

Patterns of Ongoing Activity and the Functional Architecture of the Primary Visual Cortex

Joshua A. Goldberg,^{1,3,4*} Uri Rokni,^{1,3} and Haim Sompolinsky^{1,2}

¹Racah Institute of Physics and The Interdisciplinary Center for Neural Computation

The Hebrew University
Jerusalem 91904
Israel

²Department of Molecular and Cellular Biology
Harvard University
Cambridge, Massachusetts 02138

Summary

Ongoing spontaneous activity in the cerebral cortex exhibits complex spatiotemporal patterns in the absence of sensory stimuli. To elucidate the nature of this ongoing activity, we present a theoretical treatment of two contrasting scenarios of cortical dynamics: (1) fluctuations about a single background state and (2) wandering among multiple “attractor” states, which encode a single or several stimulus features. Studying simplified network rate models of the primary visual cortex (V1), we show that the single state scenario is characterized by fast and high-dimensional Gaussian-like fluctuations, whereas in the multiple state scenario the fluctuations are slow, low dimensional, and highly non-Gaussian. Studying a more realistic model that incorporates correlations in the feed-forward input, spatially restricted cortical interactions, and an experimentally derived layout of pinwheels, we show that recent optical-imaging data of ongoing activity in V1 are consistent with the presence of either a single background state or multiple attractor states encoding many features.

Introduction

Mounting experimental evidence demonstrates that the ongoing spontaneous activity in cortex in the absence of sensorimotor processing can exhibit complex spatiotemporal structures (Arieli et al., 1995, 1996; Contreras and Steriade, 1995; Tsodyks et al., 1999; Lampl et al., 1999; Destexhe et al., 1999; Kenet et al., 2003). One possibility is that this activity represents some background state to which the cortical network relaxes in the absence of sensory inputs. Alternatively, it has been proposed that ongoing activity represents the system’s dynamic switching between a set of intrinsic cortical states (Arieli et al., 1995, 1996; Tsodyks et al., 1999; Kenet et al., 2003; Ringach, 2003).

Optical imaging of the primary visual cortex (V1) using voltage-sensitive dyes (VSD) is an important tool for addressing this issue (Arieli et al., 1995, 1996; Tsodyks

et al., 1999; Shoham et al., 1999; Fitzpatrick, 2000; Kenet et al., 2003). Studies using VSD produce two kinds of optical images: evoked maps and spontaneous ones. The highly reproducible evoked maps are generated by averaging the optical image over a period of presentation of a salient stimulus, such as a moving grating with a fixed orientation. In contrast, a spontaneous map is a single image recorded in the absence of oriented visual stimuli. In order to characterize the nature of ongoing activity in V1, Grinvald and colleagues studied the relationship between evoked and spontaneous maps (Tsodyks et al., 1999; Kenet et al., 2003). Interestingly, they found that instantaneous spontaneous maps are often highly correlated with the maps evoked by oriented gratings. By analyzing the dynamics of the ongoing activity, they conclude that the spontaneous cortical activity reflects a dynamic switching between a set of intrinsic states, many of which correspond closely to orientation maps (Ringach, 2003).

Intrinsically generated states of a dynamical system are known as “attractor states.” The presence of these states is believed to underlie various phenomena in cortical dynamics. Ermentrout and Cowan proposed that under the influence of drugs the network of V1 settles into attractor states which are perceived as visual hallucinations (Ermentrout and Cowan, 1979). These states are manifested as patterns across V1 whose structure is determined by the underlying large-scale cortical architecture (Ermentrout and Cowan, 1979; Bressloff et al., 2002). Ben Yishai et al. have argued that the local circuitry in V1 generates attractor states that amplify the response to weakly oriented visual stimuli (Ben Yishai et al., 1995). Similar attractor states are thought to be evoked by transient sensory stimuli in cortical networks that underlie working memory (Zhang, 1996; Compte et al., 2000; Brody et al., 2003). The interpretation of Kenet et al. is novel in that it suggests that even in the absence of any visual stimulation the ongoing activity in the anesthetized visual cortex wanders between a set of intrinsic cortical states (Tsodyks et al., 1999; Kenet et al., 2003; Ringach, 2003).

Motivated by the above experimental and theoretical results, we address the nature of the ongoing activity in cortex by inquiring (1) what statistical and dynamical characteristics of ongoing activity are indicative of spontaneous switching between underlying intrinsic brain attractors; (2) how the structure of the ongoing activity in V1 is related to the feed-forward architecture of the inputs to this area and to the large-scale horizontal intracortical connections; and (3) whether the experimental results regarding ongoing activity in V1 support the hypothesis of switching among cortical states or, conversely, what alternative interpretations can account for these results.

Results

Statistical Structure of Spontaneous Cortical Maps

In this work, we examine two mutually exclusive hypotheses regarding the nature of the spontaneous spatio-

*Correspondence: jgoldberg@utsa.edu

³These authors contributed equally to this work.

⁴ Present address: Department of Biology, University of Texas at San Antonio, 6900 North Loop, 1604 West, San Antonio, Texas 78249.




Figure 1. Histogram of the Raw and Spike-Triggered Similarity Indices from the Experimental Optical Imaging Data

(A) All-times histogram of the SI.

(B) Histogram of the spike-triggered SI. These histograms were pooled from 12 recording sessions from five cats (adapted with permission from Tsodyks et al., *Science* 286, 1943–1946. Copyright 1999 American Association for the Advancement of Science).

temporal activity in cortex. According to the *single state hypothesis*, the spontaneous activity represents the dynamics of a single background state. The source of activation is local noise with approximately Gaussian statistics. This noise is subsequently filtered by the underlying cortical connectivity and dynamics. According to the *multiple state hypothesis*, the spontaneous activity can be described as noise-driven spontaneous switching between different attractor states of the system. Furthermore, the set of different attractors resembles the states that are generated by the relevant sensory stimuli. In the framework of this hypothesis, we assume that the attractor states of the system reside in a low-dimensional continuous manifold. For example, in the context of V1, the putative attractor states include the cortical states evoked by oriented gratings, which can be parameterized by a single angle variable (Kenet et al., 2003).

We seek low-dimensional statistics of the high-dimensional spontaneous activity that can differentiate the single from the multiple state hypotheses. We are particularly interested in a statistic that measures the similarity between the spontaneous maps and an evoked map. Following Tsodyks et al. (1999), we use the similarity index (SI), which is the correlation coefficient between the spontaneous map and a map evoked by a drifting grating of a given orientation. An example of the distribution of the SI obtained by Tsodyks et al. (1999) is shown in Figure 1A. Another way of characterizing the spontaneous activity is to quantify to what degree the activity of a single neuron is correlated with the global SI measure. This can be achieved by measuring the distribution of the SI triggered on the spikes of a single neuron. The bias of the spike-triggered SI, defined as the mean of this distribution normalized by the standard deviation of the SI distribution, is a measure of the correlation between the spiking of the neuron and the SI (see Equation [S2] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>).

An example of an empirical distribution of the spike-triggered SI is shown in Figure 1B. Using the SI and the spike-triggered SI, we can detect three features that distinguish the two hypotheses.

(1) Dimensionality of Spontaneous Maps

In the single state hypothesis, the fluctuations of distant pixels in the spontaneous map are weakly correlated such that the set of patterns of spontaneous activity span a high-dimensional space. This high dimensionality should be reflected in a small SI. Additionally, the small correlation between the global measure SI and the activity of a single neuron should be reflected in a small bias in the spike-triggered SI. In contrast, in the multiple state hypothesis, the network wanders in a low-dimensional subspace of states, which resembles the set of evoked states. Consequently, the SI and the bias in the spike-triggered SI are expected to be high.

(2) Form of Distributions

According to the single state hypothesis, the SI averages over many independent portions of the image, and therefore its statistics should be approximately Gaussian. However, in the multiple state hypothesis, the form of the statistical distribution of the SI depends on the particular structure of the low-dimensional attractor. In some neuronal systems, the evoked states code for an angle variable (e.g., orientation of stimuli or head direction) and therefore reside on a closed contour, implying that the subset of such states is highly nonlinear. Wandering among such states could generate highly non-Gaussian statistics.

