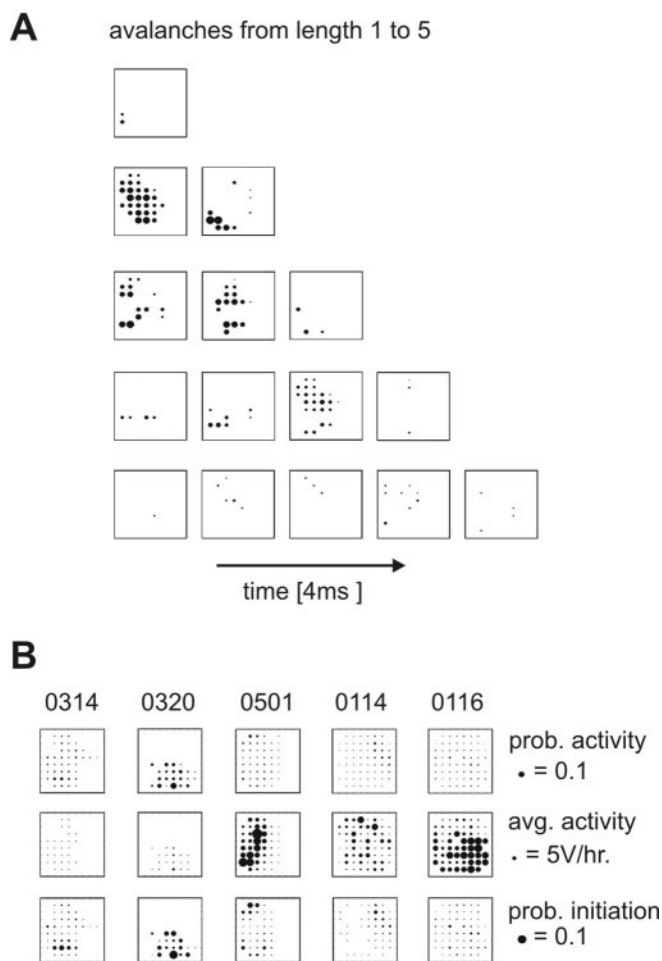


Figure 2. Summary of spatial activity distributions for all five networks in the current study. *A*, Five examples of single neuronal avalanches produced by the same culture. As before, the size of the dot indicates amplitude recorded at that electrode. Note that activity is not often wave-like and may involve large or small amplitudes. *B*, Statistics of activity for five different cultures. Numeric culture names are given at the column heads, whereas types of statistics and dot scales are given at the end of each row. Probability of activity (prob. activity): how likely that an electrode was to be active at any time, regardless of amplitude; average activity (avg. activity): how much a particular electrode contributed to the total activity (probability of activity times amplitude); probability of initiation (prob. initiation): how often that electrode participated in the initiation of an avalanche.

statistically significant activity patterns. Third, we demonstrate the long-term stability of these patterns over 10 hr. Fourth, we address the diversity of significant activity patterns, showing that there is an average of 30 ± 14 distinct families of patterns per culture. Fifth, we show that these patterns have a temporal precision of ± 4 msec. Finally, we show that the methods used to extract activity patterns are robust.

General description of cultures and spontaneous activity

The present results summarize data from five organotypic cortex–striatum–substantia nigra triple cultures recorded after 28 ± 3 DIV (mean \pm SD). For the analysis of cortical neuronal activity, spontaneous LFPs (Fig. 1*A*) were recorded continuously from the cortex for at least 10 hr in each culture at a 1 kHz sampling rate. A typical LFP consisted of a negative population spike superimposed on a positive envelope (Fig. 1*A*). This general time course was consistent over time at individual electrodes and across electrodes. On average, 49 ± 15.3 electrodes per culture

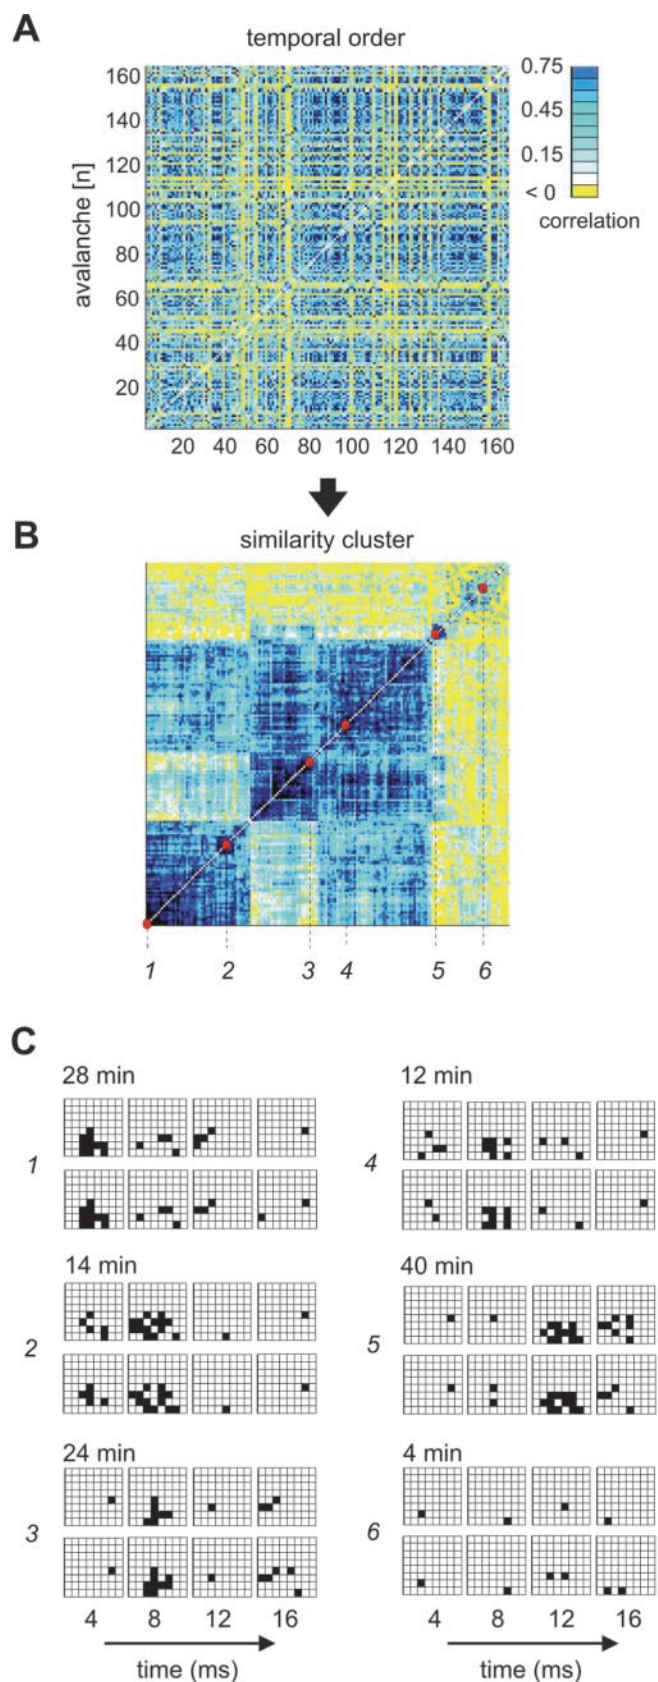


Figure 3. Sorted correlation matrix reveals similar avalanches. *A*, Unsorted correlation matrix of all avalanches of length 4 recorded in 1 hr from one culture (bin width, 4 msec). Avalanches are numbered in order of appearance along the margins of the matrix. Diagonal elements have been removed for clarity. Note that blue (high correlation) and yellow (low correlation) regions are intermingled. *B*, Same correlation matrix as in *A*, now sorted in order of

were active [i.e., had negative population spikes that crossed threshold (-3 SD)]. For additional quantification, the maximal negative excursion of a suprathreshold LFP within a 20 msec local search window (see Materials and Methods) was called an event, and its time of occurrence and amplitude were recorded. Events occurred at an average rate of 0.16 ± 0.07 Hz (active electrodes for all cultures), and in all cultures, interevent interval histograms (IEI) showed that activity had a refractory period that always exceeded 20 msec at single electrodes. Previous work has shown that activity at single electrodes is not periodic but has IEIs that vary in length (Tateno et al., 2002; Beggs and Plenz, 2003).

