Efficient coding in heterogeneous neuronal populations

Mircea I. Chelaru and Valentin Dragoi*

Department of Neurobiology and Anatomy, University of Texas-Houston Medical School, Houston, TX 77030

Edited by Dale Purves, Duke University Medical Center, Durham, NC, and approved August 22, 2008 (received for review August 7, 2008)

A ubiquitous feature of neuronal responses within a cortical area is their high degree of inhomogeneity. Even cells within the same functional column are known to have highly heterogeneous response properties when the same stimulus is presented. Whether the wide diversity of neuronal responses is an epiphenomenon or plays a role for cortical function is unknown. Here, we examined the relationship between the heterogeneity of neuronal responses and population coding. Contrary to our expectation, we found that the high variability of intrinsic response properties of individual cells changes the structure of neuronal correlations to improve the information encoded in the population activity. Thus, the heterogeneity of neuronal responses is in fact beneficial for sensory coding when stimuli are decoded from the population response.

computational modeling | information coding | sensory cortex | visual cortex | neuronal populations

It is generally believed that neurons within the same cortical area vary widely in their response properties, such as mean firing rate, receptive field location and size, and stimulus selectivity. Even the responses of nearby neurons located within the same functional column, which encode the same stimulus property, exhibit a high degree of heterogeneity (1–5). In primary visual cortex (V1), for instance, it has been reported that neurons tuned to the same stimulus orientation exhibit a high degree of variability in their strength of orientation selectivity, peak response, and baseline activity (1, 2). Although this ubiquitous feature of neuronal responses has been reported in many sensory areas, e.g., visual (1, 2), auditory (3), somatosensory (4), and motor cortical areas (5), and in a wide range of species, such as rats (3, 4), cats (1), and monkeys (2, 5), whether it serves a purpose or is an epiphenomenon caused by an imprecise synaptic targeting during development (6) is unclear.

The idea of response heterogeneity is apparently challenged by the precision of functional architecture in specific cortical areas (7). Thus, two-photon calcium imaging studies have recently demonstrated that the functional architecture of orientation and direction selectivity in cat visual cortex (7, 8) is extraordinarily precise. For instance, in area 18, neurons preferring opposite stimulus directions were segregated by remarkably specific "fractures" that were one to two cells wide (7). Nonetheless, the precision of the functional architecture is not necessarily inconsistent with heterogeneous neuronal responses. Indeed, neurons can vary widely in their response properties such as baseline firing, response amplitude, and strength of stimulus selectivity without showing significant differences in their stimulus preference. This claim is supported by previous optical imaging (9, 10) and electrophysiological studies (2) confirming that orientation preference varies smoothly across the cortical surface, even though the responses of nearby neurons were highly heterogeneous.

Despite the fact that the heterogeneity of neuronal responses has been amply reported in experimental studies, its consequences were rarely examined in models of network function. However, a heterogeneous distribution of neuronal responses across the network is likely to influence the accuracy of the population code. This issue is important: the accuracy with which cortical networks encode external stimuli into patterns of neural impulses determines the accuracy of behavioral responses (11–13). In sensory systems, the stimulus-triggered responses of neurons are often characterized by bell-shaped tuning curves (1, 2, 14, 15). It has been demonstrated that population coding depends not only on the shape of the neurons' tuning curves but also on the distribution of neuronal correlations across the network (11–13, 16). Indeed, during the past decade, it has become increasingly understood that the trial-by-trial variability in neuronal responses, or "noise," is not independent, but exhibits correlations (17, 18).

In principle, a highly variable distribution of neuronal responses across the network could influence the relationship between the shape of the neurons' tuning curves and the structure of correlations in trial-by-trial response variability between neurons. Characterizing this relationship is critical for determining the amount of information contained in a population of neurons and could provide important clues for understanding whether neuronal networks perform efficient computations. However, whether the wide diversity of neuronal responses across the population influences the structure of trial-by-trial neuronal (noise) correlations and its relationship with population coding is unknown. We addressed this issue in the context of orientation selectivity in primary visual cortex (V1) by implementing a conductance-based integrate-and-fire network of cortical neurons that sharpen the broadly tuned thalamocortical afferents (19). We found that the information contained in the population response increases with the increase in the heterogeneity of the orientation tuning curve profiles of individual neurons, mainly because of a decrease in the correlation between cells. In other words, contrary to expectation, response diversity is in fact beneficial when stimulus orientation is decoded from the population response.