(3) Time Scale of Fluctuations

In the single state hypothesis, the characteristic time scales of the SI fluctuations should reflect the basic time constants of the synaptic currents, which are in the range of tens of milliseconds. In the multiple state hypothesis, because the noise driving the system is assumed to be local, it produces very slow wandering within the manifold of states. This should imply time scales of the SI fluctuations which are considerably longer than the basic time constants of the synaptic currents.

In order to demonstrate the properties of the above scenarios in cortical networks, the next section presents a simplified dynamic model of the spontaneous activity in V1. Depending on the parameter regime, this model is capable of generating the two statistical scenarios discussed above. Using the model, we demonstrate which properties of the network architecture and dynamics underlie the two scenarios.

A “Ring” Model of V1

In this study, we are interested in the large-scale structure of V1. Therefore, the basic component of the model is an orientation column located within a patch of cortex. We assume that in the absence of oriented stimuli the lateral geniculate nucleus (LGN) provides a stimulus-independent feed-forward drive to each column in the form of Gaussian noise with mean T , variance σ_n^2 , and a time constant τ . This noise is the source of excitation that gives rise to the ongoing spontaneous activity. During the presentation of a ψ -oriented stimulus (e.g., a drifting grating), the stimulus-evoked input from the LGN

to a column located at a point \vec{r} includes an additional oriented term $L[1 + \epsilon \cos 2(\theta(\vec{r}) - \psi)]$, where L is the overall contrast of the input and ϵ is the depth of the orientation modulation of the LGN input. $\theta(\vec{r})$ is the preferred orientation (PO, between -90° and 90°) of the LGN input to this column. In addition, the column also receives recurrent input from other cortical columns. The long-range recurrent interactions in V1 are known to be tuned for orientation (Gilbert and Wiesel, 1983; Ts'o et al., 1986; Malach et al., 1993; Weliky and Katz, 1994; Das and Gilbert, 1995; Bosking et al., 1997). Thus, we assume that the strength of the interaction between two columns depends linearly on the cosine of the difference between their POs, with amplitude given by the parameter λ (Equation [1] in the Experimental Procedures). The dynamics of the network are described by a first order threshold-linear rate model with a relaxation time constant τ_0 (Equation [3] in the Experimental Procedures).

Because optical images generated by VSD correspond to membrane potential changes (Tsodyks et al., 1999; Grinvald et al., 1999; Fitzpatrick, 2000; Sharon and Grinvald, 2002), we assume that the optical signal represents the instantaneous synaptic input. The spontaneous activity map of the V1 model at time t is the N -dimensional vector $\vec{M}^{sp}(t) = h(\vec{r}, t) - \langle h(\vec{r}, t) \rangle_t$, where $h(\vec{r}, t)$ is the total synaptic input to the orientation column at \vec{r} , consisting of the noisy LGN input as well as the recurrent cortical inputs (Equation [4] in the Experimental Procedures), and we subtract the temporal average of the spontaneous activity. The map evoked by a ψ -oriented stimulus, \vec{M}_ψ^{ev} , is defined as the temporal average of $h(\vec{r}, t)$ during the presentation of the stimulus after subtracting its spatial average. The SI of the spontaneous map with an evoked map with orientation ψ is denoted $\rho_\psi(t)$ and computed by $\rho_\psi(t) = \vec{M}^{sp}(t) \cdot \vec{M}_\psi^{ev} / \|\vec{M}^{sp}(t)\|$, where $\|\dots\|$ denotes the norm of a vector. We also compute ρ_ψ triggered on the spikes of a neuron with a PO ψ , denoted ρ_ψ^{spike} . The spikes are generated by implementing a time-varying Poisson process with a rate determined by the time evolution of the average discharge rate $m(\vec{r}, t)$ of a column with PO ψ (appearing in Equation [3] in the Experimental Procedures).

We begin the study of the nature of ongoing activity in our network model of V1 by considering a limiting case where the POs are distributed uniformly in the network. In addition, we assume that the spontaneous LGN input is spatially white noise and that, as mentioned above, the cortical interactions depend only on the POs of the interacting units. These assumptions imply that the only relevant architectural variable of a column is its PO θ . All columns can therefore be ordered in a "ring" architecture in the space of the variable θ , termed the orientation space (Figure 2A).

Evoked States in the Ring Model

The time-averaged evoked states are characterized by a hill-shaped profile in orientation space of the columns' firing rates centered about one of the columns (Figure 2A). The profile of the time-averaged input $h(\theta)$ into the columns corresponding to the hill-shaped profile of firing-rates is itself cosine shaped about the same column (Figure 2B). Each such profile corresponds to an evoked map \vec{M}_ψ^{ev} .




Figure 2. Schematic Representation of Evoked and Spontaneous Maps According to the Single and Multiple State Hypotheses

(A) In the ring model, the columns are arranged in a circle in orientation space according to their PO θ . Note that the circle ranges from 0° to 180° of orientation. The mean rate of each column is represented by its height above this circle. In the evoked state, the profile of firing rates in the network is hill shaped, centered about some orientation (here this orientation is zero).

(B) The time-averaged synaptic inputs $h(\theta)$ as a function of θ corresponding to the evoked state in panel (A). The dashed line marks the threshold.

(C) The set of evoked maps resides along a circle embedded within the two-dimensional subspace, termed the subspace of evoked maps. This subspace is spanned by two evoked maps \vec{M}_0^{ev} and $\vec{M}_{\pi/4}^{ev}$. The vector $h(\theta)$ in panel (B) corresponds to \vec{M}_0^{ev} and is marked by a green arrow. The spontaneous maps (represented by the colored vectors) reside in the full N -dimensional space. In the single state scenario, the projection of the spontaneous map (blue vector) onto the subspace of evoked maps (blue dashed line) is small, whereas its component within the other $N - 2$ dimensions is large. In the multiple state scenario, the spontaneous map (red vector) is very close to one of the evoked maps. Consequently, its projection within the subspace of evoked maps (red dashed line) is large and it has a relatively small orthogonal component within the other dimensions.

The set of all evoked maps delineates a circle within the full N -dimensional space of optical images. This circle is confined to a subspace spanned by two basis vectors, \vec{M}_0^{ev} and $\vec{M}_{\pi/4}^{ev}$, termed the subspace of evoked maps (Figure 2C).

Spontaneous States of the Ring Model

In Figure 3A, we present an analytical solution of the phase diagram of the model (see the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>), which describes the nature of the spontaneous state as a function of two parameters: (1) the gain of the orientation-modulated recurrent cortical interactions, λ , and (2) the ratio between the mean of the spontaneous LGN drive and its standard deviation, T/σ_n . The phase diagram of the ongoing activity in the network exhibits three qualitatively




Figure 3. Network States of the Ring Model

(A) Phase diagram of the model. Abscissa: λ , gain of cortical interactions. Ordinate: T/σ_n is the ratio of the mean LGN drive to its standard deviation. *H* denotes the homogeneous phase. *M* denotes the marginal phase. *I* denotes the amplitude instability phase. (B) Snapshot of the input $h(\theta)$ to the columns in the network as a function of their PO, θ , in the homogeneous phase (simulated for the point in the phase diagram in panel (A) marked by a diamond, $\lambda = 1.2$ and $T/\sigma_n = -0.5$). (C) Snapshot of the input $h(\theta)$ in the marginal phase (calculated for the point in the phase diagram marked by an asterisk, $\lambda = 1.2$ and $T/\sigma_n = 2$).

different regions denoted *H*, *M*, and *I*. The *H* region is a parameter regime where the network is in a single homogeneous state, corresponding to a flat profile in orientation space of the synaptic input $h(\theta, t)$ plus some local noise (Figure 3B). As λ is increased, the homogeneous state becomes unstable, and the network crosses over to the *M* region, which is termed the marginal phase (Ben Yishai et al., 1995). In this region, $h(\theta, t)$ has a noisy cosine-shaped profile in orientation space (Figure 3C). When the cortical gain is further increased beyond $\lambda = 2$, the system crosses into the *I* region, which is a regime of instability, where the network undergoes runaway excitation and diverges.

