Hierarchical Interaction Structure of Neural Activities in Cortical Slice Cultures

Gustavo S. Santos,¹ Elakkt D. Gireesh,² Dietmar Plenz,² and Hiroyuki Nakahara^{1,3}

¹Laboratory for Integrated Theoretical Neuroscience, RIKEN Brain Science Institute, Wako, Saitama 351-0198, Japan, ²Section on Critical Brain Dynamics, Laboratory of Systems Neuroscience, National Institute of Mental Health, Bethesda, Maryland 20892, and ³Department of Computational Intelligence and Systems Science, Tokyo Institute of Technology, Yokohama 226-8501, Japan

Recent advances in the analysis of neuronal activities suggest that the instantaneous activity patterns can be mostly explained by considering only first-order and pairwise interactions between recorded elements, i.e., action potentials or local field potentials (LFP), and do not require higher-than-pairwise-order interactions. If generally applicable, this pairwise approach greatly simplifies the description of network interactions. However, an important question remains: are the recorded elements the units of interaction that best describe neuronal activity patterns? To explore this, we recorded spontaneous LFP peak activities in cortical organotypic cultures using planar, integrated 60-microelectrode arrays. We compared predictions obtained using a pairwise approach with those using a hierarchical approach that uses two different spatial units for describing the activity interactions: single electrodes and electrode clusters. In this hierarchical model, short-range interactions within each cluster were modeled by pairwise interactions of electrode activities and long-range interactions were modeled by pairwise interactions of cluster activities. Despite the relatively low number of parameters used, the hierarchical model provided a more accurate description of the activity patterns than the pairwise model when applied to ensembles of 10 electrodes. Furthermore, the hierarchical model was successfully applied to a larger-scale data of ~60 electrodes. Electrode activities within clusters were highly correlated and spatially contiguous. In contrast, long-range interactions were diffuse, suggesting the presence of higher-than-pairwise-order interactions involved in the LFP peak activities. Thus, the identification of appropriate units of interaction may allow for the successful characterization of neuronal activities in large-scale networks.

Introduction

Multineuron firing patterns constitute a basis of neural coding, and different proposals have been made concerning their functions (Abeles et al., 1993; Aertsen and Arndt, 1993; Singer and Gray, 1995; Richmond and Gawne, 1998; Salinas and Sejnowski, 2001; Pillow et al., 2008). Recent studies in the retina (Schneidman et al., 2006; Shlens et al., 2006) and in cortical cultures (Tang et al., 2008) suggested that the probabilities of multineuron spike activities [or multielectrode local field potential (LFP) peaks] can be mostly predicted by pairwise interactions between binary activities of elements (neurons or electrodes) and, consequently, that higher-than-pairwise neuronal interactions can be effectively ignored. In principle, a network with N elements can have multiple orders of interactions, from second-order (pairwise), third-order, up to N th-order interactions (Martignoni et al., 2000; Amari, 2001; Nakahara and Amari, 2002; Montani et al., 2009). Therefore, if generally applicable, the so-called pairwise

model in the above-mentioned studies significantly simplifies the description of instantaneous network interactions (Nirenberg and Victor, 2007). Several important issues, however, merit further examination.

First, there is the scalability issue of the pairwise model. The pairwise model was previously applied to subsets of ~10 elements [a notable exception being Shlens et al. (2009)], primarily for the following two reasons: (1) the model's performance criterion was not applicable to larger-scale data due to the limitation of the number of samples (Paninski, 2003). This was addressed in a recent study (Shlens et al., 2009), in which the evaluation criterion was modified for larger-scale data (this modification was adopted in the present study); and (2) direct estimation of the pairwise model becomes computationally prohibitive as the number of elements increases. The model's performance in small subsets does not necessarily extrapolate to a larger set (Roudi et al., 2009); therefore, a model with better scalability is desirable.

Second, the pairwise model assumes that any activity pattern is explained by the cumulative effect of pairwise interactions between elements. It may be the case, however, that not all pairwise interactions are always necessary (Shlens et al., 2006, 2009). Here, we examined whether long-range interactions need to be represented by every pair of elements. More generally stated: can other simplified models adequately capture the interaction structure?

Third, closely related to this issue, we considered another fundamental question that is often overlooked: what are the appropriate units for describing network interactions? The units are

Received Dec. 3, 2009; revised Feb. 16, 2010; accepted May 10, 2010.

This work was supported by KAKENHI Grant 21300129 (Japan Society for Promotion of Science) and KAKENHI Grant 20020034 (Ministry of Education, Culture, Sports, Science and Technology). We thank T. Higashi and M. Arisaka for assistance with programming and analysis, Dr. T. Ozaki for discussion in the early stage of this work, Drs. S. Amari, S. Yu, H. Hirase, T. Fukui, J. Teramae, and H. Shimazaki for discussion and helpful comments on an early draft of this manuscript, and two anonymous reviewers for their helpful comments.

Correspondence should be addressed to Hiroyuki Nakahara, Laboratory for Integrated Theoretical Neuroscience, RIKEN Brain Science Institute, Hirosawa 2-1, Wako, Saitama 351-0198, Japan. E-mail: hiro@brain.riken.jp.

DOI:10.1523/JNEUROSCI.6141-09.2010

Copyright © 2010 the authors 0270-6474/10/308720-14\$15.00/0

usually defined as the particular elements recorded in a given experiment (e.g., single neurons or LFP activity at electrodes). The recorded elements, however, may not be the only units. Depending on the underlying interaction structure, there may be other appropriate units that allow for a better model.

To address these issues, we constructed a hierarchical model of network interactions that uses units of interactions at two spatial levels. We used this model to examine the spontaneous negative LFP (nLFP) peak activities of cortical cultures. At a larger spatial level, the unit of interaction is the “cluster activity” (determined by nLFP peak activities among electrodes in each cluster), and long-range interactions are modeled as pairwise interactions between cluster activities. At a local level (within clusters), the unit of interaction is the nLFP peak activity at each electrode, and short-range interactions are modeled as pairwise interactions between electrode activities. Our findings indicated that, despite its relative parsimony, the hierarchical model captures network interactions more accurately than the pairwise model.

Materials and Methods

Recordings. Details of the recordings were provided previously (Beggs and Plenz, 2003, 2004; Gireesh and Plenz, 2008). Briefly, coronal slices from the somatosensory cortex (350 μm thick) and ventral tegmental area (VTA; 500 μm thick) were obtained from rat pups (postnatal day 0–2, Sprague Dawley; Taconic Farms) and placed on poly-D-lysine-coated, 8×8 microelectrode arrays (MEA, 30 μm electrode diameter; 200 μm interelectrode distance, Multi Channel Systems). Each MEA covered ~ 1 mm of the cortical depth of a coronal slice and a VTA slice was positioned at the ventral border of the cortex, juxtaposing the white matter. Coculturing of slices was described in detail previously (Plenz and Kitai, 1998; Karpik and Plenz, 2002). In short, a sterile chamber was attached to the MEA, which allowed for cultivation and repeated recording from single cocultures for many weeks. After plasma/thrombin-based adhesion of the tissue to the MEA, 600 μl of standard culture medium was added (50% basal medium, 25% HBSS, 25% horse serum; all from Sigma-Aldrich) and the MEAs were affixed to a rocking tray inside the incubator ($\pm 70^\circ$ angle, 0.005 Hz; $35.5 \pm 0.5^\circ\text{C}$).

For recordings, individual MEAs were attached to a recording head stage (Multi Channel Systems) affixed to a second tray within the incubator. This configuration allowed for repeated long-term recordings from individual cultures in culture medium under conditions identical to those during development. Recordings began after 9 d *in vitro* (DIV), when neuronal avalanches were found (Gireesh and Plenz, 2008; Stewart and Plenz, 2008). Spontaneous extracellular activity (LFP) was sampled at 4 kHz per channel for an average of 3–5 h per condition (MC-Rack; Multi Channel Systems) against the ground electrode incorporated in the MEA, which acted as the common signal ground for all electrodes. Light microscopic pictures taken at 1 DIV and during the first and second weeks of cultivation enabled reconstruction of the electrode locations in the cortical culture. Data from 12 cultures (10–16 DIV) were used in this study.

Extracting nLFP peak activity. Recordings were low-pass filtered at 200 Hz with a second-order Butterworth filter for the extraction of LFPs, which exhibited sharp negative peaks indicative of population spikes (Beggs and Plenz, 2003). The threshold was set at -8 SDs of the signal to detect significant negative peaks, and the time of the maximum excursion below the threshold was recorded as the time of the nLFP peak activity (see Fig. 1A), as previously described (Beggs and Plenz, 2003, 2004). nLFP peaks were binned at a fixed interval, counting at most one peak per electrode per bin and discarding the amplitude information of the peaks, as done in a previous application of the pairwise model to LFP activities in cortical cultures (Tang et al., 2008). Instantaneous activity was defined as the microelectrode binary activity occurring in a single time bin on the array (Fig. 1C). Three different bin sizes were examined (4, 10, and 20 ms). The choice of bin sizes was justified as follows: 4 ms was the average interval time between successive nLFPs on the array at an interelectrode distance of 200 μm (Beggs and Plenz, 2003), and 10 and 20

ms were the bin sizes used in previous studies of network correlations (Schneidman et al., 2006; Shlens et al., 2006; Tang et al., 2008).

The pattern of the nLFP peak activity in a bin is represented as a random binary vector, $\mathbf{X} = (x_1, x_2, \dots, x_n)$, where for each electrode i , the corresponding element x_i has a value of 1 or 0, indicating the presence or absence of an nLFP peak, respectively. The underlying true probability of an activity pattern $\mathbf{X}$ is denoted as $P(\mathbf{X})$; the true probability is not known. We can only empirically observe samples in our datasets, and we denote the empirical probability observed in the data by $\hat{P}(\mathbf{X})$.

Models of instantaneous activity. In this study, we examined three models of instantaneous activity: an independent model, a pairwise model, and a hierarchical model. The first two models belong to a class of maximum entropy models, and have been applied to datasets similar to ours (Tang et al., 2008). The focus of the present study was to introduce the third model, the hierarchical model, and to examine its validity using our datasets. Indeed, the model can also be cast onto the framework of the maximum entropy model, because the maximum entropy formulation can also be understood through the intrinsic relation between dual information geometric coordinates; that is, the relation between expectation and natural (log-linear) coordinates for probability distributions of discrete variables (Amari, 2001, 2003; Nakahara and Amari, 2002) (see Appendix). In the following sections, we summarize the first two models as maximum entropy models and then separately describe the third model.