At low temporal resolution, events appeared on all active electrodes simultaneously forming synchronized activity epochs (Fig. 1*B*), in agreement with previous reports of synchronized bursting in dissociated and cortical slice cultures (Crain, 1966; Calvet, 1974; Gutnick et al., 1989; Maeda et al., 1995; Gopal and Gross, 1996; Kamioka et al., 1996; Plenz and Aertsen, 1996*b*; Plenz and Kitai, 1996). Our recent analysis revealed that these epochs are composed of neuronal avalanches, successive but distinct periods of neuronal activity with inherently complex spatio-temporal structure (Fig. 1*C,D*) (Beggs and Plenz, 2003) (see Materials and Methods). Over the course of 1 hr, each culture produced an average of 4736 ± 2769 neuronal avalanches at a temporal resolution of 4 msec. As described in Materials and Methods and previously (Beggs and Plenz, 2003), a bin width of 4 msec captured the inherent conduction velocity of neuronal activity and also allowed the dynamics of the critical state to be observed with optimum resolution.

We first demonstrate the variability of these neuronal avalanches by analyzing their initiation site distribution as well as the average activity across each network (Fig. 2). The electrodes that were active at the start of each avalanche (in the first frame) were termed the initiation sites. Rather than being concentrated at one or two electrodes, the initiation sites in all cultures were distributed, showing some regional differences, but indicating that there was no single focus from which activity originated (Fig. 2*B*). The largest probability of initiation at any electrode was 0.11, indicating that it was more likely to be inactive than to initiate. In a similar manner, the probability that a given electrode was active was also distributed throughout the MEA grid with modest variability (Fig. 2*B*). The average activity at an electrode was taken to be the product of the average field potential amplitude and the probability of activation at that site, and also showed regional variations (Fig. 2*B*).

Statistical significance of activity patterns

To investigate whether there was structure in the spontaneous activity produced by these cultures, all avalanches of the same length recorded in 1 hr were checked against each other for similarities. Similarity between avalanches was quantified by a correlation coefficient or a measure of Boolean overlap (see Materials and Methods). When all combinations of correlation coefficients were calculated, they were color-coded and placed into a corre-

←
similarity by clustering algorithms. Note that blue regions are now concentrated along the diagonal, forming groups of avalanches with high mutual correlations. The red dots indicate regions of highly similar avalanches (numbered 1–6). *C*, Highly similar avalanches revealed by clustering. Each number shows a pair of avalanches that were clustered together at red dots in the sorted similarity matrix. The frames are shown as 8×8 grids, and suprathreshold electrodes are indicated by black squares. Differences in time of occurrence between avalanches are given in min. Note that avalanches show high spatial and temporal similarity despite occurring many minutes apart.




Figure 4. Examples of avalanche patterns from a single culture that are highly repeatable even after many hours. *A*, Avalanche pattern of length 4 was spontaneously reproduced with high similarity by the culture seven times over a period of 7.14 hr. *B*, Another avalanche pattern repeated six times in 6.10 hr. *C*, Sparse avalanche pattern reproduced four times over 7.82 hr. As before, suprathreshold electrodes are indicated by black pixels.

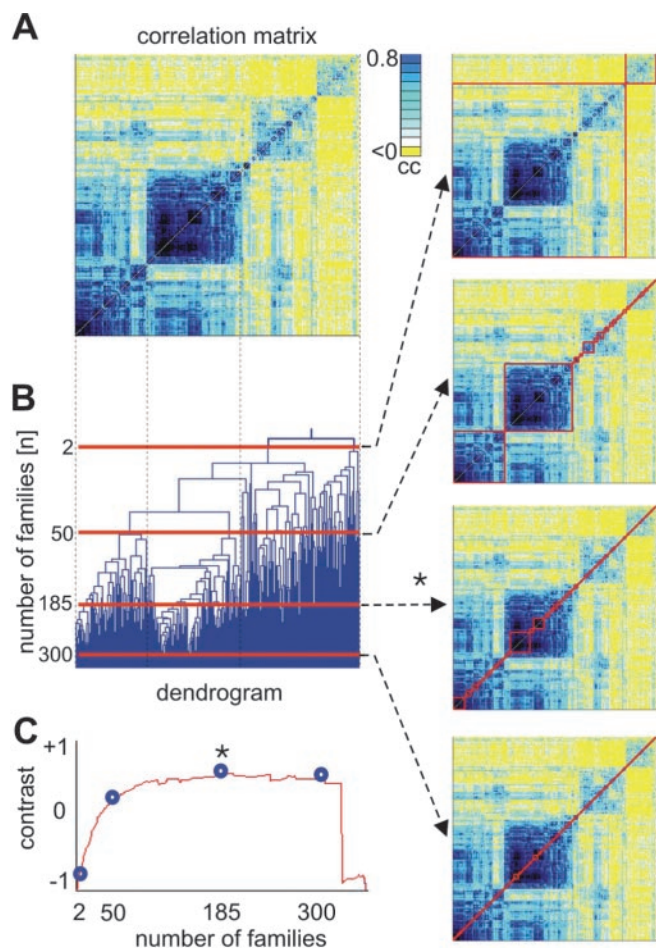


Figure 5. Levels of the dendrogram correspond to different sets of families in the correlation matrix. *A*, Correlation matrix is positioned directly above its dendrogram. *B*, Dendrogram of correlation matrix (avalanche length, 4; bin width, 4 msec). At the top of the dendrogram, a single branch signifies that all avalanches are in one family. Just below this, the dendrogram divides into two branches, representing a set of two families. Branching continues further down the dendrogram until every avalanche is in its own family. The number of times each horizontal red line crosses the dendrogram indicates the number of families at that level (2, 50, 185, and 300). *B*, Right, Correlation matrix can be partitioned into family sets that correspond to horizontal lines on the dendrogram. Dotted arrows proceeding from each level of the dendrogram indicate how the correlation matrix is partitioned into families of avalanches at each level. Red boxes surround each family. The top matrix shows two large families, corresponding to the two branches in the dendrogram that are crossed by horizontal red line number 2. The next red line crosses 50 branches of the dendrogram, and the corresponding dotted arrow points to the matrix with 50 families surrounded by red boxes. The bottom matrices correspond to groupings with more, but smaller, families. Note that all family groupings merely represent different ways of partitioning the same correlation matrix. *C*, Contrast function plots contrast (compares density of correlations within red boxes to density outside of red boxes; see Materials and Methods) against the number of families, and contrast peaks at 185 families (asterisk), indicating optimum grouping into families for this matrix. The blue circles on the contrast curve correspond to the red lines drawn on the dendrogram and their corresponding groupings of families shown in matrices at right.

lation matrix (Fig. 3*A*). Visual inspection of correlation matrices from actual data revealed a large number of positive correlations, suggesting that many of the avalanches produced were unusually similar to each other. Clustering algorithms (see Materials and Methods) were used to rearrange a correlation matrix from temporal order (Fig. 3*A*) to order of similarity (Fig. 3*B*). When presented in order of similarity, several large squares of high correlation values appeared on the diagonal of the matrix, suggesting that the avalanches could be grouped into several families that