In order to study the statistical structure of the spontaneous maps in the regimes of the model, we simulated the network dynamics and calculated the value of the SI and the projections of the spontaneous maps onto the vectors \vec{M}_0^{ev} and $\vec{M}_{\pi/4}^{ev}$.

Spontaneous Maps in the *H* Regime

The *H* regime of the network dynamics corresponds to the single state scenario discussed above. This scenario is depicted schematically by the blue vector in Figure 2C: this vector has a small projection onto the subspace of evoked maps (dashed blue line) and a large component within the orthogonal space of all other $N - 2$ dimensions. The network simulation shows that the scatter plot of the projections on the subspace of evoked maps forms a single radial-symmetric Gaussian cloud centered about the origin (Figure 4A). The distribution of the SI is Gaussian shaped and narrow (Figure 4B). The width of the distribution scales as $1/\sqrt{N}$ (data not shown), because the SI involves averaging over N weakly correlated columns. To study the effect of the cortical gain λ , we consider the case where the bias T in the noisy input is large relative to the standard deviation σ_n of the noise, implying that the

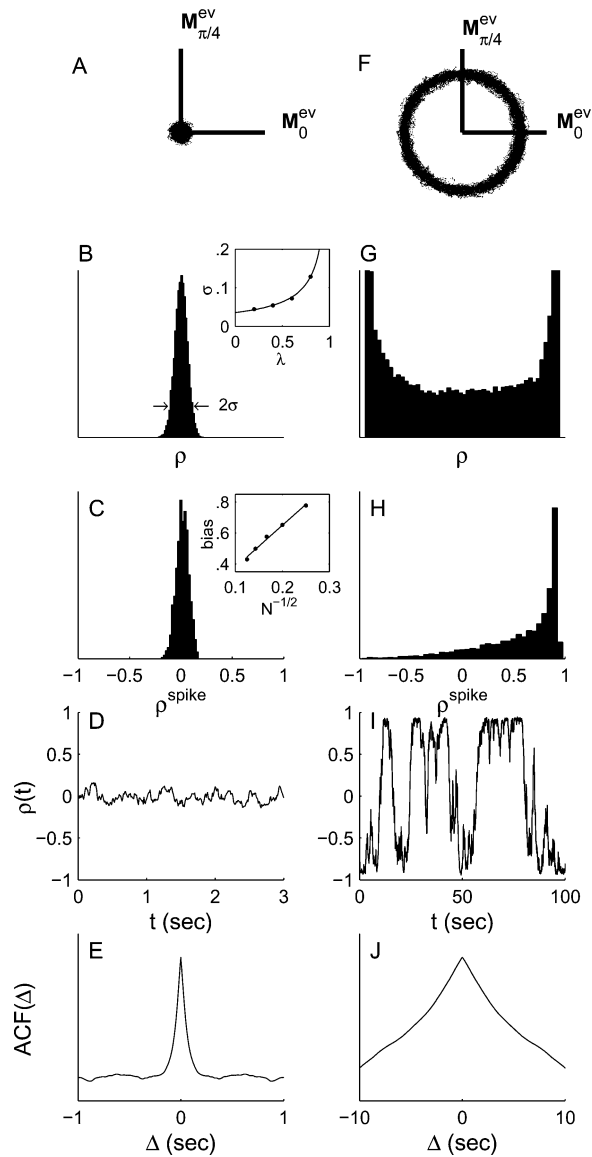


Figure 4. Spontaneous Map Fluctuations in the *H* and *M* Regimes of the Ring Model

(A–E) Fluctuations in the *H* regime (for the point in parameter space marked by a diamond in Figure 3A). (F–J) Fluctuations in the *M* regime (for the point marked by an asterisk). (A and F) Scatter plot of the projections of the spontaneous maps onto the evoked maps \vec{M}_0^{ev} and $\vec{M}_{\pi/4}^{ev}$. (B and G) Distribution of the SI ρ . (Inset) The width σ of the distribution in the linear region of the *H* regime as a function of the cortical gain λ . Solid line, analytical solution (see Equation [S6] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>); points, simulation results. $\sigma_n = 1$ and $T = 3$. (C and H) Distribution of the spike-triggered SI ρ^{spike} . (Inset) The bias in the distribution of ρ^{spike} as a function of $1/\sqrt{N}$ in the *H* regime, where N is the number of columns in the network. The bias is defined as the mean of this distribution normalized by the standard deviation of the distribution of ρ . Points, simulation results; solid line, linear fit. (D and I) Temporal evolution of the SI $\rho(t)$ as a function of time. (E and J) Autocorrelation function (ACF) of $\rho(t)$ as a function of the time lag Δ . Note the different time bases in panels (D) versus (I) and (E) versus (J).

dynamics of the columns are effectively linear. For this linear regime, we calculate analytically the width of the distribution of the SI, denoted by σ , as a function of λ (Equation [S6] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>). The width of the distribution of the SI increases with λ , as shown by this analytical calculation and verified by the network simulations (Figure 4B, inset). Due to the broadly tuned orientation-modulated structure of the cortical interactions (Equation [1] in the Experimental Procedures), increasing λ amplifies the component of the ongoing activity within the subspace of evoked maps. The bias in the distribution of ρ_{ψ}^{spike} is very small (Figure 4C), indicating a small correlation between the global SI measure and the spiking of an individual neuron, as expected in the single state hypothesis. The bias scales as $1/\sqrt{N}$ (inset), because the SI involves averaging over N weakly correlated columns, and the spiking of the neuron is strongly correlated only with its own column.

The deflections of $\rho_{\psi}(t)$ from zero last at most a few hundred milliseconds (Figure 4D). In addition, the autocorrelation function (ACF) of the SI exhibits a narrow peak with a width of ~ 50 ms (Figure 4E). The analytical calculation of the ACF in the linear regime of the model (Equation [S5] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>) demonstrates that the time scale of the fluctuations in the SI is determined by the larger of the following two time scales: that of the noise τ and that of the network $\tau_{\lambda} = \tau_0/(1 - \lambda)$. As λ is increased, the network time scale increases. However, unless λ is extremely close to unity, τ_{λ} remains on the order of tens of milliseconds. In summary, the spontaneous activity in the H regime demonstrates all the properties expected from the single state hypothesis, namely, fast and high-dimensional Gaussian fluctuations.

Spontaneous Maps in the M Regime

In this regime, the multiple state hypothesis is borne out. The noisy cosine-shaped profile of $h(\theta, t)$ is very similar to one of the evoked maps with a relatively small component within the orthogonal $N - 2$ dimensional space (red vector in Figure 2C). The local noise in the LGN input causes the location of the peak of this profile, denoted $\phi(t)$, to perform a slow random walk in orientation space, thereby distributing homogeneously in time over all orientations. This can be directly visualized by considering the highly non-Gaussian distribution of projections of the spontaneous maps onto the subspace of evoked maps (Figure 4F). This distribution outlines the set of attractors in the dynamics, which is termed the “ring attractor” (Ben Yishai et al., 1995). The SI $\rho_{\psi}(t)$ is itself proportional to $\cos 2(\psi - \phi(t))$, and so its distribution acquires the bimodality of the distribution of the cosine of a uniformly distributed angle (Figure 4G). Furthermore, the distribution of the SI in this regime is very broad and is independent of the number N of columns in the network. In addition, the bias in the distribution of ρ_{ψ}^{spike} is substantial (Figure 4H). Finally, fluctuations of the SI are slow (Figure 4I), as captured by the ACF of the SI in this regime. The random walk of $\phi(t)$ in orientation space generates an ACF with a very long time scale τ_W (Figure 4J). We show below (Figure 5G) that τ_W scales linearly with the number N of columns in the network. This is a result of the assumption about the local nature of the noise driving the system. Such noise generates very weak

fluctuations along the low-dimensional attractor manifold, with a variance of order $1/N$ (Renart et al., 2003). Thus, the system requires time that is proportional to N to generate an appreciable (order unity) drift from one attractor state to another. In conclusion, the behavior of the ongoing activity in the M regime is truly indicative of a scenario of slow wandering among a set of cortical states.