Independent and pairwise models as maximum entropy models. “Maximum entropy” refers to the criterion for selecting a probability distribution $Q(\mathbf{X})$ given certain constraints (Jaynes, 1957), namely the distribution with the highest Shannon entropy, $H(Q) = -\sum_{\mathbf{X}} Q(\mathbf{X}) \log[Q(\mathbf{X})]$, among all candidate distributions $Q(\mathbf{X})$ satisfying the imposed constraints. Generally speaking, maximum entropy models refer to those chosen by the above criterion. In the current study, the random variable of interest is a random binary vector $\mathbf{X} = (x_1, x_2, \dots, x_n)$. Accordingly, the so-called “pairwise maximum entropy model” is related to the Ising model (from statistical mechanics) with pairwise interaction terms, and the “independent maximum entropy model” corresponds to that without any interaction terms but only first-order terms. There have been a tremendous number of studies, especially in computational neuroscience, using potential relations or applications of pairwise models to explain functions of both artificial and biological neural networks, e.g. (Amari, 1977; Hopfield, 1982; Hopfield and Tank, 1986; Hertz et al., 1991). More recently, the pairwise model was shown to be capable of characterizing most neural interactions and correlations in various tissues and preparations (Schneidman et al., 2006; Shlens et al., 2006; Tang et al., 2008).

The pairwise model is expressed as

$$\log Q_{\text{pair}}(\mathbf{X}) = \sum_i \theta_i x_i + \sum_{i,j} \theta_{ij} x_i x_j - \psi, \quad (1)$$

where ψ is a normalization factor, θ_i s are the model parameters relating to the first-order term, and θ_{ij} s are the second-order or pairwise interaction terms (Amari, 2001, 2003; Nakahara and Amari, 2002). The independent model is expressed as

$$\log Q_{\text{ind}}(\mathbf{X}) = \sum_i \theta_i x_i - \psi, \quad (2)$$

and contains only the first-order term. It can also be expressed as

$$Q_{\text{ind}}(\mathbf{X}) = \prod_{i=1}^n \hat{P}(x_i), \text{ where } \hat{P}(x_i) \text{ is the marginal probability of activity at electrode } i.$$

Both models are convenient in that their model parameters can be estimated uniquely under the maximum entropy criterion (Jaynes, 1957). Either estimated model can be expressed under the maximum entropy criterion as follows:

$$Q_{\text{model}}(\mathbf{X}) = \arg \max_{Q(\mathbf{X})} \left[H(Q) + \sum_{l=1}^N \theta_l (E_Q[F_l(\mathbf{X})] - k_l) \right], \quad (3)$$

where the constraints are specified using the Lagrange multipliers θ_l with the linear functions $F_l(\mathbf{X})$ (i.e., x_i for the first-order and $x_i x_j$ for the second-order terms), and the corresponding expected values k_l [i.e., first-

order $E_p(x_i) = \hat{P}(x_i)$, and second-order $E_p(x_i x_j) = \hat{P}(x_i x_j)$. $Q_{\text{pair}}(\mathbf{X})$ uses both first and second-order terms, but $Q_{\text{ind}}(\mathbf{X})$ uses only the first-order terms. To specify and solve this convex optimization problem, we used CVX, a software package for specifying and solving convex programs (Grant and Boyd, 2008, 2009). All analyses were performed in Matlab (MathWorks). It should be noted that the number of parameters in the maximum entropy model is equal to the number of nonredundant constraints placed on this optimization problem.

Structure of hierarchical model. A brief summary of the hierarchical model's structure is provided here to facilitate understanding of the subsequent subsections (see Results for details about model structure). The hierarchical model assumes that the underlying interaction structure of $P(\mathbf{X})$ is characterized using units for interactions at two levels: the electrode cluster level and the single electrode level within a cluster. At the cluster level, the interactions are assumed to be characterized by pairwise interactions of cluster activities. At the electrode level, interactions within a cluster are assumed to be pairwise and interactions across clusters are assumed to be independent, given the activity of each cluster. Together, these assumptions allow for a mathematical expression of the hierarchical model shown in Equation 10.

Cluster activity. To use the hierarchical model described above, we must first define cluster activity. Cluster activity is the output of a cluster; similar to electrode activity, it can be represented in different ways (e.g., as a real-value or thresholded-binary-value variable). We set the minimum property of cluster activity as a scalar measure monotonically reflecting the level of activity among its member electrodes. We adopted three measures of cluster activity: linear, log, and binary. If we assume that cluster I contains n electrodes, then $\mathbf{X}_I = (x_1, x_2, \dots, x_n)$. The linear

measure is a count of the active electrodes, expressed by $c_I^{\text{lin}} = \sum_{i=1}^n x_i$; the log measure takes a logarithmic transformation of the linear measure, $c_I^{\text{log}} = [\log_2(1 + c_I^{\text{lin}})]$, where the brackets indicate the floor function; and the binary measure indicates whether any electrode in the cluster is active, $c_I^{\text{bin}} = 1(c_I^{\text{lin}} > 0)$, where $1()$ is an indicator function that returns 1 if the proposition inside the bracket is true and 0 if it is false. For example, a cluster with 10 active electrodes has a corresponding cluster activity given by $c_I^{\text{lin}} = 10$, $c_I^{\text{log}} = 3$, and $c_I^{\text{bin}} = 1$. Henceforth, references to the binary/log/linear measures of cluster activity are shortened to binary/log/linear (cluster) activity. The probabilities of cluster activities are denoted and estimated in the same way as for electrodes. For example, the probability that cluster I has log activity level y is given by $P(c_I^{\text{log}} = y)$, the probability of simultaneous activities in clusters I and J is given by $P(c_I^{\text{log}} = y, c_J^{\text{log}} = z)$, and the probability distribution of array-wide activity patterns is given by $P(\mathbf{C})$, where $\mathbf{C} = (c_1, c_2, \dots, c_K)$ is the set of K clusters spanning a given electrode array.

There are advantages and disadvantages associated with each type of cluster activity. The linear cluster activity most faithfully reflects the overall number of active electrodes in a cluster, but is computationally the most expensive. The hierarchical model with linear cluster activity has a number of parameters comparable to that of the pairwise model, which for n electrodes is given by $n + \frac{n(n-1)}{2} = \frac{n(n+1)}{2}$. Hence, the hierarchical model using linear activity offers little gain in dimensionality over the pairwise model. However, it does contain higher-order interaction terms that are nonexistent in the pairwise model, while imposing further constraints on some pairwise interactions (see Appendix). By contrast, the binary activity is the most parsimonious and computationally least expensive, as it represents only the existence or nonexistence of any active electrode, and thus provides a minimum amount of information about the underlying electrodes' activities. In principle, the hierarchical model with binary activity has a significantly lower dimensionality than the pairwise model; its structural dimensionality is given by

$$\sum_i \frac{n_i(n_i + 1)}{2} + \frac{K(K-1)}{2} + \sum_i 1(n_i > 2), \quad (4)$$

where we assume K clusters, each of which has n_i electrodes ($\sum_i n_i = n$). The hierarchical model with log activity is intermediate in terms of both computational simplicity and representational capacity. In the present study, we used linear activity only to identify clusters when we wished to use as much of the available information as possible. We used either binary or log cluster activity when evaluating the hierarchical model's performance with the identified clusters, because either was sufficient for the model's high performance and use of either significantly decreases the size of the model compared with the pairwise model.

Estimation of the hierarchical model's parameters. Once the clusters are identified (as described in the next section), the hierarchical model formulation (Eq. 10) decomposes the original $P(\mathbf{X})$ -estimation problem into two subproblems (electrode and cluster level) (see Fig. 2C), as follows: (1) estimating the electrode activity distribution $P(\mathbf{X}_I)$ in each cluster I , which is approximated by $Q_{\text{pair}}(\mathbf{X}_I)$ in this formulation; and (2) estimating the cluster activity distribution $P(\mathbf{C})$ over the array, which is approximated by $Q_{\text{pair}}(\mathbf{C})$ in the formulation. The estimation methods for the pairwise model (as described in Independent and pairwise models as maximum entropy models, above) are applicable to each subproblem. After solving these subproblems, the hierarchical model $Q_{\text{hier}}(\mathbf{X})$ is obtained using the conditional independent assumption, as given in Equation

10: $Q_{\text{hier}}(\mathbf{X}) = Q_{\text{pair}}(\mathbf{C}) \prod_{I=1}^K Q_{\text{pair}}(\mathbf{X}_I | c_I)$. Note that $Q_{\text{pair}}(\mathbf{X}_I | c_I) = Q_{\text{pair}}(\mathbf{X}_I) / P(c_I)$, where $P(c_I)$ is marginalized from $Q_{\text{pair}}(\mathbf{X}_I, c_I)$; it can, in practice, be effectively replaced by $\hat{P}(c_I)$, as long as $Q_{\text{pair}}(\mathbf{X}_I)$ fits reasonably well with $\hat{P}(\mathbf{X}_I)$.

In the case of log or linear cluster activity, the estimation of the pairwise model $Q_{\text{pair}}(\mathbf{C})$ becomes an instance of a maximum entropy problem with $m > 2$ discrete-alphabet sizes (Amari, 2001; Ince et al., 2009). The basic framework remains the same, except that the constraints now depend on the probabilities of cluster activities c_I taking multiple values. The set of first- and second-order probabilities must include combinations of all nonzero levels of cluster activity. For example, given two clusters $\mathbf{X}_1 = (x_1, x_2)$ and $\mathbf{X}_2 = (x_3)$ and using linear activity c_I^{lin} , both the first-order constraints $\hat{P}(c_1^{\text{lin}} = 1)$, $\hat{P}(c_1^{\text{lin}} = 2)$, and $\hat{P}(c_2^{\text{lin}} = 1)$, and the second-order constraints $\hat{P}(c_1^{\text{lin}} = 1, c_2^{\text{lin}} = 1)$, and $\hat{P}(c_1^{\text{lin}} = 2, c_2^{\text{lin}} = 1)$ should be applied. The maximum entropy problem still follows the general form of Equation 3, but includes the m -ary alphabet constraints as explained above. After solving the maximum entropy problem and obtaining the Lagrange multipliers, $Q_{\text{pair}}(\mathbf{C})$ is given by the following:

$$\log Q_{\text{pair}}(\mathbf{C}) = \sum_I \sum_k \Theta_I^k 1(c_I = k) + \sum_{I,J} \sum_{k,l} \Theta_{IJ}^{kl} 1(c_I = k, c_J = l) - \psi, \quad (5)$$

where the superscripts of the Lagrange multipliers Θ_I^k and Θ_{IJ}^{kl} indicate the values of cluster activity in each cluster I and J .