Results

We have implemented a dynamic model that is similar to the conductance-based integrate-and-fire model published by Somers *et al.* (19). The model describes neuronal interactions in layer 4 of V1 between 1,008 excitatory and 252 inhibitory orientation-selective cells that receive excitatory feedforward inputs from lateral geniculate nucleus (LGN) [see *Materials and Methods* and supporting information (SI) *Text*]. First, we ensured that model cortical cells exhibit classical receptive field properties of V1 neurons, such as sharpening of orientation selectivity, in agreement with reverse correlation studies (20, 21). For instance, Fig. 1 *A* and *B* shows that the weakly selective responses of LGN afferents to V1 excitatory neurons (*A*, an oriented stimulus was presented for 500 ms) are sharpened by V1 intracortical circuits (*B*).

Author contributions: V.D. designed research; M.I.C. performed research; and V.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

*To whom correspondence should be addressed at: Department of Neurobiology and Anatomy, University of Texas-Houston Medical School, Suite 7.166, 6431 Fannin Street, Houston, TX 77030. E-mail: v.dragoi@uth.tmc.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0807744105/DCSupplemental.

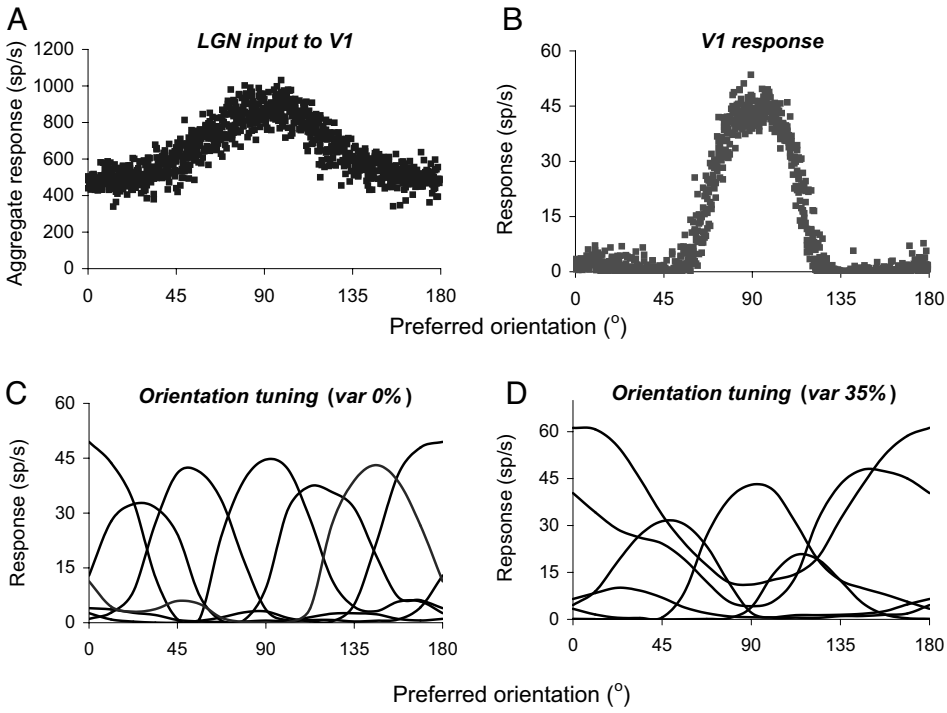

Fig. 1. Homogeneous and heterogeneous networks. (*A*) Homogeneous network—pooled response of LGN afferents to V1 excitatory neurons to a stimulus presented for 500 ms (averaged over 1,008 trials). Each point represents the average sum of the LGN spikes converging to a V1 neuron of preferred orientation indicated on the *x* axis. Neurons are ordered monotonically according to their preferred orientation. The stimulus was a vertically oriented bar of 4 deg length and 1 deg width. (*B*) Homogeneous network—mean V1 population responses to a stimulus oriented at 90°. Intracortical connections sharpen the weakly selective responses of the LGN afferents. (*C*) Homogeneous network—orientation tuning curves of six V1 neurons with preferred orientations separated by 30°. The homogeneous network had an intrinsic (small) variability due to the Poisson spike generator from LGN and the probabilistic connectivity for the thalamocortical afferents and intracortical connections. We labeled this intrinsic heterogeneity as 0% variability (or control). (*D*) Heterogeneous network—orientation tuning curves of the same neurons as in *C*. The figure shows a high degree of heterogeneity of response amplitudes and orientation tuning widths. For each neuron, the control synaptic conductances (see *SI Text*) were transformed into Gaussian-distributed random variables of mean equal to the control conductance, μ, and standard deviation, σ. The ratio between σ and μ represents the percent induced variability. *D* represents the orientation tuning curves when the induced synaptic variability was 35% (the maximum used in our study).