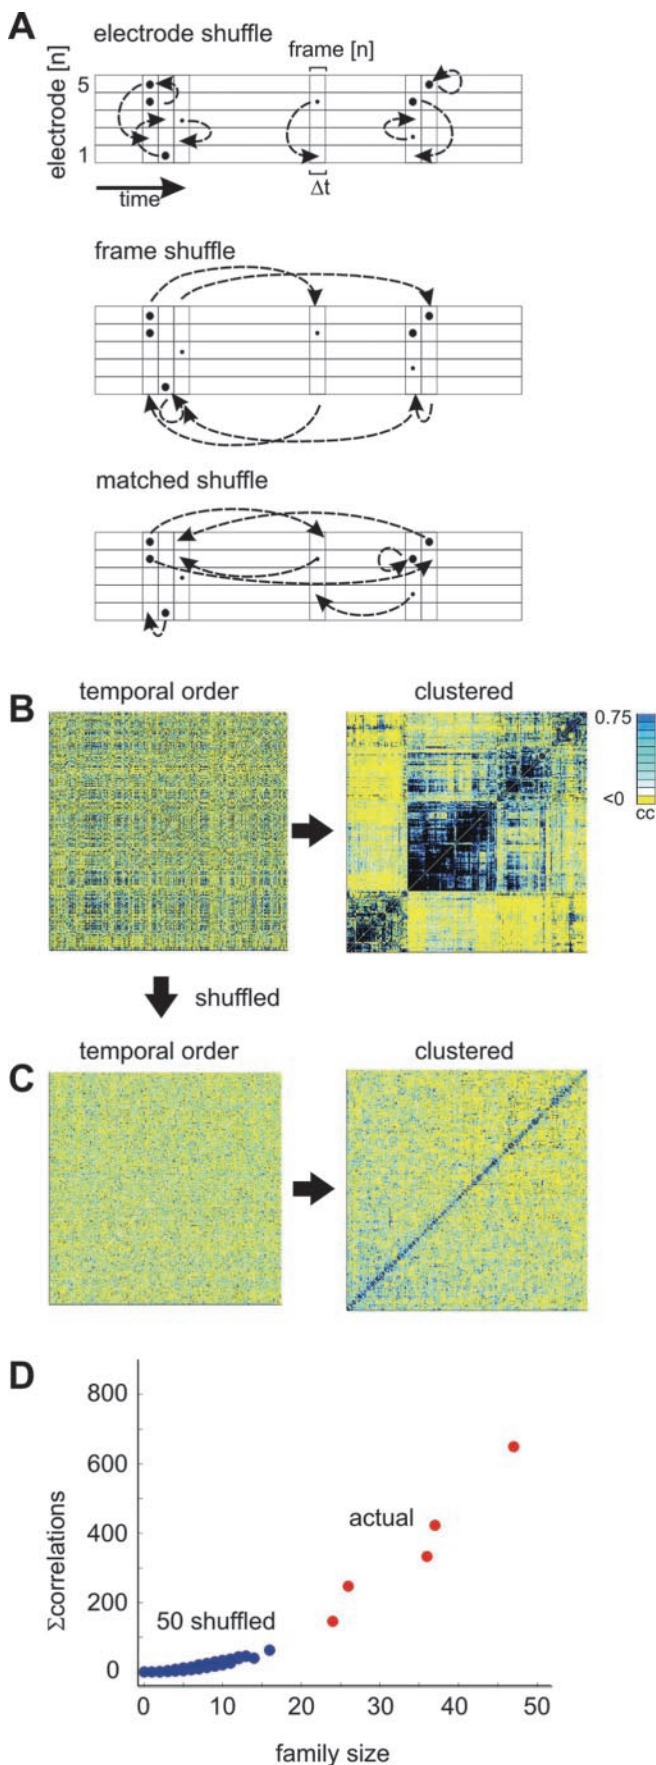


Figure 6. Shuffling reveals significant families of avalanches. *A*, Electrode shuffling of a raster randomly permutes active electrodes within each frame but does not change the times of

were highly similar within themselves. Indeed, within these families there could be found avalanches that were strikingly similar, despite the fact that they had been produced by the culture many minutes apart (Fig. 3*C*). These similarities could remarkably persist even for several hours (Fig. 4).

Within several of the large families in a sorted matrix, there were smaller subfamilies with even greater correlations (Fig. 5). The relationships between these families and subfamilies could be succinctly described by a dendrogram (Fig. 5*B*). A horizontal line drawn across the dendrogram cut across many branches, at which each branch corresponded to a family of avalanches (Fig. 5*B*). Each line that cut across the dendrogram at a different level represented a different way of grouping the avalanches into families. The optimum grouping was defined as the one that maximized the contrast between correlations within families and correlations outside of families (Fig. 5*C*) (see Materials and Methods).

Although the highly similar avalanches and the complex family structure were intriguing, it remained possible that they could have been caused by chance. To investigate whether these apparent structures produced by the cultures were statistically significant, the actual data were compared with 50 sets of shuffled data. Several different procedures of shuffling were used (Fig. 6*A*) (see Materials and Methods), and all produced qualitatively similar results. Initial visual comparison of the correlation matrices from actual data and shuffled data (Fig. 6*B, C*) clearly hinted that more high correlations and more dense family groups could be found in the actual data than in the shuffled data sets. In fact, the actual data yielded correlation matrices with significantly ($p < 0.02$) higher sums of correlations than the matrices obtained from shuffled data. When the correlation matrices from shuffled data were clustered into families at maximum contrast, they had significantly ($p < 0.02$) smaller correlation densities (sum of correlations within a family divided by the size of the family) (Fig. 6*D*). Statistically significant families were produced by actual data from all cultures at most avalanche lengths. The average number of different, but statistically significant, families produced per culture in 1 hr was 30 ± 14 . Thus, the spontaneous activity of the cortical cultures contained a wide variety of complex spatiotemporal patterns that were repeated within 1 hr far more often than would have been expected by chance.

Stability of activity patterns over many hours

Although it was apparent that there were numerous significant activity patterns within 1 hr of spontaneous activity, it remained

←
 the frames. Any electrode that has had suprathreshold activity over the course of 1 hr is considered active. Electrode shuffling creates a data set with spatial correlations expected by chance. Frame shuffling randomly permutes times of all active frames but does not change the locations of the electrodes. This creates a data set with temporal correlations expected by chance. Matched shuffling (short for activity count matched shuffling) randomly permutes the suprathreshold activity on each electrode within the times allowed by the active frames. Note that activity on each electrode is permuted independently of other electrodes, thereby disrupting spatial correlations. Matched shuffling creates a data set with spatial and temporal correlations expected by chance, while preserving the activity rates on all electrodes and the times of active frames. *B*, Left, Unsorted correlation matrix of all avalanches of length 4 from actual data. Right, Same matrix now sorted to reveal clusters along the diagonal. *C*, Left, Unsorted correlation matrix of all avalanches of length 4 from matched shuffled data. Right, Same matrix now sorted. Note how actual data has larger blocks of high correlations along the diagonal (dark blue) than shuffled data. *D*, Clusters formed from actual data (red circles) have larger size and greater sum of within-cluster correlations than clusters formed from 50 sets of shuffled data (blue circles). Families of avalanches corresponding to the red circles are therefore significant at the $p < 0.02$ level.