Clustering criterion and clustering algorithm. We identified clusters of electrodes using activities of all electrodes on each array. We selected what we considered the simplest clustering criterion for identifying the clusters: the homogeneity criterion. In general, given a cluster activity c_I , there can be multiple underlying active electrode patterns $\mathbf{X}_I$. Thus, if a cluster was determined arbitrarily, the probability of the cluster activity, $P(c_I)$, would not uniquely determine the probabilities of those underlying multiple electrode probabilities $P(\mathbf{X}_I)$. If the underlying probabilities $P(\mathbf{X}_I)$ are homogeneous, however, it is possible to determine the probabilities $P(\mathbf{X}_I)$ given $P(c_I)$. With no additional assumptions, the maximum entropy estimate of $P(\mathbf{X}_I)$, given the knowledge of $P(c_I)$, is the homogeneous distribution of electrode activities (see Fig. 2D). The homogeneous distribution $Q_{\text{hom}}(\mathbf{X}_I)$, given linear cluster activity c_I^{lin} , is theoretically provided by

$$Q_{\text{hom}}(\mathbf{X}_I) = \hat{P}\left(c_I^{\text{lin}} = \sum_i x_i\right) \times \left(\sum_i x_i\right)^{-1}, \quad (6)$$

where $\hat{P}(c_I^{\text{lin}})$ is the observed probability of linear cluster activity, the parentheses on the right indicate the binomial coefficient (or the choose

function), and all electrode patterns X_I of a given cluster activity c_I^{lin} are equally likely to occur. To quantify the level of homogeneity in a cluster I relative to its estimated distribution $Q_{\text{hom}}(X_I)$, we used a criterion called cluster score (see below). A perfectly homogeneous cluster has a score of 1, meaning that $\hat{P}(X_I)$ can be perfectly estimated from $\hat{P}(c_I^{\text{lin}})$, and deviations from homogeneity decrease the cluster score.

To minimize the loss of information in describing $P(X_I)$ using $P(c_I)$, we searched for clusters that were as homogeneous as possible. Clusters of electrodes were determined using an iterative clustering algorithm with the cluster score. The algorithm uses a backward search procedure, as it starts with the set of all electrodes and iteratively prunes one electrode at a time in the order that produces the greatest increase in the cluster score of the remaining electrodes. The electrode set was pruned until it contained 10 or fewer electrodes and its cluster score was above a predetermined threshold. Upon reaching the homogeneity threshold, the electrode set was declared to be a cluster. To be conservative, we fixed the maximum number of electrodes in a cluster at 10, mostly because the numerical calculation for fitting a pairwise model to a cluster becomes harder for larger sets. After finding a cluster, a new search was started using all the previously pruned electrodes, until all electrodes were assigned to a cluster. Sometimes one or two electrodes in a dataset remained unassigned to a cluster when the homogeneity threshold was not met; such electrodes were discarded from further analysis. Note that this clustering approach defines nonoverlapping clusters. The clustering threshold was predetermined for all datasets of a given bin size by obtaining a distribution of cluster scores from randomly sampled sets of 10 electrodes pooled from all datasets (1000 samples per dataset), and setting the threshold at the upper 5% significance level. The threshold values obtained were 0.88, 0.90, and 0.90 for datasets binned at 4, 10, and 20 ms, respectively.

Evaluating the cluster-level interactions using the electrode-level, pairwise interactions. For binary cluster activity, there is only one coefficient of interaction, Θ_{IJ} , between clusters I and J in Equation 5; the coefficient indicates the level of interaction between the two clusters. For log cluster activity, there is one coefficient Θ_{IJ}^{kl} for each combination of cluster activities—cluster I with activity k , and cluster J with activity l . To represent the interaction for each cluster pair by a single value, their weighted average, $\Theta'_{IJ} \equiv \left(\frac{\Theta_{IJ}^{kl}}{kl} \right)$, was obtained by averaging over the multiple cases of different activities (i.e., k and l), and was used for Figure 7C,D. For comparison, we obtained electrode interactions θ_{ij} of Equation 1 from the pairwise models of the random 10-electrode ensembles. For each electrode pair i and j , the interaction θ_{ij} was averaged over all ensembles containing both electrodes. Given clusters I and J with n_I and n_J electrodes, respectively, there are $n_I \times n_J$ interaction coefficients θ_{ij} for electrodes i in cluster I and j in cluster J . For each cluster pair I and J , we obtained both the mean, $\theta_{IJ,\text{mean}}$, and max, $\theta_{IJ,\text{max}}$, of the cross-cluster pairwise electrode interactions.

Model evaluation. Depending on the number of electrodes in the total electrode set X , we used two different methods to evaluate the models. The first method was used for small groups of electrodes ($n = 10$), for which the distribution $\hat{P}(X)$ could be reliably sampled from the data (Schneidman et al., 2003). This method evaluated the model distributions $Q(X)$ against the observed distribution $\hat{P}(X)$ (via the Kullback–Leibler divergence), and the value was normalized by the divergence of the independent model $Q_{\text{ind}}(X)$; the resulting F ratio was defined by:

$$F = 1 - \frac{D_{\text{KL}}[\hat{P}(X) \| Q(X)]}{D_{\text{KL}}[\hat{P}(X) \| Q_{\text{ind}}(X)]} \quad (7)$$

where $D_{\text{KL}}(\hat{P} \| Q) = \sum_X \hat{P}(X) \log_2 \frac{\hat{P}(X)}{Q(X)}$. This F ratio was originally proposed to evaluate the accuracy of the pairwise model $Q_{\text{pair}}(X)$ by setting $Q(X) = Q_{\text{pair}}(X)$ in Equation 7 (Schneidman et al., 2003, 2006), denoted here as F_{pair} . F_{pair} is strictly contained within the interval between 0 and 1 because the generalized Pythagorean theorem holds (Amari, 2001), namely, $D_{\text{KL}}[\hat{P}(X) \| Q_{\text{ind}}(X)] = D_{\text{KL}}[\hat{P}(X) \| Q_{\text{pair}}(X)] + D_{\text{KL}}[Q_{\text{pair}}(X) \| Q_{\text{ind}}(X)]$. Thus, F_{pair} can be viewed as representing the accuracy of the model $Q_{\text{pair}}(X)$ using its departure $D_{\text{KL}}[\hat{P}(X) \| Q_{\text{pair}}(X)]$

in relation to, or normalized by, the departure of the independent model $D_{\text{KL}}[\hat{P}(X) \| Q_{\text{ind}}(X)]$. An F_{pair} of 0 indicates that the pairwise model does not add any additional meaningful information to the independent model, whereas an F_{pair} of 1 indicates that the pairwise model is identical to the empirical distribution. Thus, as F_{pair} approaches 1, the pairwise model is considered to sufficiently capture the empirical distribution (Schneidman et al., 2006). To comparably evaluate the hierarchical model, we set $Q(X) = Q_{\text{hier}}(X)$ in Equation 7 and called the resulting F ratio F_{hier} . Note, because the generalized Pythagorean theorem no longer holds, F_{hier} is no longer strictly bounded at the lower limit by 0. Still, it represents the normalized ratio of the departure of the model in reference to that of the independent model, so it conveniently represents a degree of accuracy of the model for comparison with the pairwise model. In practice, we found no case in which F_{hier} was not within the $[0, 1]$ range. The cluster score was devised in the same manner; we set $Q(X) = Q_{\text{hom}}(X)$ in Equation 7, for a given set of electrode activities within a cluster I , thus given by the following:

$$\text{score} = 1 - \frac{D_{\text{KL}}[\hat{P}(X_I) \| Q_{\text{hom}}(X_I)]}{D_{\text{KL}}[\hat{P}(X_I) \| Q_{\text{ind}}(X_I)]} \quad (8)$$

Note that the homogeneous probability $Q_{\text{hom}}(X)$ is easy to estimate (see Eq. 6), even when the cluster is a full array. This is an advantage of using the homogeneity criterion to identify clusters.

The method explained above (Eq. 7) was previously used only when relatively small subsets of elements (spike or LFP activities; typically ~ 10) were evaluated. A second method of evaluating the models was used for entire electrode arrays (Shlens et al., 2009). This was because directly evaluating F_{pair} with a larger number of elements is contaminated by a strong bias that is imposed by limited sampling in the estimation of the divergence between empirical and model distributions (Nemenman et al., 2002; Paninski, 2003). When using the entire array, $P(X)$ can no longer be sampled adequately, thereby precluding the direct use of the original F ratio for entire electrode arrays. Thus, we followed the method introduced by Shlens et al. (2009), which consists of marginalizing the statistics of both the data and the model before estimating the divergences; the probabilities in Equation 7 are replaced by the probabilities of the fraction of electrodes active in a time bin. This is called a modified F ratio, or F^c , and can be expressed by:

$$F^c = 1 - \frac{D_{\text{KL}}[\hat{P}(f_e) \| \hat{Q}(f_e)]}{D_{\text{KL}}[\hat{P}(f_e) \| \hat{Q}_{\text{ind}}(f_e)]} \quad (9)$$

where f_e indicates the fraction of active electrodes of the total number of electrodes and $\hat{P}(f_e)$ is the empirical distribution of f_e . Note that $\hat{P}(f_e)$ can be estimated adequately even when $\hat{P}(X)$ cannot. $\hat{Q}(f_e)$ and $\hat{Q}_{\text{ind}}(f_e)$ are the corresponding marginal distributions for the models $Q(X)$ and $Q_{\text{ind}}(X)$, respectively, from which we generated samples in simulations to estimate the marginal distributions. To evaluate the cluster-level activity in a large-scale dataset, we devised an F ratio similar to F^c , referred to as F^c . It replaces f_e with f_c , which indicates the fraction of active clusters (i.e., clusters that have at least one active electrode) in Equation 9. Likewise, the empirical distribution $\hat{P}(f_c)$ can be reliably sampled from the data, and $\hat{Q}(f_c)$ and $\hat{Q}_{\text{ind}}(f_c)$ from models $Q(X)$ and $Q_{\text{ind}}(X)$.

In accordance with previous studies (Schneidman et al., 2006; Tang et al., 2008), we report in the main text the results of each model obtained using the same samples that were used for estimating the model's parameters. This approach, however, potentially leads to so-called overfitting (Hastie et al., 2001); because the same samples were used for both estimating (training) the model's parameters and evaluating (testing) the model's performance, the tested performance might appear better than the true performance of the model with respect to the unknown true interaction structure of the data (Hastie et al., 2001). We therefore also examined the results using cross-validation, which is an approach that is more resistant to overfitting (Hastie et al., 2001; Shlens et al., 2006, 2009). Overall, the cross-validated results (reported in the supplemental material, available at www.jneurosci.org) support the major findings described in the main text.