Encoding Multiple Features in V1


The ring attractor in the above model encodes a single feature, namely, orientation. We now consider a more general case in which the functional architecture of V1 gives rise to cortical states that encode multiple features (e.g., orientation, direction, spatial frequency, etc.). We distinguish two scenarios in which V1 may represent these features: a unary representation or a combinatorial representation of multiple features.

Unary Representation of Multiple Features

In this scenario, different features are encoded by distinct manifolds of states, hence each intrinsic cortical state represents only one feature. As an example, we construct a network architecture that generates two separate ring attractors (see Experimental Procedures) representing two angle-like features: orientation, denoted θ^o , and direction, denoted θ^d (Figure 5A). For the multiple manifolds to have appreciable effect on the nature of the fluctuations of the spontaneous activity, the system must have a mechanism for spontaneous transitions between these manifolds. These transitions involve global changes in the instantaneous state and thus are only rarely triggered by the local noise present in our system. For this reason, we added a global noise term to the dynamics to cause the network to alternate between states of θ^o and states of θ^d (see Experimental Procedures). We measure the similarity of an instantaneous spontaneous state of the network to the orientation and direction attractors by the two similarity indices ρ_o and ρ_d , respectively. Simulating this model (in the parameter regime where the two attractor manifolds exist) and calculating the time evolution of the corresponding SIs demonstrates that the system switches between manifolds: during times in which one SI attains a large absolute value, the other is close to zero (Figure 5B). The reason one SI is small while the other is large is because the states in the direction attractor are orthogonal to the orientation ones. The resulting distribution of ρ_o is thus trimodal (Figure 5C). If yet more features are encoded in the network, the resulting distribution of the SI with a given feature will have essentially the same shape. However, the central peak will be larger as the network spends more time occupying states of the other features.

Combinatorial Representation of Multiple Features

In this scenario, a cortical state can simultaneously encode several features (Basole et al., 2003), and all possible combinations of features are represented by a continuous multi-dimensional attractor manifold. As before, we begin by considering only one feature, say spatial frequency, in addition to orientation. Following Bressloff and Cowan (Bressloff and Cowan, 2002, 2003), we model the V1 interactions such that the geometry of the two-feature attractor manifold is that of a surface of a three-dimensional sphere (see Experimental Procedures), with the longitudinal angle θ representing orientation and the latitudinal angle ϕ spatial frequency (Figure 5D). The phase diagram of this model



(H) Analytical solution of the probability density function (pdf) of the SI, ρ , in the M regime of a $d + 1$ sphere of cortical states encoding d features at the large N limit (see Equation [S7] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>). The number of features is indicated below each pdf.

is identical to that of the ring model (Figure 3A). However, the statistics of the spontaneous fluctuations are affected by the spherical structure of the attractor manifold in the M regime. This is evident from the projection of the spontaneous maps onto the subspace of evoked maps of orientation, i.e., the equatorial plane (Figure 5E). The resulting distribution of the SI with any given state of the network is essentially flat and spans nearly the whole range between -1 and 1 (Figures 5F). Thus, the additional feature slightly reduces the width of the distribution of the SI (compare Figures 5F and 4G). As expected in general for attractor manifolds both in the ring model and in the three-dimensional spherical model, the time scale τ_w of the ACF scales linearly with N (Figure 5G). Nevertheless, the slope is smaller in the latter model (crosses) than in the former (points), because the network state has an additional dimension to wander along and can thus drift away faster from any given state.

Next we consider an attractor manifold that encodes an arbitrary number d of features in a combinatorial fashion, corresponding to a d -dimensional sphere. In this case, the distribution of the SI can be calculated analytically in the limit of large N (Equation [S7] in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>). For d equals 1 or 2, we recover the bimodal and flat distributions of the ring and three-dimensional spherical models, respectively. As the number of features is increased, the distribution becomes narrower and more Gaussian-like (Figure 5H). The time scale τ_w of wandering among states is proportional to N/d and thus decreases with the number of features encoded on the sphere. This is due to the variance of the projection

Figure 5. Encoding Multiple Features in V1 (A–C) The unary multiple feature scenario. (D–H) The combinatorial multiple feature scenario.

(A) The cortical states of orientation and direction are arranged on two distinct rings. (B) The temporal evolution of the SIs of the network state with a state of orientation (ρ_o , thin line) and a state of direction (ρ_d , thick line). (C) The distribution of ρ_o . (D) The cortical states of orientation and spatial frequency are arranged on a sphere, such that the longitudinal angle θ represents orientation and the latitudinal angle ϕ represents spatial frequency. (E) Scatter plot of the projections of the spontaneous maps onto the evoked maps M_0^{ev} and $M_{\pi/4}^{ev}$ in the M regime, corresponding to the point in parameter space marked by an asterisk in Figure 3A. (F) Distribution of the SI, ρ , in the M regime (for the point in parameter space marked by an asterisk in Figure 3A). (G) The decay time τ_w of the ACF as a function of N , in the M regimes of the ring model (points) and the three-dimensional spherical model (crosses), calculated from network simulations. Solid lines, linear fit. τ_w was extracted by fitting an exponential decay to the ACF for lags between 0 and 0.5 s.

of the local noise onto the surface of the sphere that increases as d/N with the dimensionality of the sphere.

Two-Dimensional Model of V1

In this section, we extend our study to a 2D model. This allows us to consider several hitherto neglected aspects. First, our analysis of the network size dependence of the variance and the time scale of the SI fluctuations were based on the assumption that the noise driving the network was spatially uncorrelated. Using a 2D model, we will study the important effects of spatial correlations in the input from the LGN. Additionally, a 2D model will allow us to incorporate spatially restricted cortical interactions. For the 2D model, we revert to cortical interactions that encode a single feature of orientation and use a cortical map of POs $\theta(\vec{r})$ which was extracted from the optical imaging data (Figure 6A, courtesy of T. Kenet, A. Arieli, and A. Grinvald). This map corresponds to a $3 \text{ mm} \times 3 \text{ mm}$ portion of an empirical “angle map” (Bonhoeffer and Grinvald, 1993). This size is within the range used in recent studies of ongoing cortical activity and contains about 15 pinwheels (Muller et al., 2000). The observed cortical angle map shows a nonuniform distribution of POs (Figure 6B), with an overrepresentation of the orientation $\theta = 0$. In order to investigate the properties of multiple orientation states, we performed small local adjustments on $\theta(\vec{r})$ of Figure 6A to preserve a uniform distribution of POs across this portion of the map (see Experimental Procedures). This transformation mimics a possible local compensation mechanism (e.g., by varying the local synaptic strength)




Figure 6. An Empirical Two-Dimensional Map of Preferred Orientations

(A) A color-coded angle map of POs (θ between -90° and 90°) derived from the optical imaging of a $3 \text{ mm} \times 3 \text{ mm}$ region of the primary visual cortex (V1) of a cat (courtesy of T. Kenet, A. Arieli, and A. Grinvald).

(B) The distribution of POs derived from the angle map depicted in panel (A).

that acts to counterbalance the above-mentioned anisotropy (Renart et al., 2003). In the next section, we will study the 2D model with the measured anisotropy.