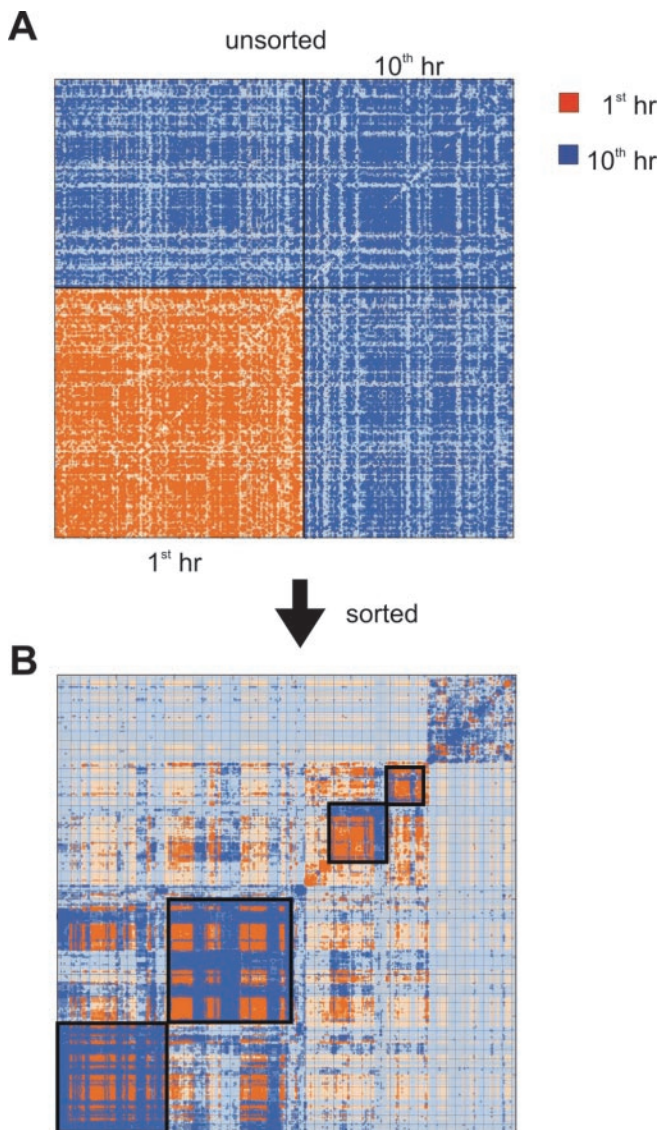


Figure 7. Clustering reveals that avalanches are similar across time. *A*, Unsorted correlation matrix contains avalanches from different times. Rasters from the first and the tenth hour were merged, and all avalanches of length 4 were extracted. Correlations among avalanches from the first hour are red, all others blue. High correlations are indicated by large dots, low correlations by small dots, and negative correlations by no dots. Because correlations are presented in temporal order, red dots are to the left. *B*, Correlation matrix sorted in order of similarity clusters red and blue dots together into families. Families outlined by black boxes are composed of approximately equal portions of red and blue, indicating that avalanches from the first hour are similar to avalanches from the tenth hour.

to be seen whether these patterns would persist over time. To address this question, a single correlation matrix was prepared that contained all the avalanches of a given length from the first hour of activity and all the avalanches of the same length from a later hour of activity (Fig. 7*A*). This merged correlation matrix was then sorted by the clustering algorithm to see whether the avalanches produced hours apart would cluster together into the same groups. By marking all of the correlations produced solely within the first hour in red, and all other correlations in blue, it became apparent that avalanches were still similar because they coalesced into clusters containing both red and blue (Fig. 7*B*). Although this result was suggestive, additional testing was needed to demonstrate that this grouping was statistically significant.

To examine whether significant correlations between ava-

lanches were preserved over time, merged matrices were again created. Each merged matrix had three regions: (1) a square in the bottom left containing all the correlation values produced by avalanches found entirely within the first hour; (2) a square in the top right containing all the correlation values produced by avalanches found entirely within a later hour; and (3) two rectangles on either side of the diagonal, each representing the correlation values produced by the interaction of avalanches from the first hour with avalanches from a later hour (Fig. 8*A*). In most merged matrices, the high correlations found in the rectangular interaction area appeared to be as prevalent as the correlations found in either hour alone. This suggested that the avalanches produced at later times remained similar to the original set of avalanches. To test whether these correlations were significant, they were compared with the correlations produced in the interaction areas of 50 shuffled data sets. As shown in Figure 8*B*, the correlations produced by the actual data were significantly greater than those produced by the shuffled data at all time differences tested. Thus, even 10 hr later, the culture was still producing complex spatio-temporal patterns that resembled those produced in the first hour.

Although there were significant correlations preserved over time for all times tested, the fraction of correlations that were significant within an interaction area did decline somewhat after longer times. This fraction reflects the probability that an avalanche produced in a later time period would share significant correlations with an avalanche of the same length produced in the first time period. Because this fraction retained most of its original value, most of the statistically significant avalanches produced in the first hour were likely to be produced in later hours also. To quantify the retention of similar avalanches, the fraction of significant correlations in the interaction area was plotted over time for all cultures (Fig. 9*D*). The average retention over 5 hr was $84 \pm 16\%$ and over 10 hr was $78 \pm 16\%$ ($n = 5$).

Another index of retention was the mutual information between distributions of correlation values produced in the first hour and later hours (Fig. 9*A*). Plots of the mutual information showed an average retention of $100 \pm 16\%$ after 10 hr (Fig. 9*B, C*). This high retention indicates that the shape of the distribution of correlation values remained virtually unchanged over time.

Diversity of activity patterns

Figure 10 shows the average number of statistically significant families and the average number of avalanches in these families for all cultures. Avalanche length 1 was excluded from consideration because of its brief temporal extent, and avalanche lengths >9 occurred so infrequently that they generally failed to produce significant families. In most cultures, significant families were found in avalanche lengths 3–9, and, on average, each culture produced 30 ± 14 statistically significant families in 1 hr. The total number of avalanches that were grouped into significant families was, on average, 697 ± 437 , meaning that the average significant family contained 23 avalanches. Electrodes that were active in one significant family could also be active in others, indicating that significant families were not merely independently active groups of electrodes.

Temporal precision of activity patterns

To probe the temporal precision of the activity patterns, the time points from the original data binned at 1 msec in raster form were randomly jittered by different amounts (see Materials and Methods). The average correlation within a correlation matrix was