We controlled the degree of response heterogeneity within the V1 network by modulating the variance of the conductance distributions for the thalamocortical and recurrent intracortical synapses (see *Materials and Methods*). This is consistent with experimental (18) and modeling (19, 22, 23) studies indicating that visual cortical responses depend critically on the integration of excitatory and inhibitory thalamocortical and intracortical inputs by individual cells. Response heterogeneity was implemented by replacing the model fixed conductance values with Gaussian-distributed random variables (see *Materials and Methods*). This implementation allowed us to vary the ratio of the standard deviation and mean synaptic conductance, and thus increase the degree of response heterogeneity (referred to as "induced variability") between 0 (homogeneous networks) and 35% (this value maximizes the degree of response heterogeneity in the network). Increasing the degree of heterogeneity beyond 35% results in a severe loss of orientation selectivity for >60% of the neurons and a non-Gaussian distribution of the response properties (see *Materials and Methods*). This loss of orientation selectivity beyond 35% induced variability is caused by an alteration of the balance between excitation and inhibition (controlled by the strength of intracortical synapses; see *Materials and Methods*).

In agreement with previous findings, we found that orientation-selective V1 responses of neighboring cells exhibit a high degree of variability in orientation bandwidth, peak firing rate, and baseline activity (1, 2). Indeed, we found that increasing the degree of response heterogeneity (Fig. 1 *C–D*) by increasing the synaptic variability, increased the variance of the neurons' response amplitude (Fig. 2*A*) and orientation selectivity index (OSI) (Fig. 2*B*) and altered the distribution of the tuning curve slopes of individual neurons (Fig. 2*C*). However, the orientation preference of neurons remained virtually unaltered (Fig. 2*D*; $P > 0.1$; Wilcoxon signed-rank test). This result further supports the idea that response heterogeneity is not inconsistent with the precision of functional architecture (7, 8).

The degree of response variability in the model was comparable to that shown by neurons in cat area 17. Thus, across a population of 109 neurons (9, 10), the standard deviation of response amplitude was found to be $\sigma_{resp} = 22.1$ Hz, whereas the standard deviation of the OSI was $\sigma_{osi} = 0.17$. These values are comparable with the variability exhibited by our model neurons: The standard deviation of response amplitude was $\sigma_{resp} = 13.1$ Hz (induced variability 25%) and $\sigma_{resp} = 18.7$ Hz (induced variability 35%), whereas the standard deviation of the OSI was $\sigma_{osi} = 0.175$ (induced variability 25%) and $\sigma_{osi} = 0.2$ (induced variability 35%). These values are also consistent with the response variability of monkey V1 neurons (2, 24).

These results bring up the issue of whether the increase in the degree of response heterogeneity influences the information encoded in the population response. A measure of the accuracy of population coding is the discrimination threshold, which is the smallest change in orientation that can be reliably discriminated by a network of cells on the basis of a single trial response. The inverse of this threshold is proportional to the square root of Fisher information, which represents the upper limit of the accuracy with which any decoding mechanism can extract information about a stimulus parameter, e.g., orientation (25).