We first discuss the effects of correlations in the LGN inputs, keeping the interactions between columns in the map solely dependent on their POs (Equation [1] in Experimental Procedures). Correlations in the LGN input are expected to be generated by the overlap of the axonal arborizations of geniculocortical neurons (Gilbert and Wiesel, 1979; Humphrey et al., 1985) and the overlap in the dendritic arborizations of V1 neurons (Braitenberg and Schuz, 1991; Kossel et al., 1995; Tyler et al., 1998). We assume that these correlations are spatially isotropic and model the noisy LGN input as spatially white Gaussian noise filtered by a two-dimensional Gaussian kernel with a length scale ξ (Equation [5] in Experimental Procedures). For a uniform distribution of POs, we find that the full model still exhibits the *H*, *M*, and *I* regimes described above for the ring architecture (Figure 3A).

The Single State *H* Regime

The main effect of the correlations of the noisy drive from the LGN is to control the number of independent degrees of freedom (DOF) in the system, denoted as N_{eff} . This is seen by considering the distribution of the SI in the *H* regime. In Figure 7A, we display the width of the resulting distribution of the SI, denoted σ , as a function of the spatial scale of the kernel, ξ , for the simple case where the cortical interactions are turned off, $\lambda = 0$. For uncorrelated noise; i.e., very small ξ , N_{eff} equals the number of columns, N , therefore for vanishing ξ the width of the distribution is very small. As ξ is in-




Figure 7. The 2D Model—Uniform Distribution of POs

(A) The width of the distribution of the SI ρ , denoted σ , in the single state regime for $\lambda = 0$ as a function of the length-scale ξ of the correlations in the LGN input.

(B) Dependence of the width of distribution of the SI σ on the strength of cortical interactions λ . Here and for the remainder of the simulations of the full model, we chose $\xi = 120 \mu\text{m}$ and $T/\sigma_n = 2$.

(C) A spontaneous map occurring in a simulation of the full model for $\lambda = 0.6$.

(D) An evoked map of zero orientation. The SI with the spontaneous map in panel (C) is 0.69.

(E) The histogram of the SI ρ for $\lambda = 0.6$.

(F) Distribution of ρ^{spike} for an ordinary column in the network.

(G) Tuning curve (arbitrary units) of an ordinary column in the network (dash-dotted line) and that of a “selective” column whose afferent cortical interactions were strengthened by a factor of 8 (solid line).

(H) Distribution of ρ^{spike} for the “selective” column.

(I) The decay time τ_w of the ACF as a function of ξ in the multiple state regime.

(J) Color-coded profile of cortical interactions with a spatial falloff impinging on a column located at the center of the patch. Length-scale of spatial falloff is 0.3 mm . red, excitation; dark blue, inhibition; light blue, zero.

creased, N_{eff} decreases, since it scales roughly as $N_{eff} \propto Area/\xi^2$, where $Area$ is the total area of the system. By analogy to the case of the independent columns, we expect that in the presence of these short-range correlations σ will be on the order of $1/\sqrt{N_{eff}}$. Indeed, we find a linear increase of σ with ξ for small values of ξ . However, for larger values of ξ , the fluctuating spatial patterns in the LGN input are only composed of low-frequency spatial modes. In contrast, the power spectrum of the evoked map has a band-pass structure due to the regularity in the spatial layout of pinwheels (Muller et al., 2000). Therefore, the correlation between the maps decreases for large values of ξ (Figure 7A). In summary, in the absence of cortical interactions, the distribution of the SI attains a maximal width of $\sigma \sim 0.11$, for the area of V1 that we are modeling, at the optimal spatial scale of LGN-driven noise correlations of $\xi \sim 200 \mu\text{m}$. As in the single state of the ring model (Figure 4B, inset), the distribution of the SI in the single state of the 2D model broadens as λ is increased (Figure 7B, which is calculated for $\xi = 120 \mu\text{m}$). However, the width of the distributions in the 2D model is larger, reflecting the reduction of the number of DOF from $N = 784$ in the case of spatially white noise to $N_{eff} \sim 100$ in the presence of LGN correlations. In Figure 7C, we depict a snapshot of a spontaneous map that occurred in a network simulation with $\lambda = 0.6$ and resembled the evoked map (Figure 7D). The distribution of the SI in this simulation is shown in Figure 7E, displaying a broad Gaussian-like shape. Finally, the temporal fluctuations of the SI are fast in this regime generating an ACF (data not shown) very similar to the one calculated in the H regime of the ring model (Figure 4E).

The distribution of ρ_{ψ}^{spike} in this regime displays only a small bias (Figure 7F), equal roughly to $1/\sqrt{N_{eff}} \approx 0.1$. However, because the empirical ρ_{ψ}^{spike} measures the correlation between a global measure (e.g., the SI) and the activity of a single neuron, it is likely to be sensitive to the properties of this neuron, notably, its selectivity to orientation. Due to the diversity in the degree of orientation selectivity displayed by V1 neurons (Ringach et al., 2002; Monier et al., 2003), we hypothesized that calculating the SI for a highly selective neuron should generate a substantial bias in the distribution of ρ_{ψ}^{spike} . In order to test this hypothesis in the model, we increased the selectivity of a column whose PO is ψ by strengthening the orientation modulated recurrent connections to this column. Amplifying these connections indeed amplified the tuning curve of this column (calculated for oriented inputs) relative to that of an ordinary column in the network (Figure 7G), thereby increasing its selectivity. Due to the structure of these connections (Equation [1] in the Experimental Procedures), the recurrent feedback to this column during the ongoing activity is proportional to ρ_{ψ} . Therefore, strengthening the feedback to this column should amplify the correlation between ρ_{ψ} and the firing rates of neurons in this column and thus generate a large bias in the distribution of ρ_{ψ}^{spike} . As expected, the amplification of recurrent connections generated a substantial bias in the distribution of ρ_{ψ}^{spike} (Figure 7H). The mean spontaneous firing rate of this column was also 40% higher than the rest of the population. Note that in these simulations, we increased the recurrent

input to a whole column because it is the basic element in the model. However, these results would still hold if the input to a single neuron, rather than a whole column, were strengthened.

The Multiple State M Regime

The behavior of the network dynamics in the M regime of the full model is similar to that of the M regime of the ring model, except that the time scale of wandering among states τ_w depends on the spatial scale ξ of the LGN correlations. As we saw in the single state regime, the LGN input has the strongest correlation with the spatial structure of the evoked map when its spatial scale is $\sim 200 \mu\text{m}$. It is therefore expected that input from the LGN with this scale will be most efficient at driving the network from one orientation state to another, leading to a minimal time scale τ_w as a function of ξ . This is confirmed by network simulations (Figure 7I), which show that the smallest attainable time scale in this regime, for the area of V1 that we are modeling, is ~ 750 ms.

Spatially Restricted Cortical Interactions

In all models described above, the coupling between columns was modulated solely by their preferred features and was thus independent of the cortical distance between them. To examine the effect of spatial restrictions of the cortical interactions, we introduced a spatial falloff to the orientation-modulated interactions in the 2D model by multiplying them with a spatially isotropic Gaussian kernel with a length scale κ (Equation [2] in the Experimental Procedures). We found that for spatial scales as small as $\kappa = 300 \mu\text{m}$ (Figure 7J), each column receives recurrent input from columns with an effectively uniform distribution of POs leading to a negligible effect on the dynamics in both the single state and multiple state regimes. For smaller spatial scales, the interactions gradually lose their orientation-modulated structure, thereby disrupting the ring attractor and generating spontaneous maps that no longer resemble evoked ones.

Nonuniform Distribution of POs

The existence of a ring attractor in the dynamics relies on the perfect uniformity of the distribution of POs. In order to study the effects of a nonuniform distribution, we invoke the empirical angle map with its nonuniform distribution of POs (Figure 6). Because the distribution of the POs is not uniform, the network dynamics do not support a marginal phase with a ring attractor even for large values of λ . However, for sufficiently large values of λ , two to three stable states may appear (Renart et al., 2003), depending on the form of the distribution of the POs. For the value of $T/\sigma_n = 2$ that we used in the network simulations, a second stable state appears as λ is increased above 1.15.