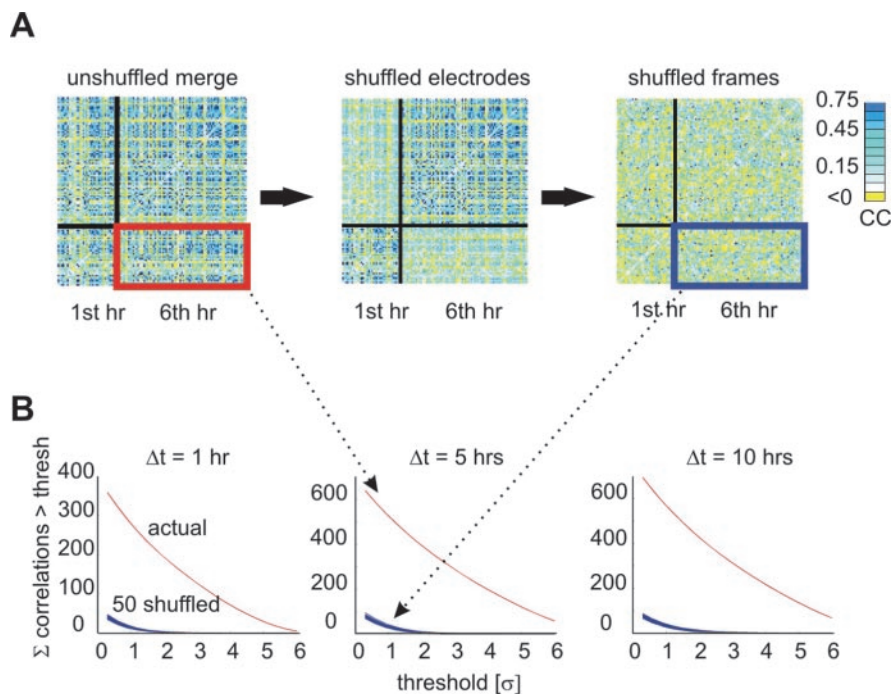


Figure 8. Significant structure across many hours. *A*, Unsorted correlation matrices contain avalanches of length 4 from first and sixth hours. The total number of avalanches in both matrices is 120. Black lines divide the matrix into correlations within the first hour (small square, bottom left), correlations within the sixth hour (large square, top right), and interaction area that contains correlations between hours (two rectangular regions on either side of the matrix diagonal). The leftmost matrix is an unshuffled merge at which the interaction area contains many high correlations (blue pixels). The middle matrix is a merge at which data from the first hour had all electrodes randomly permuted the same way (done here merely to illustrate decoupling the hours, actual electrode shuffling was more complete, as described in Materials and Methods). This remapping preserved correlations within each hour but reduced correlations between hours, indicated by reduced blue pixels in the interaction area. The rightmost matrix is a merge at which data from both hours has had shuffled electrodes and frames. This interaction area gives correlations between the first and sixth hours expected by chance. *B*, Actual data are significantly different from shuffled data. The sum of correlations in the interaction area from unshuffled data (red box) plotted as a red curve against threshold for three time differences Δt . The sum of correlations in the interaction area from shuffled data (blue box) plotted as 50 blue curves against threshold for each time difference. Note that the curve for actual data are far above curves for shuffled data. Also note changes in the vertical scale. The ratio of actual correlations to shuffled correlations is fairly constant for 10 hr and remains statistically significant.

then measured as a function of jitter. The avalanches could be considered temporally precise only if their average correlation dropped sharply with increasing jitter. On average, the cultures lost 49% of their original correlation values after a jitter of only ± 4 msec (significant at $p < 0.01$; Tukey multiple comparisons test), suggesting that these avalanches were temporally precise (Fig. 11).

Robustness of methods

The above results were obtained with the six analysis parameters set at their “standard” values (threshold, -3 SD; bin width, 4 msec; data type, amplitude; similarity measure, correlation; shuffling, matched; shift, none; avalanche length, 5). As mentioned in Materials and Methods, these parameters could be varied to produce 480 different ways to analyze the data. To demonstrate that the results were robust with respect to choice of analysis parameters, each of the six parameters was varied through its full range of nonstandard values while keeping the other parameters fixed at their standard values. This relatively simple approach required only 12 different sets of parameters, rather than the computationally prohibitive 480 sets. The fraction of correlations above significance ($p < 0.02$) for avalanche length 5 was used as an index to evaluate the effects of parameter changes. The standard set of analysis parameters produced a fraction of 0.139 ± 0.105 above significance, whereas the nonstandard analysis parameters pro-

duced, on average, a fraction of 0.154 ± 0.074 . An ANOVA revealed that the 12 parameter sets had no significant differences (ANOVA: $F_{\text{crit}} = 2.17$, $F_{\text{actual}} = 1.40$, $\alpha = 0.05$, $df_{\text{groups}} = 12$, $df_{\text{error}} = 60$; two-tailed test), indicating that the results do not significantly depend on the choice of analysis parameters. The standard set of analysis parameters is therefore justified and, if anything, is more stringent than the alternative sets of analysis parameters. To show that these results are not dependent on the choice of avalanche length, the fraction of significant correlations was also calculated for avalanche lengths 4 and 6 using standard parameters. There were no significant differences between avalanche lengths 4, 5, and 6 (ANOVA: $F_{\text{crit}} = 5.46$, $F_{\text{actual}} = 0.32$, $\alpha = 0.05$, $df_{\text{groups}} = 2$, $df_{\text{error}} = 10$; two-tailed test), consistent with the idea that avalanche length 5 can be used as a representative for other avalanche lengths.

To show more carefully that the correlations did not depend significantly on the threshold chosen, similar avalanches were extracted from 1 hr of data for thresholds at 3.25, 3.75, 4.25, and 4.75 SDs below the mean. As shown in Figure 12*A*, pairs of similar avalanches could be extracted at all of these thresholds, and the spatiotemporal patterns seen in these pairs changed very little with the threshold. To quantify this for all cultures, the average correlation value obtained from a 1 hr data set using binary correlation was plotted as a function of threshold (Fig. 12*B*). As can be seen from the figure, the average correlation value is virtually constant over this range of thresholds. Selecting threshold values

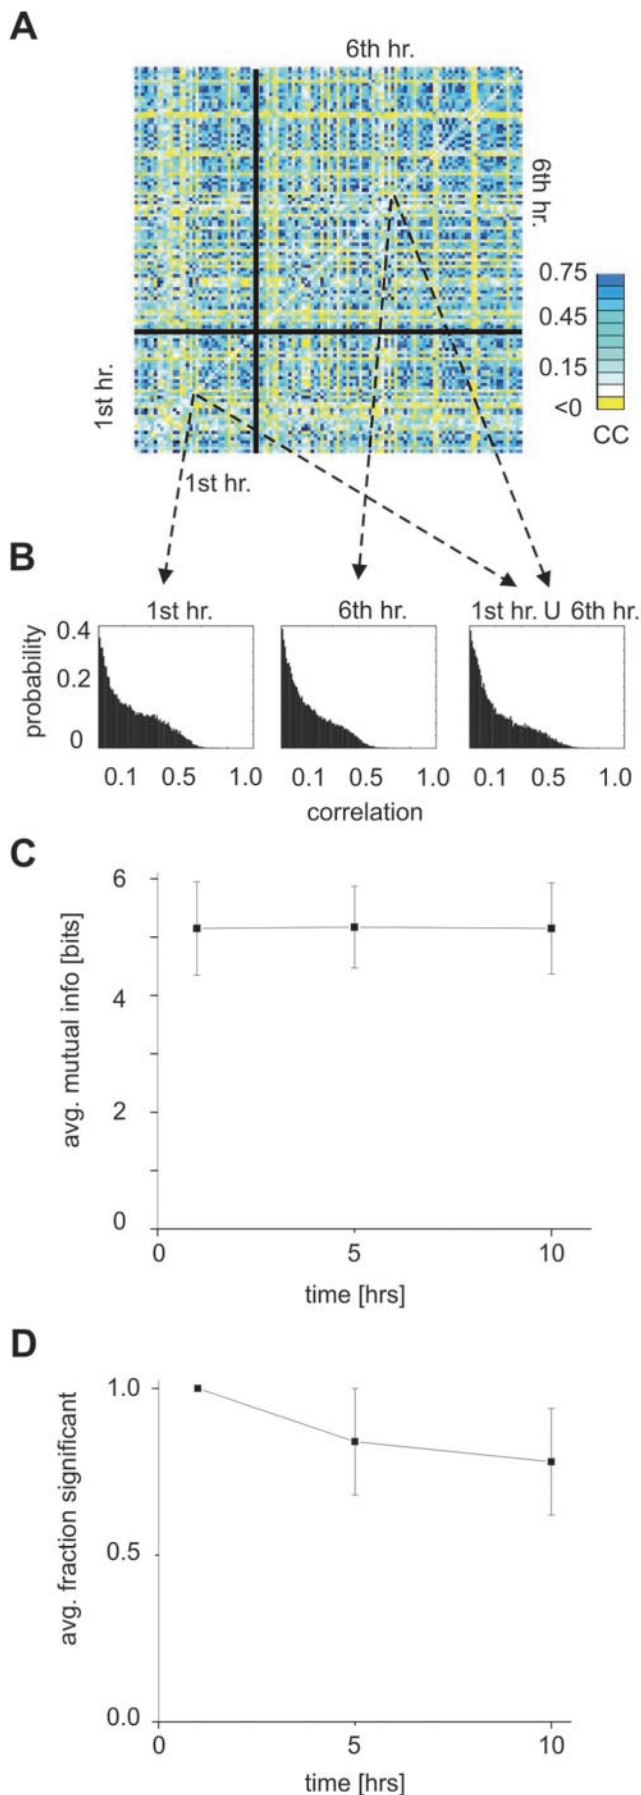
smaller than 2.75 SDs will cause changes in the correlations and in the patterns though, and this is to be expected. When the data were analyzed using ROC curves, the optimum threshold was identified as 2.86 SDs below the mean (for review, see Beggs and Plenz, 2003). At this value, the number of type I and type II errors are minimized. Thresholds smaller than 2.86 would be expected to increase type I errors by selecting waveforms that should be classified as noise.