Results

Spontaneous LFPs were recorded from 12 organotypic slice cultures of rat somatosensory cortex grown on 8×8 multichannel microelectrode arrays. Significant nLFP peaks were detected using a threshold procedure (see Materials and Methods) (Fig. 1A) (Beggs and Plenz, 2003, 2004). The nLFP peaks were indicative of population spikes (Jimbo and Robinson, 2000; Bédard et al., 2004) and occurred in bursts of approximately synchronous activity with quiescent intervals spanning several seconds (Fig. 1B), consistent with previous studies (Crain, 1966; Calvet, 1974; Gutnick et al., 1989; Maeda et al., 1995; Kamioka et al., 1996; Plenz and Aertsen, 1996; Corner et al., 2002; Eytan and Marom, 2006). The timing of the detected nLFP peaks is sufficient to reveal important spatiotemporal patterns of the activities obeying power scaling law, so-called neuronal avalanches (Beggs and Plenz, 2003, 2004; Stewart and Plenz, 2006; Gireesh and Plenz, 2008). The power laws are unchanged whether nLFPs are binarized or whether their amplitudes are considered (Beggs and Plenz, 2003), and the statistical properties of the avalanches are scale-invariant with respect to the nLFP threshold (Petermann et al., 2009). Moreover, using the timing alone of the nLFP peaks, the pairwise model was previously applied to a similar type of data (Tang et al., 2008). Therefore, in the present study, based only on the timing of the nLFP peaks (i.e., discarding their amplitudes), instantaneous nLFP peak activities were defined as the binary activities occurring in a single time bin on the array (see Materials and Methods) (Figs. 1C,D). We developed in detail a hierarchical, statistical model of the instantaneous nLFP activity. We first described the construction of the hierarchical model. We then verified the accuracy of the pairwise maximum entropy model for the nLFP activities in ensembles of 10 electrodes. Next, we compared the accuracy of the hierarchical model to that of the pairwise model on the same 10-electrode ensembles and showed that the hierarchical model achieves a superior level of accuracy despite its relative parsimony. We then assessed the performance of the hierarchical model on entire electrode arrays and described the characteristics of the electrode clusters and the cluster-level interactions captured by the model. The results reported in subsequent sections were obtained using the same samples that were used to estimate the model's parameters (see Model evaluation, Materials and Methods). We also reported the cross-validated results in the supplemental material (available at www.jneurosci.org) as a reference.

Construction of hierarchical model

Our hierarchical model treats interactions at two different spatial scales. A promising premise for constructing such a model can be found in our datasets (Fig. 2A). The spatial resolution of correlations between nLFP peak activities is relatively low, as correlations with similar magnitudes extended over multiple spatially contiguous electrodes. This suggests that such network correlations can be represented better and more parsimoniously by grouping together activities of functionally similar electrodes into larger, more appropriate units of interactions.

Figure 1. Recording of spontaneous neuronal activity from organotypic cultures of rat cortex using planar integrated MEA. **A**, Left, Culture grown on a 8×8 MEA dish. Right, LFP recorded from one electrode with the threshold set at -8 times the LFP SD (dashed line). **B**, Raster plot of significant nLFP deflections; points indicate time of maximum negative excursion below the threshold (nLFP peaks). Raster plot reveals periods of significant synchronized nLFP peak activity interspersed with quiescent periods. **C**, Raster plot with an increased time resolution. nLFP peaks are binned for subsequent analysis (in this example, $\Delta t = 4$ ms; see Materials and Methods). For visualization purposes, the size of each dot represents the voltage magnitude of nLFP peaks, which is not used in subsequent analyses. **D**, Detection of nLFP peak activity at an electrode is indicated by a "1" at the corresponding electrode; a "0" indicates no activity, illustrated on the MEA. This figure is reproduced with modifications from Beggs and Plenz (2004) (copyright 2004 Society for Neuroscience).

If we call the larger-scale unit a cluster (e.g., a cluster of electrodes, or the nLFP peak activities recorded at a cluster of electrodes) (Fig. 2B), then the hierarchical model uses units for interactions at two spatial scales or two levels, the electrode cluster level (global) and the single electrode level within a cluster (local). An activity of a cluster or a cluster activity is assumed to be determined by the activities recorded by the electrodes that constitute the cluster. To build the hierarchical model as parsimoniously as possible, we used the approach of the pairwise model at both levels; interactions at the cluster level are modeled by the pairwise model of clusters and interactions at the electrode level within a cluster are modeled by the pairwise model of electrodes. We assumed that a given electrode activity in a cluster is influenced by another electrode's activity in two ways: by the pairwise interaction with the other electrode if both electrodes are in the same cluster, and indirectly via cluster activity if the other electrode is in another cluster. This is a conditional independence assumption across clusters (Fig. 2C); more formally, electrode activity across clusters is statistically independent given the activity of each cluster, and electrodes in different clusters interact indirectly only via their cluster activities. To use such a hierarchical model, clusters must first be identified. Although there are various ways to identify clusters, we chose a practical and simple approach: we identified clusters using the activity homogeneity criterion (see below). A measure of cluster activity must also be defined. We tested three such measures: binary, log, and linear cluster activity. All types of cluster activity indicate the overall magnitude of electrode activities in a cluster, without specifying which electrodes are active (see below).

In mathematical terms, we denote instantaneous nLFP activity as $X = (x_1, x_2, \dots, x_n)$, where $x_i = (1, 0)$ indicates the detection or absence, respectively, of an nLFP peak at electrode i (Fig. 1D). The hierarchical model is denoted by Q_{hier} , and the pairwise model is denoted by Q_{pair} . Cluster activity in cluster I is denoted

Figure 2. Schematics of hierarchical model. **A**, An example of the spatial structure of activity correlations on the array obtained from one representative dataset. Contours indicate the correlation coefficients of array electrodes (open circles) relative to one reference electrode (solid circle). Contours are color-coded as indicated on the right. **B**, Schematic of two spatial levels of activity on the array. Left, Electrodes as the spatial units of activity. Solid circles indicate active electrodes and arrows indicate interactions between active pairs. Right, Clusters of electrodes as the spatial units of activity. Two clusters are indicated by different shades of gray, and arrows indicate the interaction between the clusters. **C**, Schematic of the hierarchical model on a nine-electrode array with two clusters. Top layer, Array-wide global activity is represented by the cluster as the unit for interaction; long-range correlations are described by cluster interactions (top arrow). Bottom layer, Local activity (within each cluster) is represented by electrodes as the unit. Short-range interactions (small arrow) occur only within each cluster. Thus, electrode activity is assumed to be independent across clusters when conditioned on cluster activity (downward arrows). **D**, Schematic of homogeneous property of a cluster. When the activities are homogeneous, all patterns of activities corresponding to the same level of cluster activity occur with the same probabilities. Therefore, probabilities do not depend on the identity of the active electrodes. Diagram illustrates all combinations of two active electrodes (corresponding to linear cluster activity $c_{lin} = 2$) for a four-electrode cluster (top), and their equal probabilities under the homogeneity assumption (bottom).

by c_I and the vector of all cluster activities is denoted by $C = (c_1, c_2, \dots, c_K)$, where the total number of clusters is K . The cluster activities are modeled by the pairwise model, denoted by $Q_{\text{pair}}(C)$, regardless of the type of cluster activity used (see Estimation of the hierarchical model's parameters, Materials and Methods). In each cluster I , there is a different number of electrodes and the pattern of the electrode activity in cluster I is denoted by X_I . The electrode activity in the cluster is then modeled by the pairwise model denoted by $Q_{\text{pair}}(X_I)$. With the conditional assumptions stated above, the hierarchical model $Q_{\text{hier}}(X)$ is given by

$$Q_{\text{hier}}(X) = Q_{\text{pair}}(C) \prod_{I=1}^K Q_{\text{pair}}(X_I | c_I), \quad (10)$$

where $Q_{\text{pair}}(X_I | c_I)$ is the conditional probability of X_I given the cluster activity c_I .

Thus, the hierarchical model assumes the short-range or local interactions of activities at electrodes within a cluster the same way as the original pairwise model of electrodes. In contrast, the long-range or global interactions of activities are modeled using the pairwise model of clusters; they can be viewed as a rather diffuse interaction summarizing electrode-level interactions. As mentioned above, clusters are identified by the homogeneity criterion (see Clustering criterion and clustering algorithm, Materials and Methods). Based on this criterion, we applied an iterative, backward search procedure to identify clusters. Once clusters were identified, the hierarchical model [at both $Q_{\text{pair}}(C)$ and $Q_{\text{pair}}(X_I)$] could be fitted with data in nearly the same man-

ner as the pairwise model. The homogeneity criterion identifies groups of electrodes so that the probability of activity patterns in the group depends as little as possible on which electrodes the activities have occurred. In perfectly homogeneous clusters, the probability of the electrodes' activities depends solely on the overall magnitude of the cluster activity (Fig. 2D). Among the types of cluster activity, information about the electrodes' activities in a cluster is minimal when using binary activity (indicates whether any electrode is active), maximal with linear activity (indicates the number of active electrodes), and intermediate with log activity (applies a logarithmic transformation of the linear activity; see Cluster activity, Materials and Methods). Selection of the type of activity depends on the balance between computational simplicity and information-representing capacity. We report the results for the simpler binary and log cluster activity in this work, as they were sufficient for describing the network interactions with high accuracy.

Performance of pairwise model when applied to small electrode ensembles

We first examined the accuracy of the pairwise model using small electrode ensembles (randomly sampled 10-electrode ensembles), largely following the approach taken in previous studies (Schneidman et al., 2006; Shlens et al., 2006; Tang et al., 2008) (see Materials and Methods). Three choices of bin sizes were examined: 4, 10, and 20 ms; we present mainly the results for the 20 ms bins to facilitate comparison with previous studies (results for the other bin sizes are provided in Tables 1, 2, and 3). The accuracy of the pairwise model $Q_{\text{pair}}(X)$ was measured using the F ratio (Schneidman et al., 2003) (see Model evaluation, Materials and Methods). The F ratio used for $Q_{\text{pair}}(X)$, or F_{pair} , is bounded by 0 and 1, the higher value indicating greater accuracy. It was previously demonstrated that the pairwise model attains a high accuracy in describing the observed distribution $\hat{P}(X)$: 0.90–0.99 for neural spike activity in retina (Schneidman et al., 2006; Shlens et al., 2006), 0.85 ± 0.06 for cortical spikes, and 0.93 ± 0.04 for cortical LFPs (Tang et al., 2008).

The accuracy of the pairwise model applied to nLFP peak activities of cortical slice cultures was also quite high and was consistent with that reported by Tang et al. (2008). As shown by the example of a randomly selected, 10-electrode ensemble (Fig. 3A), the probabilities estimated by the pairwise model are much closer to the equality line than those estimated by the independent model, implying that they are nearly equal to the observed probabilities. This observation was quantitatively confirmed using the F ratio (e.g., $F_{\text{pair}} = 0.938$ for the data in Fig. 3A) for all datasets (Fig. 3B). The mean of the F ratios for each dataset was quite high, and the average for all datasets was $F_{\text{pair}} = 0.924 \pm 0.030$ (200 random ensembles per dataset, averaged over all 12 datasets; the range of the mean over datasets was 0.89–0.96) (Fig. 3B).