The nonuniformity of the distribution of the POs has several effects on the single state regime. First, the long time average of the pattern of inputs $h(\vec{r}, t)$ is not homogeneous but displays a peak in orientation space near $\theta = 0$ (data not shown), due to the overrepresentation of this angle in the distribution of POs (Figure 6B). Second, the scatter-plot of the projections of the spontaneous maps onto the subspace of evoked maps is oval shaped (Figure 8A), indicating that the fluctuations are mildly anisotropic in this subspace: the fluctuations parallel to M_0^{ev} are broader than those parallel to $M_{\pi/4}^{ev}$, due




Figure 8. The 2D Model—Nonuniform Distribution of POs
(A) The scatter plot of the projections of the spontaneous maps onto the evoked maps M_0^{ev} and $M_{\pi/4}^{ev}$ for $\lambda = 0.6$.
(B) Dependence of σ on the orientation ψ of the evoked map for $\lambda = 0.6$.
(C) The distribution of ρ for $\lambda = 1.5$.

to the overrepresentation of $\theta = 0$. This anisotropy is subsequently reflected in the dependence of the width of the distribution of the SI on the orientation ψ of the evoked map used to calculate it (Figure 8B). As λ is increased into the two-state regime, the high gain amplifies the effect of the nonuniformity of the distribution of POs and generates a highly non-Gaussian SI distribution (Figure 8C).

Discussion

We have presented several network models of the large-scale architecture and ongoing dynamics of V1. The idealized ring model exhibited two qualitatively different regimes of ongoing activity as a function of the effective gain in the network: a single state regime, characterized by high-dimensional Gaussian-like fluctuations of synaptic activity around a single background state; and a multiple state regime, in which the system wanders between many attractor states that form a ring-shaped manifold. The ring attractor encodes a single feature, namely, orientation. Next, we considered generalizations of the ring attractor to attractors that encode multiple visual features in either a unary or a combinatorial fashion (Basole et al., 2003). Finally, the full 2D model, which included spatial correlations in the LGN input and spatial decay of cortical interactions, also exhibited both regimes when the distribution of the POs was uniform. We found however that even a slight nonuniformity in

the distribution of POs replaced the ring attractor with only a small number of states (Renart et al., 2003).

The seminal studies of Grinvald and colleagues provide a novel perspective on the study of the structure of ongoing activity in V1 (Tsodyks et al., 1999; Kenet et al., 2003). In light of the present models, we are in a position to inquire which of the scenarios—the single state or the multiple state—better describes their empirical results concerning the statistics of the ongoing activity.

Time Scale of the SI Fluctuations

The temporal fluctuations of the empirical SI are fast, on the order of ~ 80 ms (Kenet et al., 2003). This is consistent with the single state regime where the fluctuations are on the order of tens to one hundred milliseconds, reflecting the basic time constants of the dynamics, moderately increased by the cortical connections. On the other hand, in the multiple state regime, the time scale of the fluctuations is long and scales with the size of the network. The presence of correlations in the LGN input dictates a time scale that is on the order of a second long (Figure 7I), for a system whose size is $3 \text{ mm} \times 3 \text{ mm}$. If the system were larger (e.g., the whole V1), the time scale could be many seconds long. In the combinatorial multiple feature scenario, which generalizes the ring model to a higher dimensional sphere, the time scale of the fluctuations decreases as the reciprocal of the number of features. This means that in order to attain the empirical time scale the number of features required is at least ten.

The Shape of the Distribution of the SI

The empirical distribution of the SI is Gaussian-shaped (Figure 1A), which agrees with the predictions of the single state scenario with moderate cortical feedback (Figures 4B and 7E). In contrast, the multiple state ring attractor scenario predicts a bimodal distribution (Figure 4G). In the combinatorial multiple feature scenario, the distribution becomes Gaussian-like as the number of encoded features is increased above five (Figure 5H). In the Supplemental Data (<http://www.neuron.org/cgi/content/full/42/3/489/DC1>), we show that the predicted shape of the distribution of the SI in the different scenarios is relatively insensitive to the measurement noise.

We have shown that in the unary multiple features scenario the statistics of the ongoing fluctuations resemble that of a “mixture density”: at any given instance, the instantaneous map originates either from a manifold resembling orientation maps or from a manifold of states which code for other features that are uncorrelated with the orientation maps. This behavior results in an SI distribution with a pronounced non-Gaussian shape (Figure 5C). Another “mixture” scenario suggested by Kenet et al. (2003) is that the system spends most of the time in a single background state but occasionally switches to one of multiple cortical states (perhaps due to nonstationarity in the cortical gain). We modeled this scenario in the isotropic model by adding a global noise term to the dynamics that causes the network to alternate between the single state regime and the multiple state regime. The resulting distribution was trimodal, resembling the one generated in the unary multiple feature

scenario (Figure 5C). The narrow central peak is due to the time spent in the single state regime, and the remainder of the distribution is due to the time spent in the multiple state regime. In conclusion, the “mixture” scenarios cannot account for the Gaussian-shaped empirical distribution.

The Width of the Distribution of the SI

the empirical width is up to ~ 0.2 (Figure 1A). The single state of the ring model predicts a very narrow width, on the order of $1/\sqrt{N}$, where N is the number of imaged columns. Introduction of spatial correlations in the LGN input yields an effective number N_{eff} of DOF considerably smaller than N . This, combined with a moderate cortical feedback, yields SI distributions in the single state regime with widths similar to the empirical ones (Figure 7E). Cortical interactions encoding more features would narrow the distribution further in the single state regime, requiring stronger cortical gain (λ) to reproduce the empirical width. On the other hand, the multiple state regime of either the single or the unary multiple feature models predicts a width that spans the full range of -1 to 1 (Figures 4G and 5C). In the combinatorial multiple feature model, the distribution becomes narrower as the number of encoded features is increased (Figure 5H), yielding a width of ~ 0.2 for over ten features. In the Supplemental Data (<http://www.neuron.org/cgi/content/full/42/3/489/DC1>), we show that measurement noise reduces the width of the distribution of the SI. However, in light of the high signal-to-noise ratio attained with the improved VSDs used by Grinvald and colleagues (Grinvald et al., 1999; Shoham et al., 1999), we expect the effect of measurement noise to be small.

Importantly, the single state and the multiple state scenarios make very different predictions regarding the dependence on image size. In the single state regime, increasing the area of the imaged patch increases the number of independent degrees of freedom in the image and thus reduces the width of the SI. In contrast, in the multiple state scenario, in which the states are assumed to be low dimensional and global, the width of the SI is predicted to be independent of the size of the imaged patch.

Bias in the Distribution of the Spike-Triggered SI

Our analysis predicts a small bias in the distribution of $\rho_{\psi}^{\text{spike}}$, on the order of $1/\sqrt{N_{\text{eff}}}$ in the single state regime, in contrast to a substantial bias in the multiple state regime (Figure 4H). Indeed, the bias in the empirical distribution is substantial (Figure 1B). However, this bias also depends on the degree of orientation selectivity of the neuron used to calculate it. Thus, in the single state of the 2D model, a neuron with a large modulation depth of response to oriented stimuli that also exhibits a high spontaneous rate can generate a bias comparable to the empirical one (Figure 7H). Indeed, in their work, Tsodyks et al. indicated that the spike-triggered distribution of the SI was calculated for neurons that both exhibited a high spontaneous firing rate and were strongly selective to orientation (Tsodyks et al., 1999). In summary, the single state hypothesis predicts that the majority of the neurons will generate relatively small biases in the

distribution of $\rho_{\psi}^{\text{spike}}$ and, furthermore, that this bias will decrease if ongoing maps of a larger size are used.