Discussion

Recent work (Beggs and Plenz, 2003) has shown that cultured cortical networks, as well as acute cortical slices, operate near a critical point at which they spontaneously produce neuronal avalanches that may optimize information transmission. The main finding of this report is that these neuronal avalanches are highly repeatable and can be clustered into many statistically significant families of activity patterns that satisfy several requirements of a memory substrate. These patterns display stability over 10 hr, occur on average in 30 ± 14 distinct forms in each culture, and have a temporal precision of ± 4 msec.

Previous work

Advances in technology over the past 20 years have allowed cultures of cortical neurons to be grown on MEAs, often for periods



of several weeks (Pine, 1980; Gross et al., 1982; Corey et al., 1991; Potter, 2001; Segev et al., 2001). Previous experiments have shown that spontaneous activity in cultured neural networks can propagate in a wave-like manner (Maeda et al., 1995), that plasticity can be induced in such networks through electrical stimulation (Jimbo et al., 1999; Shahaf and Marom, 2001; Tateno et al., 2002; Eytan et al., 2003), and that different concentrations of calcium can produce correlations in convolved representations of network activity (Segev et al., 2004).

Neuronal avalanches represent highly diverse, yet repeatable, activity patterns different from epileptic activity

Several measures indicated that the spontaneous activity in the cortical slice cultures used here was healthy (i.e., non-epileptic). The wide variety of patterns produced, their remarkable stability over 10 hr, their multiple initiation sites, their rapid activation velocity, their non-wave-like propagation, and the fact that most avalanches only activated a subset of available electrodes all stand in sharp contrast to reports of epileptic activity. Epilepsy in the cortex can be characterized by stereotypic activity patterns with little variety (Albowitz et al., 1990), reduced activity in tissue after a seizure (Gorji et al., 2001), a single or few initiation sites (foci) (Reid and Palovcik, 1989; Biella et al., 1996), relatively slow and wave-like propagation (Chervin et al., 1988), or large depolarizations that would affect all electrodes within a recording area (Chagnac-Amitai and Connors, 1989). Because the cortical slice cultures reported here did not show any of these signs, we conclude that avalanches reflected healthy activity. In addition, we demonstrated in previous work that application of a GABA_A antagonist to these cultures was shown to produce activity patterns that propagated in a wave-like manner and activated all electrodes significantly more often than under control conditions (Beggs and Plenz, 2003). These data indicate that inhibition was intact in the untreated cortical slice cultures reported here.

Field potentials rather than spikes

Whereas it is relatively common to record spikes in cultured neurons using MEAs (Maeda et al., 1995; Kamioka et al., 1996; Jimbo et al., 2000; Segev et al., 2001; Shahaf and Marom, 2001), the analysis presented here focuses exclusively on patterns produced by LFPs. Because the sharp negative deflection of an LFP contains information about synchronized spiking in a population of neurons near the electrode (Bove et al., 1998), LFP measurements may be more representative of the state of a network than measurements obtained from single spikes.

←

Figure 9. The same correlations between avalanches are preserved over time. *A*, Merged correlation matrix from Figure 8 is shown with correlations from the first hour (bottom left box) and correlations from the sixth hour (top right box) marked out by black lines. *B*, Correlations that were significant are plotted as three distributions. Horizontal axes give correlation values, and vertical axes give probability of observing that correlation value. The three plots give the distribution of significant correlations observed in the first hour, the sixth hour, and the distribution produced by the union of the correlations produced by the first and sixth hours. Note the similar shapes of the distributions, indicating that the same patterns continued to correlate with each other over time. *C*, The average mutual information from the avalanches showed no significant decline after 10 hr. Mutual information was calculated from the entropies of the distributions of significant correlations. This graph indicates that the significant avalanches that did repeat were essentially unchanged in form over time. *D*, The fraction of significant correlations between avalanches from different hours declined slightly after 10 hr. Data are from avalanches of length 5 ($n = 5$ cultures; matched shuffling). Here, the average retention for the group was 81% after 5 hr and 73% after 10 hr. Values are represented as the fraction of what was retained after 1 hr. This graph indicates that the number of significant avalanches decreased somewhat with time.




Figure 10. Diversity of significant families and avalanches. *A*, average number of significant families of avalanches plotted as a function of avalanche length. Note general decline in significant families for longer avalanches. Nevertheless, the average culture produced a total of 30 ± 14 distinct significant families, indicating that the networks could support a diversity of significant patterns. *B*, Log of average number of avalanches found in significant families plotted as a function of avalanche length. The dashed line is best linear fit. Again, note that the number of avalanches in significant families declines with increasing avalanche length. Still, the average significant family contained 23 avalanches, indicating that the network revisited these activity states frequently.




Figure 11. Temporal precision of avalanches. Plot of average correlation ($n = 5$ cultures) against size of Gaussian jitter (see Materials and Methods). Original raster data were binned at 1 msec, and all avalanches of length 5 were extracted and used to construct correlation matrices and to obtain average correlation value. The time points within the original rasters were then jittered, and average correlation values were again extracted. Note that average correlation declines sharply with jitter, losing 49% of its original value after a jitter of 4 msec. The difference between the first two data points and the third data point was statistically significant (Tukey test; $p < 0.01$). The Tukey test was applied after a one-way ANOVA showed that there were significant differences between the points. This result demonstrates that the avalanches produced by spontaneous activity are temporally precise. Error bars indicate SDs.

Indeed, several recent reports indicate that bursts of action potentials, like the kind that could produce LFPs, as well as LFPs themselves, often contain more information about the state of a network than do single action potentials (Tsodyks et al., 1999; Butts and Rokhsar, 2001; Krahe et al., 2002; Pesaran et al., 2002). For these reasons, an analysis of LFP data in the organotypic cortex culture is expected to provide insights into intrinsic cortical network activity. The methods used here, however, need not be restricted to LFP data and could be applied to the analysis of spikes in the future.