Table 1. Accuracy of pairwise and hierarchical models applied to 10-electrode ensembles

Bin size (ms)	F_{pair}	$F_{\text{hier}}(\text{bin.})$	$F_{\text{hier}}(\text{log.})$	Model comparisons	
4	0.819 ± 0.075	0.886 ± 0.034	0.887 ± 0.041	$p < 1 \times 10^{-5*}$	$p = 0.22^{**}$
10	0.871 ± 0.042	0.911 ± 0.026	0.917 ± 0.027	$p < 1 \times 10^{-5*}$	$p < 1 \times 10^{-5**}$
20	0.924 ± 0.030	0.928 ± 0.026	0.938 ± 0.023	$p < 1 \times 10^{-5*}$	$p < 1 \times 10^{-5**}$

Entries indicate means $\pm$ SD of F -ratios obtained with the pairwise model (F_{pair}) and the hierarchical model with binary cluster activity [$F_{\text{hier}}(\text{bin.})$] and with log cluster activity [$F_{\text{hier}}(\text{log.})$]. Results described in the text were obtained with bin size = 20 ms. Model comparisons report p values of paired t test applied to the F ratios over all ensembles. $*F_{\text{hier}}(\text{bin.}) > F_{\text{pair}}$, $**F_{\text{hier}}(\text{log.}) > F_{\text{hier}}(\text{bin.})$.

Hierarchical model captures network correlations in small electrode ensembles better than the pairwise model

To examine the hierarchical model, we first identified clusters of electrodes on the entire array using our clustering criterion and algorithm. Figure 4*A* shows the clusters obtained on one example array. We examined the model using 10-electrode ensembles to allow for direct comparison of the performance with that of the pairwise model applied to the same data as a benchmark and also with previously reported results of the pairwise model. Cluster identity for electrodes in the 10-electrode sample was based on clustering obtained from the entire array, and the hierarchical model was then fitted and evaluated. Using 200 random ensembles of 10 electrodes per dataset, we first applied the hierarchical model using binary cluster activity. The hierarchical model achieved a higher level of accuracy than the pairwise model (Fig. 4). Figure 4*B* shows one example ensemble for which the hierarchical model had a higher F ratio: $F_{\text{hier}} = 0.928$ compared with $F_{\text{pair}} = 0.893$ of the pairwise model. The difference in F ratios over all ensembles (Fig. 4*C*) was highly significant (paired t test, $p < 1 \times 10^{-5}$), with $F_{\text{hier}} = 0.928 \pm 0.026$ for the hierarchical model and $F_{\text{pair}} = 0.924 \pm 0.030$ for the pairwise model. A comparison of the F ratios within each dataset (Fig. 4*D*) revealed that the hierarchical model obtained significantly higher F ratios for 7 of the 12 datasets, statistically equivalent F ratios for two datasets, and less accurate F ratios for the remaining three datasets (paired t test, $p < 0.05$). The relative advantage of the hierarchical model was even greater for other bin sizes (4 and 10 ms) (Table 1). Moreover, the hierarchical model performed this well despite using significantly fewer parameters than the pairwise model (Fig. 4*E*), suggesting that a hierarchical structure of interactions exists on the array.

To confirm the above observations, we performed two control tests using the same random ensembles of 10 electrodes. In the first test, we shuffled the clusters identified on each array by randomly swapping electrodes across clusters (e.g., the correct, original affiliation of electrode i in cluster I and electrode j in cluster J was swapped with electrode j in cluster I and electrode i in cluster J ; all electrodes were swapped by a random permutation). Thus, this test distorted the cluster structures originally identified on the entire arrays, but preserved the cluster sizes (the number of electrodes in each cluster). The hierarchical model performed significantly worse with the shuffled clusters than with the original clusters (paired t test, $p < 1 \times 10^{-5}$ for the comparison of F ratios over all ensembles). In the second test, we ignored the clusters originally identified over the entire array and instead

Figure 3. Pairwise model captures network correlations in random ensembles of 10 electrodes. **A**, Estimated probabilities for a representative 10-electrode ensemble. Each point corresponds to a probability estimate for a pattern of activity X (ordinate), plotted against empirical probabilities $\hat{P}(X)$ (abscissa). The black line indicates equality. Estimates by the pairwise model ($Q_{\text{pair}}(X)$, black circles) are significantly closer to the empirical probabilities than estimates by the independent model ($Q_{\text{ind}}(X)$, gray circles). **B**, Distributions of F ratios obtained by the pairwise model (F_{pair}) for each of the $n = 12$ datasets (200 random ensembles per dataset; datasets ordered by mean F_{pair}). Boxes represent the interquartile range (lower edge is the first quartile, horizontal line is the median, and upper edge is the third quartile), and whiskers extend from the upper and lower quartiles to indicate the sample maximum and minimum, respectively, that are not outliers. Outliers (gray crosses) are sample values lying > 1.5 times the interquartile range lower than the first quartile or higher than the third quartile.

applied the clustering algorithm to each 10-electrode ensemble by itself. We posited that the clusters obtained in this way, though not random, were inferior to the original clusters because information about electrodes outside the ensemble was not available. As expected, the clusters obtained using this approach were mostly different from the original ones (data not shown). Again, the performance of the hierarchical model with these clusters was significantly worse than that of the original hierarchical model (paired t test, $p < 1 \times 10^{-5}$ for the comparison of F ratios). These results were the same for all bin sizes. The findings from these two control tests suggest that the original cluster organization identified on the whole array was essential for the high level of accuracy attained by the hierarchical model and support the idea that a hierarchical interaction structure underlies the array-wide activity.

We then examined the hierarchical model using log cluster activity of the same 10-electrode ensembles (Fig. 5). Although binary cluster activity is the simplest measure, it contains less information than the other measures of cluster activity (see Cluster activity, Materials and Methods). Thus, we examined whether a richer measure of cluster activity would further improve the model's performance. The average F ratio achieved was significantly higher using log cluster activity ($F_{\text{hier}} = 0.938 \pm 0.023$; paired t test over all ensembles against the model using binary activity, $p < 1 \times 10^{-5}$) (Fig. 5). The improvement was significant in every dataset (Fig. 5*B*) and the F ratio of the hierarchical model was higher than that of the pairwise model in 10 of the 12 datasets (paired t test, $p < 0.05$). The increased complexity of log cluster activity is evident in the greater number of parameters required to fit the model (Fig. 5*C*), but it is still lower than that of the 55 parameters required for the pairwise model. The increased effec-

Figure 4. Hierarchical model using binary cluster activity was more accurate than the pairwise model at predicting activity probabilities in 10-electrode ensembles. **A**, Spatial distribution of color-coded clusters on the 8×8 MEA for one dataset (here the superficial-to-deep layers in cortical structure approximately correspond to the upper-to-lower rows in the MEA). Missing electrodes did not have significant nLFP peak activity and were not considered for analysis. Uncolored electrodes were not included in a cluster (see Materials and Methods). **B**, Estimated probabilities for a 10-electrode ensemble. Each point corresponds to a probability estimate for a pattern of activity X (ordinate) plotted against empirical probabilities $\hat{P}(X)$ (abscissa). Predictions by the hierarchical model $Q_{\text{hier}}(X)$ (black circles, $F_{\text{hier}} = 0.928$) are better than predictions by the pairwise model $Q_{\text{pair}}(X)$ (gray circles, $F_{\text{pair}} = 0.893$) for large patterns (patterns of smaller probabilities). The solid line indicates equality. **C**, Overall results from 12 datasets (200 ensembles per dataset). Each point corresponds to the F ratio of the full pairwise model (F_{pair}) and the hierarchical model (F_{hier}) for one ensemble. **D**, Model comparison in each dataset by the fraction of F ratios $F_{\text{hier}}/F_{\text{pair}}$ (datasets ordered by increasing mean fractions). Distributions are depicted in the same box plot format as those in Figure 3B. The hierarchical model had greater accuracy for most datasets (fraction > 1). In the dataset in which the hierarchical model had the worst performance, it still obtained at least 95% of the accuracy of the pairwise model on average. **E**, Breakdown of results by the number of clusters in the hierarchical model. Bottom, The number of parameters in the hierarchical model varied with the number of clusters; error bars are $\pm$ SD at each cluster count; the pairwise model required 55 parameters (dashed line). Top, The accuracy of the hierarchical model (F_{hier}) increased with an increase in the number of clusters. Note that the accuracy increased from two to four clusters despite the decrease in the number of parameters. The dashed line indicates the mean accuracy of the pairwise model (F_{pair}), shaded region indicates the mean $\pm$ SD. Results pooled from 12 datasets (200 ensembles per dataset).

tiveness of log relative to binary cluster activity depended on the bin size; the improvement in accuracy was smaller (but still significant) for 10 ms bins, but not statistically significant for 4 ms bins (Table 1). For all bin sizes, however, the hierarchical model with log cluster activity was confirmed to be a more accurate model of activity than the pairwise model.

Hierarchical model captures large-scale activity patterns on electrode arrays

After determining the accuracy of the hierarchical model on the 10-electrode ensembles and assessing the importance of the cluster structure, we next applied the hierarchical model to entire electrode arrays. To measure the accuracy of the model over the entire array, we used a modified F ratio based on the observed fractions of units active in a time bin (Shlens et al., 2009) (see Model evaluation, Materials and Methods for details). In the hierarchical model, the fractions of active units can be observed at two levels: as the fractions of active clusters at the cluster level, and as the fractions of active electrodes at the electrode level (Fig. 6A, an example is shown). At the cluster level, both hierarchical models (with binary and log cluster activity) fitted equally well with the observed fractions of active clusters. By comparison, the fractions of active clusters predicted by the independent model deviated strongly from the observed fractions. This was confirmed by examining the modified F ratio (F^c ratio) at the cluster level. Averaged over all datasets, the F^c ratios of the hierarchical models with binary and log cluster activity were 0.990 ± 0.010 and 0.996 ± 0.004 , respectively, and the difference between them was not significant ($p = 0.10$, paired t test). Summary results from all datasets at the cluster level are shown in Figure 6B (left panel; for log cluster activity only).