Nonuniformity of the Distribution of the POs

In order for the model to support a multiple state scenario with a continuous attractor, the distribution of preferred features needs to be fine tuned. Our study indicates, as does the recent study by Renart et al., that even a slight inhomogeneity in the distribution of POs disrupts the “ring attractor” and replaces it with a small number of states (Renart et al., 2003). This raises the question of whether and to what degree the distribution of POs in V1 is uniform. In the cat, early studies found fairly uniform distributions (Bonhoeffer and Grinvald, 1993; Muller et al., 2000), whereas recent studies have found an overrepresentation of the horizontal and vertical orientations (Li et al., 2003; Wang et al., 2003). Additionally, the cortex may possess mechanisms to compensate for possible deviations from uniformity (Renart et al., 2003). The single state regime of the 2D model predicts that deviations from uniformity should be observable in the statistics of the ongoing activity: (1) the temporal mean of the ongoing activity should have a finite projection on the subspace of evoked maps; and (2) the width of the distribution of the SI should change as a function of the orientation of the stimulus used to generate it (Figure 8B). This latter prediction was in fact borne out in the recent study of Kenet et al. (2003).

Other Statistical Measures

According to the multiple state scenario, a large portion of the variance in the spontaneous maps resides within a low-dimensional subspace. In contrast, according to the single state scenario, the total amount of variance in the spontaneous maps that is attributable to fluctuations within this subspace is small, less than 10% in our 2D model (data not shown). The single state hypothesis therefore predicts that applying the principal component method for uncovering a low-dimensional linear structure in a set of high-dimensional vectors (Jolliffe, 1986) to the set of spontaneous maps will not single out the subspace of evoked maps. On the other hand, we have found that the self-organized feature map (FM) algorithm used by Kenet et al. (2003) is not a sensitive measure of this issue, yielding similar results for both the single state and the multiple state scenarios. The FM algorithm detects the dimensions in the data with maximal variance, even when this variance is small in comparison with the total variance (see details in the Supplemental Data at <http://www.neuron.org/cgi/content/full/42/3/489/DC1>). In this work, we focused on the implication of the system’s architecture and gain on the statistics of the ongoing activity. Analyzing the statistics of the evoked activity requires further modeling assumptions about the nature of the stimulus related inputs. This issue is addressed in the Supplemental Data.

In summary, we considered four scenarios for ongoing activity in V1: the single state, the ring attractor, the combinatorial multiple feature scenario, and the mixture scenario. Of these, only the single state scenario and the combinatorial multiple feature scenario, encoding at least ten features, are consistent with the empirical time scales of the fluctuations and the Gaussian-shaped

distribution of the SI. We propose that these scenarios can be distinguished by studying how the width of the SI distribution is affected by the size of the area of V1 being imaged: only in the single state scenario does this width decrease with the size of the imaged area. In addition, the combinatorial multiple state scenario, in which several features are represented by a continuous attractor manifold, requires substantial fine tuning. The degree of fine tuning required increases as the number of multiple features is increased. The overrepresentation of cardinal orientations observed recently in the ongoing activity (Kenet et al., 2003) indicates a lack of such fine tuning.

Experimental Procedures

V1 Model Details

Architecture

The strength of interaction between orientation columns located at points \vec{r} and \vec{r}' on the cortical surface is modeled as

$$J_{\vec{r},\vec{r}'} = \frac{2\lambda}{N} \cos 2(\theta(\vec{r}) - \theta(\vec{r}')) \quad (1)$$

where λ represents the strength of the cortical interactions and N is a normalization by the number of columns in the network. In the network simulations we used $N = 784$. $\theta(\vec{r})$ is the PO of the column located at point \vec{r} .

In the 2D model, we also considered an architecture that modulates these interactions (Equation [1]) as a function of the distance on the cortical surface between the interacting columns

$$J_{\vec{r},\vec{r}'} = 2\lambda A(\|\vec{r} - \vec{r}'\|) \cos 2(\theta(\vec{r}) - \theta(\vec{r}')) \quad (2)$$

where $A(x) \propto \exp(-x^2/2\kappa^2)$ in which κ is the length-scale of the interactions. $A(\|\vec{r} - \vec{r}'\|)$ is normalized such that the sum over \vec{r}' with periodic boundary conditions is unity.

Empirical Map of POs

The map $\theta(\vec{r})$ of the POs appearing in Figure 6A is a portion of an empirical "angle map" (Bonhoeffer and Grinvald, 1993). To generate this map, we used eight maps evoked in V1 of a cat by the presentation of eight gratings of equally spaced orientations (courtesy of T. Kenet, A. Arieli, and A. Grinvald). The PO of each pixel was determined by fitting a cosine to the modulation of its value across the evoked maps. The angle map was coarse grained to contain 28×28 pixels corresponding to a $3 \text{ mm} \times 3 \text{ mm}$ area of cortex. Strictly speaking, this empirical map does not represent the PO of the LGN input because presumably it is affected by cortical interactions. However, in our model, it can be shown that the differences between the map of POs of the LGN input and that of the output are small, and we therefore neglect them.

In order to generate a ring attractor in the full model, we homogenized the nonuniform distribution of POs (Figure 6B) as follows. Let $\{\theta_k\}_{k=1}^N$ be the set of all empirical values of POs sorted from smallest to largest. We generated the homogenized map by replacing the PO θ_k with $(2k/N - 1) \times 90 \text{ deg}$; $k = 1, \dots, N$, while retaining its original location on the two-dimensional map. The correlation coefficient between the original and transformed angle map is 0.999.

Dynamics

The mean firing rate of neurons in the column located at point \vec{r} , at a given time t , is denoted by $m(\vec{r}, t)$. We assume that the dynamics of $m(\vec{r}, t)$ are governed by the following equation

$$\tau_0 \frac{d}{dt} m(\vec{r}, t) = -m(\vec{r}, t) + [h(\vec{r}, t)]_+; \quad [x]_+ = \max(x, 0). \quad (3)$$

In these dynamics, the temporal evolution of $m(\vec{r}, t)$ is determined by a linear decay term and a threshold-linear activation function of the input $h(\vec{r}, t)$ to this column; τ_0 is the rate relaxation time constant of the columns, set to 10 ms. Note that the overall scale of $m(\vec{r}, t)$ is arbitrary. During the ongoing activity, the input $h(\vec{r}, t)$ to a given column is composed of two contributions: (1) a feed-forward input

from the LGN, denoted $\eta(\vec{r}, t)$, which is Gaussian noise with mean T and variance σ_n^2 . The noise is temporally correlated with an exponential decay and a time constant τ set to 50 ms; (2) the recurrent feedback from the other columns in V1, mediated by a matrix J of interactions (Equation [1] or [2]). Thus, the expression for the input to the column located at \vec{r} during the ongoing activity is given by

$$h(\vec{r}, t) = \sum_{\vec{r}'} J_{\vec{r},\vec{r}'} m(\vec{r}', t) + \eta(\vec{r}, t) \quad (4)$$

Spatial Correlations in the LGN Input

In the 2D model, spatial correlations were generated in the LGN input by filtering spatially white Gaussian noise through the spatially isotropic filter

$$F_{\vec{r},\vec{r}'}^{\text{LGN}} \propto \exp\left(-\frac{\|\vec{r} - \vec{r}'\|^2}{2\xi^2}\right) \quad (5)$$

normalized such that

$$\sum_{\vec{r}'} (F_{\vec{r},\vec{r}'}^{\text{LGN}})^2 = 1$$

so that the total variance of the input into each column is preserved. ξ is the length-scale of the correlations. We used periodic boundary conditions for the filtering.

The Unary Multiple Feature Scenario

To generate this scenario, we assume that (1) each cortical column, in addition to its PO, now denoted θ^d , has also a preferred direction (PD), denoted θ^d ; (2) the PDs are uniformly distributed (between -180° and 180°) and the two preferred angles are spatially uncorrelated; and (3) the cortical interactions include an additional term that is modulated by the difference in PDs of the interacting columns

$$J_{\vec{r},\vec{r}'} = \frac{2\lambda}{N} [\cos 2(\theta^d(\vec{r}) - \theta^d(\vec{r}')) + \cos(\theta^d(\vec{r}) - \theta^d(\vec{r}'))] \quad (6)$$

To induce alternations between states of θ^d and states of θ^d , we replaced the constant bias T in the LGN input with a spatially uniform Gaussian process. In the simulation (Figures 5B and 5C), we used a process with a mean of 2.5, a standard deviation of 5, and a time scale of 100 ms. The standard deviation of the local noise σ_n was 1 and $\lambda = 1.8$.