Importance of the critical state

The critical state in cortical networks is characterized by avalanche size distributions that follow a power law and by a branching parameter that is near unity (Beggs and Plenz, 2003). Other

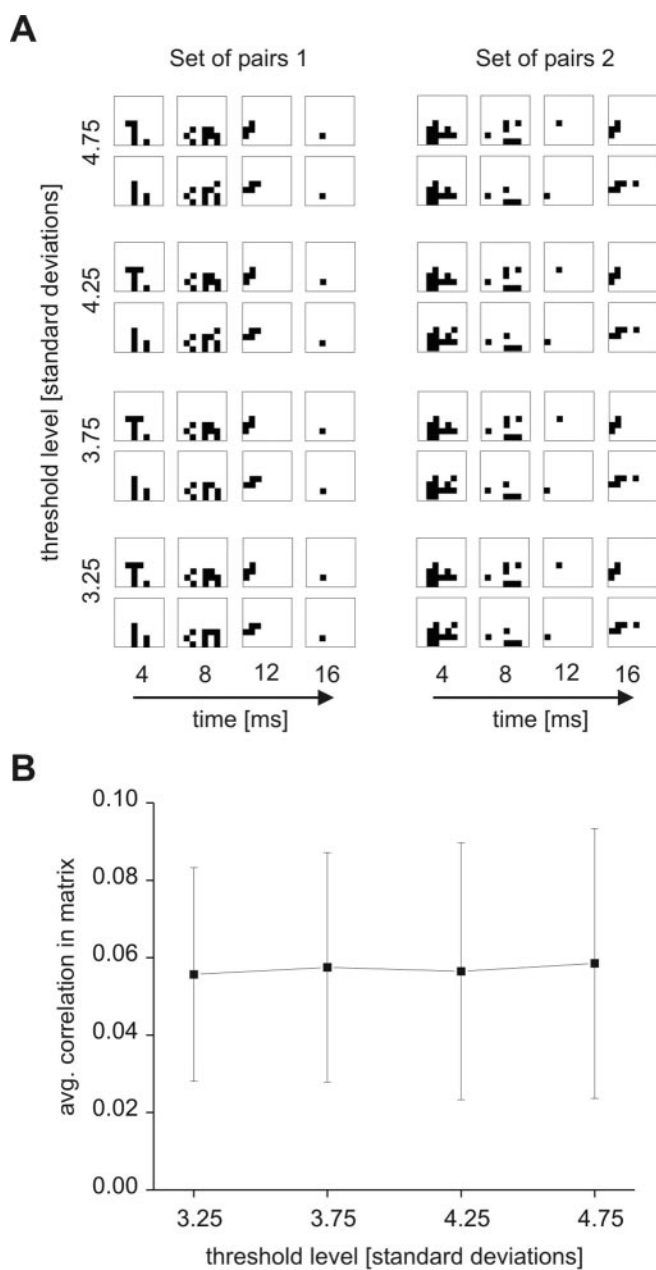


Figure 12. Patterns were robust with respect to thresholding. *A*, Pairs of avalanche patterns that were found to be similar (binary data) at different thresholds. Patterns of four frames are shown, with black pixels representing electrodes that were over threshold. A pair of similar patterns is shown at each threshold, and each column contains one set of pairs followed over different thresholds. Note that the two sets of pairs change only slightly as a function of threshold. *B*, Average correlation values for avalanches of length 4 remained constant over several thresholds. Binary correlation values for avalanches of length 4 from all cultures ($n = 5$) are shown.

physical phenomena, like nuclear chain reactions, also exhibit power law distributions of event sizes and a branching parameter near unity when they are tuned to the critical state (Harris, 1989). In neural networks, several factors indicate that the critical state is important for information processing. First, cortical slice cultures arrive at the critical state without being driven by any external synaptic input, indicating that the network becomes critical through self-organization. Second, evidence for the critical state is also seen in acute cortical slices, suggesting that this state is also operative *in vivo* (Beggs and Plenz, 2003). Third, network simu-

lations show that the critical state is optimal for information transmission through the cortical network (Beggs and Plenz, 2003). Fourth, spontaneous activity in the critical state occurred in discrete avalanches that were typically only several tens of milliseconds long and had well defined temporal starting points and ending points. These avalanches had rapid dynamics, commensurate with the fast processing times observed in the cortex (Tovee, 1994). For all of these reasons, we hypothesized that activity observed in the critical state would suggest a natural time scale at which to search for repeating activity patterns. Accordingly, the analysis presented here was all done on networks that were in the critical state and at time scales that were optimized for the resolution of neuronal avalanches.

Long-term stability

Statistically significant families of avalanches have properties that satisfy several requirements thought to be necessary for mnemonic function. Long-term stability is one such property, and two different measures showed that the significant families observed here possessed stability over several hours. The fraction of correlations between avalanches that were significant declined only 22% even after 10 hr. This indicated that 78% of the significant avalanches produced by the cortical networks after 10 hr continued to be similar to the significant avalanches produced in the first hour of data. Moreover, the distribution of correlations that these avalanches had with each other remained remarkably constant over time. Mutual information was used to measure this constancy and declined <2% after 10 hr. In other words, the significant families that existed in the first hour of recording were predictive of (i.e., they contained information about) where avalanches would exist in similarity space in later hours of recording. This long-term stability suggests that significant families could be used by cortical circuits to store information for long-term memory.

Diversity

Other laboratories have observed statistically significant repeating patterns of neural activity in behaving animals (Nadasdy et al., 1999; Dave and Margoliash, 2000; Louie and Wilson, 2001; Hoffman and McNaughton, 2002) and in an isolated cortical network (Cossart et al., 2003), yet the diversity of these repeating patterns has remained poorly understood. Early theoretical work showed that distributed storage mechanisms in artificial neural networks could allow multiple patterns to be stored in the same network (Hopfield, 1982; Cohen and Grossberg, 1983; McClelland and Rumelhart, 1988). What evidence is there that the cortical networks observed here met this expectation? Using the methods of correlation-based clustering combined with a global contrast function, these data indicate that the average culture had 30 ± 14 different statistically significant families and that most cultures produced significant families at all avalanche lengths. Because there is no standard way to solve large clustering problems, our results could depend somewhat on the quality of the clustering method chosen. For some data sets, simulated annealing was used. Inspection suggested that this method was able to extract more clusters, but because it was computationally expensive, it was not consistently used. Thus, the number of significant families given here might actually be an underestimate. Each of the statistically significant families found in our data were distinct from other families within the same culture with respect to length, number of electrodes activated, or the spatiotemporal pattern of activation that occurred. Therefore, these data clearly

suggest that the cortical networks could indeed store multiple repeating activity patterns.

Temporal precision

Significant families of avalanches were also temporally precise. The temporal precision of neural activity is a subject that has received much attention in the literature (Abeles et al., 1993; Mainen and Sejnowski, 1995; Rieke et al., 1997; Shadlen and Newsome, 1998; Diesmann et al., 1999; Oram et al., 1999; Hahnloser et al., 2002; Buonomano, 2003), and the extent to which temporal precision plays a role in information processing in central associative areas of cortex remains controversial. Although the data reported here from cortical slice cultures do not directly address the precision of single action potentials, they offer some insights to timing in a local cortical network. As mentioned before, the LFPs produced by the network are timed to within ± 4 msec. This means that a population of spiking neurons (necessary to produce a LFP) is tightly synchronized at precise times in the cultures. Moreover, this precision holds true not only for a single LFP but also for patterns of LFPs involving up to 60 electrodes. The data therefore demonstrate that a cortical network can produce complex and precise spatiotemporal patterns from synchronized populations of spiking neurons. The slice culture is not only capable of producing such complex patterns; this is what it does naturally, in the absence of any external synaptic input. These data suggest that the “default mode” of the cortex is to produce complex, precisely timed patterns. Such precision meets the expectation that stored sequences of neural activity should be temporally precise.

Conclusion

The data described in the present report show that neuronal avalanches in isolated cortical slice cultures have features that fulfill several requirements of a substrate for memory. In addition, previous work indicated that simulated networks tuned to the critical state optimized information transmission through neuronal avalanches. Taken together, these studies suggest that neuronal avalanches play an important role in cortical networks for both information transmission and information storage.