At the electrode level, the performance of the hierarchical model using log cluster activity captured most neural correlations in our datasets (Fig. 6, right panels). As expected, the fitted independent model deviated strongly from the observed fractions. The hierarchical model performed well using log cluster activity, but not as well with binary cluster activity. The use of binary cluster activity tended to underestimate the probabilities of large activities, whereas the use of log cluster activity accurately predicted all but the very largest fractions. This observation was quantitatively verified using the modified F ratio (F^c ratio) at the electrode level. Averaged over all datasets, the F^c ratio was 0.996 ± 0.004 for log cluster activity and only 0.972 ± 0.015 for binary cluster activity (the F^c ratio for log cluster activity being significantly more accurate than the F^c ratio for binary cluster activity; $p < 0.01$, paired t test). The degree of accuracy achieved was comparable to that previously reported (Shlens et al., 2009), where a pairwise-adjacent model achieved 98–99% accuracy for spike activity in the primate retina. Summary results for all 12 datasets (Fig. 6B, right panel) confirmed that the hierarchical model using log cluster activity accurately

ratio for binary cluster activity; $p < 0.01$, paired t test). The degree of accuracy achieved was comparable to that previously reported (Shlens et al., 2009), where a pairwise-adjacent model achieved 98–99% accuracy for spike activity in the primate retina. Summary results for all 12 datasets (Fig. 6B, right panel) confirmed that the hierarchical model using log cluster activity accurately

Figure 5. Hierarchical model using log cluster activity was even more accurate than that using binary activity in 10-electrode ensembles. Figure formats of **A–C** are the same as those of Figure 4C–E, respectively.

predicted the fraction of observed active electrodes. Changing the bin size from 20 to 10 or 4 ms did not significantly affect the results (Tables 2 and 3). These results confirmed the effectiveness of the hierarchical model in capturing the large-scale activity observed on the arrays and suggested that log cluster activity is better suited for large-scale datasets.

Structure captured by the hierarchical model

We analyzed the interaction structures and properties of the datasets revealed by applying the hierarchical model to an entire array (Fig. 7). The results reported below (bin size = 20 ms) are similar to those for 4 and 10 ms bin sizes (data not shown). The clusters identified by the clustering algorithm tended to cover spatially contiguous areas on the array. Indeed, by comparing the distributions of spatial distances for electrode pairs in the same and in different clusters (Fig. 7A), we found that electrodes in the same cluster were located significantly closer to each other on the array

Table 2. Accuracy of hierarchical model applied to whole electrode arrays at the cluster level

Bin size (ms)	Binary	Log	Model comparison*
4	0.970 ± 0.026	0.983 ± 0.014	$p = 0.17$
10	0.993 ± 0.008	0.993 ± 0.007	$p = 0.29$
20	0.990 ± 0.010	0.996 ± 0.004	$p = 0.10$

Entries indicate means ± SD of the modified F ratios obtained with the hierarchical model using binary and log cluster activities. F ratios reflect the difference between empirical and model distributions of the ratios of active clusters in a sample. Results described in the text were obtained with bin size = 20 ms. Model comparisons report p values of paired t test applied to the F ratios over all datasets. *Log > binary.

Table 3. Accuracy of hierarchical model with log cluster activity applied to whole electrode arrays at the electrode level

Bin size (ms)	Binary	Log	Model comparison*
4	0.979 ± 0.012	0.990 ± 0.005	$p < 0.01$
10	0.975 ± 0.011	0.994 ± 0.006	$p < 0.01$
20	0.972 ± 0.015	0.996 ± 0.004	$p < 0.01$

Entries indicate means ± SD of F ratios obtained with the hierarchical model using binary and log cluster activities. F ratios reflect the difference between empirical and model distributions of the ratio of active electrodes in a sample. Results described in the text were obtained with bin size = 20 ms. Model comparisons report p values of paired t test applied to the F ratios over all datasets. *Log > binary.

than electrodes in different clusters (Kolmogorov–Smirnov test, $p < 1 \times 10^{-5}$; distributions pooled from 12 datasets). Furthermore, the activities of electrode pairs in the same cluster tended to be more strongly correlated than those of electrode pairs across clusters, as shown by the respective distributions of correlation coefficients in Figure 7B (Kolmogorov–Smirnov test, $p < 1 \times 10^{-5}$; pooled distributions). Note that this is a nontrivial characteristic of the activities within clusters (see Discussion).

We then investigated the characteristics of cluster interactions. There was a consistent relationship of cluster interaction with distance between clusters, examined two ways (see Evaluating the cluster-level interactions using the electrode-level, pairwise interactions, Materials and Methods). First, cluster pairs were classified as “adjacent” if at least one electrode from one cluster was adjacent to another electrode in the other cluster; otherwise, the cluster pair was classified as “nonadjacent.” For both binary (data not shown) and log cluster activity (Fig. 7C), cluster interactions were significantly higher for adjacent clusters than for the nonadjacent clusters (Mann–Whitney U test, $p < 0.001$; pooled distributions). Second, the physical distances between the clusters’ centers of mass (i.e., average location of member electrodes on the array) were computed and correlated against cluster interactions over all datasets. Cluster interactions were significantly anticorrelated with the distances between the centers of mass for both binary (Pearson’s $r = -0.58$, $p < 0.01$) and log (Pearson’s $r = -0.29$, $p < 0.01$) cluster activity (data not shown). Finally, we examined how cluster interactions might be seen as electrode-level, pairwise interactions across clusters if the pairwise model was fitted to the same dataset. For this purpose, we obtained the weighted mean of cluster interaction terms for each cluster pair (see Evaluating the cluster-level interactions using the electrode-level, pairwise interactions, Materials and Methods). For the fit of the pairwise model on the entire array, we substituted ensemble averages for the electrode-level pairwise interactions. Using these estimates, we obtained their mean and maximum values across each cluster pair. We then compared the weighted mean of the cluster interaction with the mean and maximum values of the electrode-level interactions. The cluster-level interactions tended to be proportional to the maximum electrode-level interactions across the two clusters. Only the maximum electrode-level interaction was consistently and signif-

Figure 6. Hierarchical model captures array-wide activity patterns. **A**, Distribution of fractions of active clusters (left) and fractions of active electrodes (right) in a representative dataset. Bars indicate the empirical probabilities, and lines are predictions made by the independent model (dashed black line), hierarchical model with binary cluster activities (solid black line), and hierarchical model with log cluster activities (solid red line). Solid black and red lines overlap in the left plot. **B**, Summary predictions by the hierarchical model with log clusters for the fraction of active clusters (left), and the fraction of active electrodes (right). Each circle corresponds to the observed and predicted probabilities of a specific fraction of the activity in a given dataset; different datasets are indicated by different colors. Gray circles indicate predictions of the hierarchical model with binary clusters.

icantly correlated to the cluster-level interaction for both binary (data not shown) and log cluster activity (Fig. 7D). The Pearson correlation between maximum electrode-level interaction and either the binary or log cluster-level interaction (Fig. 7D) was $r = 0.46$ ($p < 1 \times 10^{-5}$), whereas for mean electrode-level interactions, the correlations were $r = 0.09$ ($p = 0.18$) and $r = -0.12$ ($p = 0.12$) for the binary and log cluster interactions, respectively.

Discussion

We introduced a hierarchical model for describing neuronal activities that identifies appropriate units of interactions at different scales: at the level of single electrodes and at electrode clusters. Using the hierarchical model, we analyzed spontaneous nLFP peak activities in cortical organotypic slice cultures. By identifying clusters as substructures of the network interaction, the hierarchical model, despite its relative parsimony, performed better than the pairwise model with 10-electrode ensembles. Furthermore, the hierarchical model accurately predicted the frequency of correlated activities at both the electrode and cluster levels for the entire array. These results suggest that the hierarchical model taps into the appropriate interaction structure of the underlying network through the identified clusters.

Characteristics of clusters and cluster-level interactions

The nLFP peak activities at electrodes within each cluster were highly correlated (Fig. 7B) and spatially clustered (Fig. 7A), often covering spatially contiguous areas on the array (Fig. 4A). These

characteristics were not direct consequences of our clustering criterion—the homogeneity criterion. Homogeneity does not necessarily imply correlation; e.g., a cluster of electrodes in which all activities with the same firing rates are statistically independent from one another is a perfectly homogeneous cluster. Thus, the characteristics of the clusters represent the specific nature of the spontaneous nLFP peak activities. One important question remains: what neural substrates underlie our identified clusters? Based on visual inspection of the spatially contiguous clusters, we speculate that the clusters correspond to the underlying cortical layer structure (Fig. 4A). Although very appealing, confirmation of this speculation requires application of the hierarchical model to datasets with more detailed anatomical information than those used in the current analysis.

The criterion used to identify the clusters was homogeneity within a cluster (homogeneous activities at electrodes within a cluster); this criterion does not ensure homogeneity of interactions across clusters. Homogeneity in cross-cluster electrode-level interactions, however, is also needed so that the interactions can be properly abstracted as cluster-level interactions, and the models consequently fit the datasets well (see below and Appendix). Therefore, the high degree of accuracy attained by the model suggests a specific feature of the interactions of nLFP peak activities in our datasets; internally

homogeneous clusters tend to interact homogeneously with each other. Improving the identification of the clusters by taking into account interactions not only within, but also across, clusters may improve the applicability of the hierarchical model to other types of datasets.

Cluster interactions can be viewed as a summary of long-range interactions of activities at electrodes. Using clusters as an additional unit, the hierarchical model captures long-range interactions as rather diffuse interactions between clusters. This notion of the hierarchical model's long-range interactions is broadly consistent with the conventional understanding that nLFP peaks are indicative of collective neural activity around the recording electrode (Jimbo and Robinson, 2000; Bédard et al., 2004). nLFP peaks correlate with multiunit spike activities recorded from the same electrode, both *in vitro* (Gireesh and Plenz, 2008) and *in vivo* (Rasch et al., 2008; Petermann et al., 2009). Accordingly, the interactions among long-range LFP peak activities may not be point-wise between electrode-level interactions (i.e., pairwise interactions), but may rather be mediated via a variety of diffuse interactions among the underlying neural activities.

Relationship of cluster-level interactions to electrode-level interactions

One might view the hierarchical model as a further simplification of the pairwise model because the pairwise model is still used at the electrode and cluster levels. This view, however, misses an

important point; the hierarchical model is not a subclass, and therefore not just a simplification, of the pairwise model. The assumed underlying interaction structure differs significantly between models: cluster-level pairwise interactions include higher-than-pairwise-order interactions at the electrode level that do not exist in the original pairwise model. The hierarchical model uses upper-level (cluster) interactions to characterize whole combinations of interactions at the lower (electrode) level. An upper-level interaction between two upper-level elements mediates all the interactions between the corresponding two sets of lower-level elements. This induces higher-order interactions at the lower level. This conceptual explanation can be easily transformed into, and thus clarified by, mathematical formulations (see Appendix). Thus, the instantaneous interactions of nLFP peak activities in our datasets can also be considered to use higher-than-pairwise-order terms for long-range interactions. The high accuracy of the hierarchical model leads to an intriguing question: what kinds of higher-order interactions at the electrode level are driving the overall network interactions? We plan future investigations to answer this important question.