The Three-Dimensional Spherical Model

In the network simulations, each column was assigned two angles: θ_j and ϕ_k . θ_j is spaced equally between -90° and 90° of orientation (i.e., $\theta_j = -\pi j/n$; $j = -n/2 + 1, \dots, n/2$, where $n = \sqrt{N}$ is an integer). The angle ϕ_k (between 0° and 180°), which represents spatial frequency (Bressloff and Cowan, 2002, 2003), was chosen as follows

$$\phi_k = \arccos\left[1 - \frac{2}{n}\left(k - \frac{1}{2}\right)\right]; \quad k = 1, \dots, n.$$

This choice approximates a homogeneous density of states on the surface of the sphere. The interaction between the column with indices (j, k) and the one with indices (j', k') is given by

$$J_{j,k,j',k'} = \frac{3\lambda}{N} [\sin \phi_k \sin \phi_{k'} \cos 2(\theta_j - \theta_{j'}) + \cos \phi_k \cos \phi_{k'}]. \quad (7)$$

Acknowledgments

This research was partially supported by a grant from the US-Israeli BSF. J.A.G. and U.R. were supported by the Yeshaya Horowitz Association. We thank T. Kenet, A. Arieli, M. Tsodyks, and A. Grinvald for providing us with their optical map data (Figure 6A). We are also grateful to them for many valuable discussions and clarifications concerning their experimental results. We thank Y. Loewenstein for his critical reading of an earlier version of the manuscript.

Received: July 29, 2003

Revised: February 24, 2004

Accepted: March 18, 2004

Published: May 12, 2004

References

- Arieli, A., Shoham, D., Hildesheim, R., and Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J. Neurophysiol.* *73*, 2072–2093.
- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* *273*, 1868–1871.
- Basole, A., White, L.E., and Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. *Nature* *424*, 986–990.
- Ben Yishai, R., Bar-Or, R.L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA* *92*, 3844–3848.
- Bonhoeffer, T., and Grinvald, A. (1993). The layout of iso-orientation domains in area 18 of cat visual cortex: optical imaging reveals a pinwheel-like organization. *J. Neurosci.* *13*, 4157–4180.
- Bosking, W.H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J. Neurosci.* *17*, 2112–2127.
- Braitenberg, V., and Schuz, A. (1991). *Anatomy of the Cortex: Statistics and Geometry* (Berlin: Springer-Verlag).
- Bressloff, P.C., and Cowan, J.D. (2002). SO3 symmetry breaking mechanism for orientation and spatial frequency tuning in the visual cortex. *Phys. Rev. Lett.* *88*, 078102. Published online January 31, 2002.
- Bressloff, P.C., and Cowan, J.D. (2003). A spherical model for orientation and spatial-frequency tuning in a cortical hypercolumn. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *358*, 1643–1667.
- Bressloff, P.C., Cowan, J.D., Golubitsky, M., Thomas, P.J., and Wiener, M.C. (2002). What geometric visual hallucinations tell us about the visual cortex. *Neural Comput.* *14*, 473–491.
- Brody, C.D., Romo, R., and Kepecs, A. (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr. Opin. Neurobiol.* *13*, 204–211.
- Compte, A., Brunel, N., Goldman-Rakic, P.S., and Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* *10*, 910–923.
- Contreras, D., and Steriade, M. (1995). Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. *J. Neurosci.* *15*, 604–622.
- Das, A., and Gilbert, C.D. (1995). Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* *375*, 780–784.
- Destexhe, A., Contreras, D., and Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J. Neurosci.* *19*, 4595–4608.
- Ermentrout, G.B., and Cowan, J.D. (1979). A mathematical theory of visual hallucination patterns. *Biol. Cybern.* *34*, 137–150.
- Fitzpatrick, D. (2000). Cortical imaging: capturing the moment. *Curr. Biol.* *10*, R187–R190.
- Gilbert, C.D., and Wiesel, T.N. (1979). Morphology and intracortical projections of functionally characterized neurones in the cat visual cortex. *Nature* *280*, 120–125.
- Gilbert, C.D., and Wiesel, T.N. (1983). Clustered intrinsic connections in cat visual cortex. *J. Neurosci.* *3*, 1116–1133.
- Grinvald, A., Shoham, D., Shmuel, A., Glaser, D.E., Vanzetta, I., Shtoyerman, E., Slovlin, H., Wijnbergen, C., Hildesheim, R., Sterkin, A., and Arieli, A. (1999). In-vivo optical imaging of cortical architecture and dynamics. In *Modern Techniques in Neuroscience Research*, U. Windhorst and H. Johansson, eds. (Heidelberg: Springer-Verlag), pp. 894–969.
- Humphrey, A.L., Sur, M., Uhlrich, D.J., and Sherman, S.M. (1985). Projection patterns of individual X- and Y-cell axons from the lateral geniculate nucleus to cortical area 17 in the cat. *J. Comp. Neurol.* *233*, 159–189.
- Jolliffe, I.T. (1986). *Principal Component Analysis* (New York: Springer-Verlag).
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* *425*, 954–956.
- Kossel, A., Lowel, S., and Bolz, J. (1995). Relationships between dendritic fields and functional architecture in striate cortex of normal and visually deprived cats. *J. Neurosci.* *15*, 3913–3926.
- Lampl, I., Reichova, I., and Ferster, D. (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* *22*, 361–374.
- Li, B., Peterson, M.R., and Freeman, R.D. (2003). Oblique effect: a neural basis in the visual cortex. *J. Neurophysiol.* *90*, 204–217.
- Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc. Natl. Acad. Sci. USA* *90*, 10469–10473.
- Monier, C., Chavane, F., Baudot, P., Graham, L.J., and Fregnac, Y. (2003). Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* *37*, 663–680.
- Muller, T., Stetter, M., Hubener, M., Sengpiel, F., Bonhoeffer, T., Godecke, I., Chapman, B., Lowel, S., and Obermayer, K. (2000). An analysis of orientation and ocular dominance patterns in the visual cortex of cats and ferrets. *Neural Comput.* *12*, 2573–2595.
- Renart, A., Song, P., and Wang, X.J. (2003). Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* *38*, 473–485.
- Ringach, D.L. (2003). Neuroscience: states of mind. *Nature* *425*, 912–913.
- Ringach, D.L., Shapley, R., and Hawken, M.J. (2002). Orientation selectivity in macaque V1: diversity and laminar dependence. *J. Neurosci.* *22*, 5639–5651.
- Sharon, D., and Grinvald, A. (2002). Dynamics and constancy in cortical spatiotemporal patterns of orientation processing. *Science* *295*, 512–515.
- Shoham, D., Glaser, D.E., Arieli, A., Kenet, T., Wijnbergen, C., Toledo, Y., Hildesheim, R., and Grinvald, A. (1999). Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. *Neuron* *24*, 791–802.
- Ts'o, D.Y., Gilbert, C.D., and Wiesel, T.N. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* *6*, 1160–1170.
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* *286*, 1943–1946.
- Tyler, C.J., Dunlop, S.A., Lund, R.D., Harman, A.M., Dann, J.F., Beazley, L.D., and Lund, J.S. (1998). Anatomical comparison of the macaque and marsupial visual cortex: common features that may reflect retention of essential cortical elements. *J. Comp. Neurol.* *400*, 449–468.
- Wang, G., Ding, S., and Yunokuchi, K. (2003). Difference in the representation of cardinal and oblique contours in cat visual cortex. *Neurosci. Lett.* *338*, 77–81.
- Weliky, M., and Katz, L.C. (1994). Functional mapping of horizontal connections in developing ferret visual cortex: experiments and modeling. *J. Neurosci.* *14*, 7291–7305.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.* *16*, 2112–2126.