References

- Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 70:1629–1638.
- Albowitz B, Kuhnt U, Ehrenreich L (1990) Optical recording of epileptiform voltage changes in the neocortical slice. *Exp Brain Res* 81:241–256.
- Beggs JM, Plenz D (2001) Diversity and stability of elementary activity patterns that spontaneously form in isolated cortex. *Soc Neurosci Abstr* 27:64.7.
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23:11167–11177.
- Biella G, Forti M, de Curtis M (1996) Propagation of epileptiform potentials in the guinea-pig piriform cortex is sustained by associative fibres. *Epilepsy Res* 24:137–146.
- Bolz J (1994) Cortical circuitry in a dish. *Curr Opin Neurobiol* 1994 4:545–549.
- Bove M, Martinoia S, Verreschi G, Giugliano M, Grattarola M (1998) Analysis of the signals generated by networks of neurons coupled to planar arrays of microtransducers in simulated experiments. *Biosens Bioelectron* 13:601–612.
- Buonomano DV (2003) Timing of neural responses in cortical organotypic slices. *Proc Natl Acad Sci USA* 100:4897–4902.
- Butts DA, Rokhsar DS (2001) The information content of spontaneous retinal waves. *J Neurosci* 21:961–973.
- Calvet M-C (1974) Patterns of spontaneous electrical activity in tissue cultures of mammalian cerebral cortex vs. cerebellum. *Brain Res* 69:281–295.

- Caesar M, Schuz (1989) Maturation of neurons in neocortical slice cultures. A light and electron microscopic study on in situ and in vitro material. *J Hirnforschung* 33:429–443.
- Caesar M, Bonhoeffer T, Bolz J (1989) Cellular organization and development of slice cultures from rat visual cortex. *Exp Brain Res* 77:234–444.
- Chagnac-Amitai Y, Connors BW (1989) Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. *J Neurophysiol* 61:747–758.
- Chervin RD, Pierce PA, Connors BW (1988) Periodicity and directionality in the propagation of epileptiform discharges across neocortex. *J Neurophysiol* 60:1695–1713.
- Cohen MA, Grossberg S (1983) Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Trans Syst Man Cyber* 13:815–826.
- Corey JM, Wheeler BC, Brewer GJ (1991) Compliance of hippocampal neurons to patterned substrate networks. *J Neurosci Res* 30:300–307.
- Cossart R, Arvonov D, Yuste R (2003) Attractor dynamics of network UP states in the neocortex. *Nature* 423:283–288.
- Crain SM (1966) Development of “organotypic” bioelectric activity in central nervous tissues during maturation in culture. *Int Rev Neurobiol* 9:1–43.
- Dave AS, Margoliash D (2000) Song replay during sleep and computational rules for sensorimotor vocal learning. *Science* 290:812–816.
- Diesmann M, Gewaltig MO, Aertens A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Eytan D, Brenner N, Marom S (2003) Selective adaptation in networks of cortical neurons. *J Neurosci* 23:9349–9356.
- Gopal KV, Gross GW (1996) Auditory cortical neurons in vitro: cell culture and multichannel extracellular recording. *Acta Otolaryngol* 116:690–696.
- Gorji A, Scheller D, Straub H, Tegmeier F, Kohling R, Hohling JM, Tuxhorn I, Ebner A, Wolf P, Werner Panneck H, Oettel F, Speckmann EJ (2001) Spreading depression in human neocortical slices. *Brain Res* 906:74–83.
- Gotz M, Bolz J (1992) Formation and preservation of cortical layers in slice cultures. *J Neurobiol* 23:783–802.
- Gross GW, Williams AN, Lucas JH (1982) Recording of spontaneous activity with photoetched microelectrode surfaces from mouse spinal neurons in culture. *J Neurosci Methods* 5:13–22.
- Gutnick MJ, Wolfson B, Baldino F (1989) Synchronized neuronal activities in neocortical explant cultures. *Exp Brain Res* 76:131–140.
- Hahnloser RH, Kozhevnikov AA, Fee MS (2002) An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419:65–70.
- Harris TE (1989) *The theory of branching processes*. New York: Dover.
- Hoffman KL, McNaughton BL (2002) Coordinated reactivation of distributed memory traces in primate neocortex. *Science* 297:2070–2073.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554–2558.
- Jimbo Y, Robinson HP (2000) Propagation of spontaneous synchronized activity in cortical slice cultures recorded by planar electrode arrays. *Bioelectrochemistry* 51:107–115.
- Jimbo Y, Tateno T, Robinson HP (1999) Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J* 76:670–678.
- Jimbo Y, Kawana A, Parodi P, Torre V (2000) The dynamics of a neuronal culture of dissociated cortical neurons of neonatal rats. *Biol Cybern* 83:1–20.
- Kamioka H, Maeda E, Jimbo Y, Robinson HPC, Kawana A (1996) Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neurosci Lett* 206:109–112.
- Karpiak V, Plenz D (2002) Preparation and maintenance of organotypic cultures for multielectrode array recordings. In: *Current protocols in neuroscience* (Taylor GP, ed), pp 6.15.1–6.15.5. New York: Wiley.
- Krahe R, Kreiman G, Gabbiani F, Koch C, Metzner W (2002) Stimulus encoding and feature extraction by multiple sensory neurons. *J Neurosci* 22:2374–2382.
- Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29:145–156.
- Maeda E, Robinson HPC, Kawana A (1995) The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J Neurosci* 15:6834–6845.
- Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506.
- McClelland JL, Rumelhart DE (1988) *Explorations in parallel distributed processing*. Cambridge, MA: MIT.
- Nadasdy Z, Hirase H, Czurko A, Csicsvari J, Buzsáki G (1999) Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* 19:9497–9507.
- 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.
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811.
- Pine J (1980) Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J Neurosci Methods* 2:19–31.
- Plenz D, Aertens A (1996a) Neuronal dynamics in cortex-striatum cocultures. I. Anatomy and electrophysiology of single neuron types. *Neuroscience* 70:861–891.
- Plenz D, Aertens A (1996b) Neuronal dynamics in cortex-striatum cocultures. II. Spatio-temporal characteristics of neuronal activity. *Neuroscience* 70:893–924.
- Plenz D, Kitai ST (1996) Generation of high-frequency oscillations in local circuits of rat somatosensory cortex cultures. *J Neurophysiol* 76:4180–4184.
- Plenz D, Kitai ST (1998) “Up” and “down” states in striatal medium spiny neurons simultaneously recorded with spontaneous activity in fast-spiking interneurons studied in cortex-striatum-substantia nigra organotypic cultures. *J Neurosci* 18:266–283.
- Potter SM (2001) Distributed processing in cultured neuronal networks. *Prog Brain Res* 130:49–62.
- Reid SA, Palovcik RA (1989) Spontaneous epileptiform discharges in isolated human cortical slices from epileptic patients. *Neurosci Lett* 98:200–204.
- Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W (1997) *Spikes: exploring the neural code*. Cambridge, MA: MIT.
- Segev R, Shapira Y, Benveniste M, Ben-Jacob E (2001) Observations and modeling of synchronized bursting in two-dimensional neural networks. *Phys Rev E* 64:011920.
- Segev R, Baruchi I, Hulata E, Ben-Jacob E (2004) Hidden neuronal correlations in cultured networks. *Phys Rev Lett* 92:118102.
- Shannon CE, Weaver W (1949) *The mathematical theory of communication*. Urbana, IL: University of Illinois.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Shahaf G, Marom S (2001) Learning in networks of cortical neurons. *J Neurosci* 21:8782–8788.
- Tateno T, Kawana A, Jimbo Y (2002) Analytical characterization of spontaneous firing in networks of developing rat cultured cortical neurons. *Phys Rev E* 65:051924.
- Tovee MJ (1994) Neuronal processing. How fast is the speed of thought? *Curr Biol* 4:1125–1127.
- Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946.
- Weliky M, Kandler K, Fitzpatrick D, Katz LC (1995) Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron* 15:541–552.