In a broader context, we propose that identifying the appropriate units for interactions is essential for simplifying the description of network interactions. The units of interaction should match the interaction structure of the underlying networks under investigation. For each neural system under investigation, a model of appropriate interaction structures with appropriate units must be developed by comparing alternative models. Each model represents a class of assumed interaction structures. When the assumed structure better matches the underlying structure, a model with even a smaller number of parameters may more faithfully reflect the underlying structure, and thus fit the data better than a model with a larger number of parameters. In our case, the hierarchical organization with the additional larger-scale unit, despite its relative parsimony, led to greater accuracy than the pairwise model alone. This suggests that the hierarchical model fits the underlying network structure better than the full pairwise model. However, this might not be the case in other systems. For example, a simpler subclass of pairwise models, the adjacent-pairwise model, can describe the activities of ON-OFF parasol retinal ganglion cells almost as well as the original full pairwise model (Shlens et al., 2006, 2009), indicating that the adjacent-pairwise model may better correspond to the interaction structures underlying the retinal activities. For this kind of interaction structure, the hierarchical model might not be as accurate as either the pairwise or adjacent-pairwise models. In contrast, for the type of structure examined in our datasets, the hierarchical model was more accurate than the entire class of pairwise models (Figs. 3, 4).

Computationally, diffuse higher-order terms for long-range interactions may allow for a more robust but less specific communication between the different clusters. This robustness may

Figure 7. Characteristics of homogeneous clusters in the hierarchical model. **A**, Significantly different distributions of spatial distances (μm) between pairs of electrodes within the same cluster (black) and in different clusters (gray); Kolmogorov–Smirnov test, $p < 1 \times 10^{-5}$. **B**, Significantly different distributions of correlation coefficients for electrode pairs in the same cluster (black) and in different clusters (gray); Kolmogorov–Smirnov test, $p < 1 \times 10^{-5}$. **C**, Significant difference in interactions between nonadjacent clusters (Non-adjac.) and adjacent clusters (Adjacent); Mann–Whitney U test, $p < 0.001$. Distributions are depicted in the same box plot format as in Figure 3B. **D**, Cluster interactions (with log activity) were proportional to the maximum interaction between electrodes across the corresponding clusters; Pearson's $r = 0.46$, $p < 1 \times 10^{-5}$.

help to establish specific kinds of synchronized neural patterns theorized, such as cell assemblies (Abeles, 1991; Diesmann et al., 1999; Plenz and Thiagarajan, 2007; Teramiae and Fukai, 2007). The information carried by these higher-order terms might be significant, even if their numbers or magnitudes are relatively small (Bezzi et al., 2002; Nakahara and Amari, 2002). In the present study, the hierarchical interactions were studied by using bin sizes ranging from 4 to 20 ms and by treating the nLFP peaks in the same bin as instantaneous. Thus, the very nature of the interactions looks spatial rather than temporal. Studies on neuronal avalanches, however, have also shown that the propagation of synchronized activity as reflected in successive nLFPs on the array is rather rapid (Beggs and Plenz, 2003; Gireesh and Plenz, 2008); many avalanches are shorter than 10–20 ms at the inter-electrode distance of 200 μm . Thus, for instance, using 20 ms bin size, the hierarchical interactions captured by the present study also contain the mapping of these spatiotemporal cascades into their corresponding spatial patterns. This interpretation is supported by the improved accuracy at lower temporal resolutions (Table 1). In general, we expect that in the temporal domain, different forms of dynamic attractors may contribute to computations of neural activities such as cortical and hippocampal activities (Miyashita, 2004; Machens et al., 2005; Wills et al., 2005; Jones et al., 2007; Akrami et al., 2009; Mongillo et al., 2008; Wang, 2008). Once the hierarchical model is extended to the temporal domain, it will help to identify what form of temporal higher-order interactions exists and thus reveal what type of attractors might exist in the neural system under investigation (see below).

Characteristics of hierarchical model

The hierarchical model using either binary or log cluster activity significantly decreases the number of model parameters compared with the pairwise model, offering a potentially strong advantage for applying it to large-scale datasets. When applicable, binary activity is most preferable because of its simplicity and scalability. At the level of 10 electrodes, the hierarchical model with binary activity performed better than the pairwise model (Table 1). When the dataset was enlarged to a whole array (~60 electrodes), log activity was required to sufficiently describe the activities in our datasets (Table 3). Whether the log cluster activity can accommodate even larger datasets remains to be examined. One potential concern is that log cluster activity increases the structural dimensionality of the model, but this effect was negligible in the current datasets. Using the total number of parameters of the pairwise model as a reference in each dataset (e.g., for 60 electrodes, the pairwise model requires 1830 parameters for all first- and second-order terms), with binary cluster activity, the hierarchical model required an average of 24.8% of the parameters in the reference, and only 32.9% using log cluster activity. Thus, using log cluster activity only modestly increased the number of parameters required.

Roudi et al. (2009) cautioned that the results of the pairwise model applied to a smaller number of elements may not be straightforwardly extrapolated to a larger number of elements; the pairwise model's performance can significantly deteriorate as the number of elements increases. This issue is strongly related to the intrinsic nature of the pairwise model, in which the number of parameters rapidly increases relative to the number of interacting elements. The hierarchical model may be resilient to such a deterioration due to its hierarchical and parsimonious use of pairwise interactions, akin to the adjacent-pairwise model, as discussed by Roudi et al. (2009). Further studies are needed to clarify this issue.

The current hierarchical model discards the nLFP peak amplitudes, thereby treating the peaks as binary variables. This approach significantly simplifies the model's complexity (e.g., the number of parameters) for both hierarchical and original pairwise models, thereby making their applications to the data manageable. Studies on neural avalanches have shown that the nLFP peaks amplitudes correlate with the size of the locally active neuronal group. Importantly, though, the essential statistical properties of avalanches (e.g., the power law in avalanche size distribution and the lifetime of size distributions) are scale-invariant to the amplitude threshold used to identify nLFP peaks (Petermann et al., 2009). Such scale-invariance supports the current treatment of nLFPs as binarized variables. To include the amplitudes in future analysis, hierarchical models can treat them as m -discrete different levels and use an m -size-alphabet pairwise model (Amari, 2001; Ince et al., 2009). Such models, however, have a considerably increased complexity. The technique using the log scale with the floor function, as in the present study, helps to suppress the increase in complexity. Whether inclusion of amplitude information leads to better characterization of the interaction structure and whether the model's increased complexity can still be managed under the usual sample number limitation remain to be examined.

Temporal correlations or interactions are certainly of great interest to us for future studies. The datasets we investigated in the present study contained rich spatiotemporal interactions, i.e., neuronal avalanches (Beggs and Plenz, 2003, 2004; Gireesh and Plenz, 2008; Stewart and Plenz, 2008; Petermann et al., 2009). Analyzing similar datasets by using pairwise models, Tang et al.

(2008) suggested that incorporating temporal correlations is needed to provide a full description of spatiotemporal activity patterns by the pairwise models. Given the hierarchical model's high accuracy for instantaneous activity even for the whole array, there are several potential merits for extending it into the temporal domain, and a simple extension should be straightforward and easy to implement. Such a temporal hierarchical model would provide the same simplicity and might be more accurate than a simple temporal pairwise model, particularly if the temporal long-range interactions are diffuse. Future studies will investigate how temporal interactions, including higher-order interactions, and the spatial interactions examined in the present study constitute the spatiotemporal dynamics.

Appendix

Higher-order interactions in the hierarchical model

Using an example, we demonstrate here the relationship between the hierarchical model and higher-order interaction terms. The hierarchical model in its original construction treats interactions separately at two levels: electrodes and clusters. Here we show how the interactions at the two levels can be mapped together onto the electrode-level interactions, which reveals the presence of higher-order interaction terms. For this demonstration, we use log-linear or information geometric (IG) coordinates (Nakahara and Amari, 2002) with the simplest example being the case of three elements or electrodes: $x = (x_1, x_2, x_3)$. The intrinsic dimensionality of this full probability distribution is 7, as there are a total of 8 possible patterns, and the summation of probabilities of all patterns is 1, which eliminates one of the eight dimensions (Nakahara and Amari, 2002). The full-dimensional probability is given using IG coordinates by

$$\log P(\mathbf{X}) = \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_3 + \theta_{12} x_1 x_2 + \theta_{23} x_2 x_3 + \theta_{13} x_1 x_3 + \theta_{123} x_1 x_2 x_3 - \psi, \quad (11)$$

where θ_i , θ_{ij} , and θ_{ijk} are the first, second, and third-order coefficients, respectively, of the log-linear model, and ψ is the normalization term. In contrast, the pairwise model in this case is given by

$$\log Q_{\text{pair}}(\mathbf{X}) = \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_3 + \theta_{12} x_1 x_2 + \theta_{23} x_2 x_3 + \theta_{13} x_1 x_3 - \psi. \quad (12)$$

A comparison of these two equations demonstrates that the pairwise model does not include the third-order interaction term θ_{123} . For example, let us examine the hierarchical model with linear cluster activity, where two homogeneous clusters are assumed: one is composed of (x_1, x_2) and the other is x_3 itself. The probability distribution of this hierarchical model is given by

$$Q_{\text{hier}}(\mathbf{X}) = Q_{\text{pair}}(c_1, c_2) Q_{\text{pair}}(x_1, x_2 | c_1) Q_{\text{pair}}(x_3 | c_2) = \frac{Q_{\text{pair}}(x_1 + x_2, x_3) Q_{\text{pair}}(x_1, x_2)}{P(x_1 + x_2)}, \quad (13)$$

where $c_1 = x_1 + x_2$ and $c_2 = x_3$. The interaction terms of Q_{hier} can be revealed by representing it as

$$\log Q_{\text{hier}}(\mathbf{X}) = \log Q_{\text{pair}}(x_1 + x_2, x_3) + \log Q_{\text{pair}}(x_1, x_2) - \log P(x_1 + x_2). \quad (14)$$

On the right-hand side, using the symmetry of x_1 and x_2 , we can expand two probabilities as

$$\log P(x_1 + x_2) = \theta'_1(x_1 + x_2) + \theta'_{12}x_1x_2 - \psi' \quad (15)$$

and

$$\log Q_{\text{pair}}(x_1 + x_2, x_3) = \theta''_1(x_1 + x_2) + \theta''_3x_3 + \theta''_{12}x_1x_2 + \theta''_{13}(x_1 + x_2)x_3 + \theta''_{123}x_1x_2x_3 - \psi''. \quad (16)$$

As this second probability includes the third-order interaction term θ''_{123} , we conclude that the hierarchical model also includes the third-order term. The model's structural dimensionality is the same as that of the pairwise model (6 in this example; see Cluster activity, Materials and Methods). This is easily confirmed by Equations 14–16 along with the fact that $Q_{\text{pair}}(x_1, x_2)$ contains only first- and second-order terms. The same structural dimensionality is achieved by the hierarchical model, despite the inclusion of the third-order term, because of the symmetry of the interaction between the two clusters (e.g., in the case of Equation 16, the two pairwise interactions for x_1x_3 and x_2x_3 are assumed to be the same). Also, once Q_{hier} is expressed as in the right-hand side of Equation 14, it is obvious that the dual coordinates using the expectation parameters can also represent the hierarchical model, implying that it is possible to cast the hierarchical model in a maximum entropy formulation (Amari, 2001, 2003; Nakahara and Amari, 2002).

The hierarchical model with the binary cluster activity is similarly derived; the cluster variable is now represented by $c_1 = x_1 \vee x_2$, and $\vee$ denotes the binary “OR” function. Then the model is given by

$$Q_{\text{hier}}(\mathbf{X}) = \frac{Q_{\text{pair}}(x_1 \vee x_2, x_3)Q_{\text{pair}}(x_1, x_2)}{P(x_1 \vee x_2)}, \quad (17)$$

and the component probabilities are expanded as

$$\begin{aligned} \log P(x_1 \vee x_2) &= \theta_1[x_1(1 - x_2) + (1 - x_1)x_2 + x_1x_2] - \psi \\ &= \theta_1(x_1 + x_2 - x_1x_2) - \psi \end{aligned} \quad (18)$$

and

$$\begin{aligned} \log Q_{\text{pair}}(x_1 \vee x_2, x_3) &= \theta_1(x_1 + x_2 - x_1x_2) + \theta_3x_3 \\ &+ \theta_{13}(x_1x_3 + x_2x_3 - x_1x_2x_3) - \psi. \end{aligned} \quad (19)$$

Compared with their expansions for linear activity (Equations 15 and 16, respectively), one can see that the number of coefficients has been reduced [the structural dimensionality of model $Q_{\text{hier}}(\mathbf{X})$ is now 5]. Notably, the third-order interaction term for $x_1x_2x_3$ (i.e., θ_{13} in Equation 19) still exists, but is now set equal to the two other pairwise interaction terms for x_1x_3 and x_2x_3 . Thus, the hierarchical model with binary activity still explicitly contains a higher order term than the pairwise model (second-order), but its dimensionality is effectively reduced, imposing additional constraints on the relationship of the higher-order interaction terms to the lower-order terms (first- and second-order terms). Finally, although the reduction in dimensionality achieved by using the binary activity is small in this example, it should be noted that this is only due to the simplicity of the example; the reduction would be proportionally greater with increasing cluster sizes.

References

Abeles M (1991) *Corticonics: neural circuits of the cerebral cortex*. Cambridge, UK: Cambridge UP.

Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing

patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 70:1629–1638.

Aertsen A, Arndt M (1993) Response synchronization in the visual cortex. *Curr Opin Neurobiol* 3:586–594.

Akrami A, Liu Y, Treves A, Jagadeesh B (2009) Converging neuronal activity in inferior temporal cortex during the classification of morphed stimuli. *Cereb Cortex* 19:760–764.

Amari SI (1977) Neural theory of association and concept-formation. *Biol Cybern* 26:175–185.

Amari SI (2001) Information geometry on hierarchy of probability distributions. *IEEE Trans Inf Theory* 47:1701–1711.

Amari SI, Nakahara H, Wu S, Sakai Y (2003) Synchronous firing and higher-order interactions in neuron pool. *Neural Comput* 15:127–142.

Bédard C, Kröger H, Destexhe A (2004) Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophys J* 86:1829–1842.

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

Beggs JM, 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.

Bezzi M, Diamond ME, Treves A (2002) Redundancy and synergy arising from pairwise correlations in neuronal ensembles. *J Comput Neurosci* 12:165–174.

Calvet MC (1974) Patterns of spontaneous electrical activity in tissue cultures of mammalian cerebral cortex vs. cerebellum. *Brain Res* 69:281–295.

Corner MA, van Pelt J, Wolters PS, Baker RE, Nuytink RH (2002) Physiological effects of sustained blockade of excitatory synaptic transmission on spontaneously active developing neuronal networks: an inquiry into the reciprocal linkage between intrinsic biorhythms and neuroplasticity in early ontogeny. *Neurosci Biobehav Rev* 26:127–185.

Crain S (1966) Development of “organotypic” bioelectric activity in central nervous tissues during maturation in culture. In: *International Review of Neurobiology* (Pfeiffer CC, Smythies JR, eds), pp 1–43. New York: Academic.

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.

Eytan D, Marom S (2006) Dynamics and effective topology underlying synchronization in networks of cortical neurons. *J Neurosci* 26:8465–8476.

Gireesh ED, Plenz D (2008) Neuronal avalanches organize as nested theta- and beta/gamma-oscillations during development of cortical layer 2/3. *Proc Natl Acad Sci U S A* 105:7576–7581.

Grant M, Boyd S (2008) Graph implementations for nonsmooth convex programs. In: *Recent Advances in Learning and Control* (a tribute to M. Vidyasagar) (Blondel V, Boyd S, Kimura H, eds), pp 95–110. Springer.

Grant M, Boyd S (2009) CVX: Matlab software for disciplined convex programming (web page and software). Retrieved June 5, 2009 from <http://cvxr.com/cvx>.

Gutnick MJ, Wolfson B, Baldino F Jr (1989) Synchronized neuronal activities in neocortical explant cultures. *Exp Brain Res* 76:131–140.

Hastie T, Tibshirani R, Friedman J (2001) *The elements of statistical learning: data mining, inference, and prediction*, Ed 1. New York: Springer.

Hertz J, Krogh A, Palmer RG (1991) *Introduction to the theory of neural computation*. Redwood City, CA: Addison-Wesley.

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A* 79:2554–2558.

Hopfield JJ, Tank DW (1986) Computing with neural circuits: a model. *Science* 233:625–633.

Ince RA, Petersen RS, Swan DC, Panzeri S (2009) Python for information theoretic analysis of neural data. *Front Neuroinformatics* 3:4.

Jaynes ET (1957) Information theory and statistical mechanics. *Phys Rev* 106:620–630.

Jimbo Y, Robinson HP (2000) Propagation of spontaneous synchronized activity in cortical slice cultures recorded by planar electrode arrays. *Bioelectrochemistry* 51:107–115.

Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A* 104:18772–18777.

Kamioka H, Maeda E, Jimbo Y, Robinson HP, Kawana A (1996) Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neurosci Lett* 206:109–112.

- Karpiak VC, Plenz D (2002) Preparation and maintenance of organotypic cultures for multielectrode array recordings. In: *Current protocols in neuroscience* (Crawley JN, Gerfen CR, Rogawski MA, Sibley DR, Skolnick P, Wray S, eds) pp 1, 6.15–16.15. New York: Wiley.
- Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307:1121–1124.
- Maeda E, Robinson HP, Kawana A (1995) The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J Neurosci* 15:6834–6845.
- Martignoni L, Deco G, Laskey K, Diamond M, Freiwald W, Vaadia E (2000) Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput* 12:2621–2653.
- Miyashita Y (2004) Cognitive memory: cellular and network machineries and their top-down control. *Science* 306:435–440.
- Mongillo G, Barak O, Tsodyks M (2008) Synaptic theory of working memory. *Science* 319:1543–1546.
- Montani F, Ince RA, Senatore R, Arabzadeh E, Diamond ME, Panzeri S (2009) The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philos Transact A Math Phys Eng Sci* 367:3297–3310.
- Nakahara H, Amari S (2002) Information-geometric measure for neural spikes. *Neural Comput* 14:2269–2316.
- Nemenman I, Shafee F, Bialek W (2002) Entropy and inference, revisited. In: *Advances in neural information processing systems* (Dietterich TG, Becker S, Ghahramani Z, eds). Cambridge, MA: MIT.
- Nirenberg SH, Victor JD (2007) Analyzing the activity of large populations of neurons: how tractable is the problem? *Curr Opin Neurobiol* 17:397–400.
- Paninski L (2003) Estimation of Entropy and Mutual Information. *Neural Comput* 15:1191–1253.
- Petermann T, Thiagarajan TC, Lebedev MA, Nicolelis MA, Chialvo DR, Plenz D (2009) Spontaneous cortical activity in awake monkeys composed of neuronal avalanches. *Proc Natl Acad Sci U S A* 106:15921–15926.
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999.
- Plenz D, Aertsen A (1996) Neural dynamics in cortex-striatum co-cultures. II. Spatiotemporal characteristics of neuronal activity. *Neuroscience* 70:893–924.
- 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.
- Plenz D, Thiagarajan TC (2007) The organizing principles of neuronal avalanches: cell assemblies in the cortex? *Trends Neurosci* 30:101–110.
- Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK (2008) Inferring spike trains from local field potentials. *J Neurophysiol* 99:1461–1476.
- Richmond BJ, Gawne TJ (1998) The relationship between neuronal codes and cortical organization. In: *Neuronal ensembles strategies for recording and decoding* (Eichebaum HB, Davis JL, Eichenbaum H, eds), pp 57–79. New York: Wiley.
- Roudi Y, Nirenberg S, Latham PE (2009) Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. *PLoS Comput Biol* 5:e1000380.
- Salinas E, Sejnowski TJ (2001) Correlated neuronal activity and the flow of neural information. *Nat Rev Neurosci* 2:539–550.
- Schneidman E, Still S, Berry MJ, Bialek W (2003) Network information and connected correlations. *Phys Rev Lett* 91:238701.
- Schneidman E, Berry MJ 2nd, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012.
- Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ (2006) The structure of multi-neuron firing patterns in primate retina. *J Neurosci* 26:8254–8266.
- Shlens J, Field GD, Gauthier JL, Greschner M, Sher A, Litke AM, Chichilnisky EJ (2009) The structure of large-scale synchronized firing in primate retina. *J Neurosci* 29:5022–5031.
- Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586.
- Stewart CV, Plenz D (2006) Inverted-U profile of dopamine-NMDA-mediated spontaneous avalanche recurrence in superficial layers of rat prefrontal cortex. *J Neurosci* 26:8148–8159.
- Stewart CV, Plenz D (2008) Homeostasis of neuronal avalanches during postnatal cortex development in vitro. *J Neurosci Methods* 169:405–416.
- Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, Prieto A, Petrusca D, Grivich MI, Sher A, Hottowy P, Dabrowski W, Litke AM, Beggs JM (2008) A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci* 28:505–518.
- Teramae J, Fukai T (2007) Local cortical circuit model inferred from power-law distributed neuronal avalanches. *J Comput Neurosci* 22:301–312.
- Wang XJ (2008) Decision making in recurrent neuronal circuits. *Neuron* 60:215–234.
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876.