We computed the network orientation discrimination threshold by using a linear estimate of the Fisher information (FI) (11–13, 26, 27). Stimulus orientation was decoded from the population response by using a locally optimal linear estimator trained to minimize the discrimination threshold between two adjacent stimulus orientations by maximizing Fisher information. The information contained in the population response of the regular spiking neurons was estimated with a lower bound of the FI (labeled I_{LD}) computed from the firing rates of the excitatory neurons by training a linear decoder of stimulus orientation (ref. 27; see *Materials and Methods*). The weights of the linear decoder were computed to maximize I_{LD} (see *Materials and Methods*).

We estimated the amount of information in the population response by presenting input stimuli differing by 2° in orientation (i.e., 89° and 91°). To ensure that the results do not depend on the specific values for the synaptic conductances, we ran 20 experimental sessions, i.e., 5 sessions for each level of synaptic variability. In

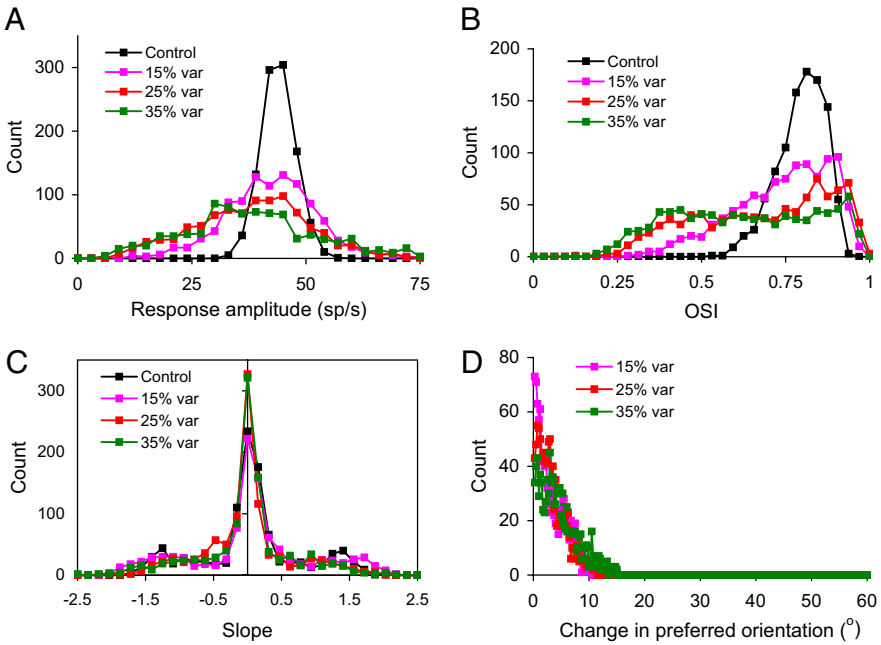

Fig. 2. Changes in the neurons' response properties due to an increase in response heterogeneity. (*A*) Distributions of response amplitudes (the difference between the maximum and the minimum of the tuning curves in spikes per second) for homogeneous and heterogeneous networks (for different levels of synaptic variability). The variance of response amplitudes was increased by 4, 9, and 12 times for the three heterogeneous networks. (*B*) Distributions of the orientation selectivity index (OSI) of individual neurons for different levels of induced variability. The variance of OSIs was increased by 5, 8, and 10 times for the three heterogeneous networks. (*C*) Distributions of the tuning curve slopes for different levels of synaptic variability. The derivatives of the tuning curves do not vary significantly with the level of induced variability. The slopes were computed from the tuning curves of all of the neurons by presenting a stimulus at a fixed, vertical orientation. Increasing the level of synaptic variability did not change the variance of the tuning curve slopes in a pronounced manner (40.4%, −23.9%, and −7.7% changes in slope variance for the three heterogeneous networks). (*D*) Distributions of the absolute changes in the neurons' preferred orientation for the heterogeneous vs. homogeneous networks at different levels of synaptic variability (orientation differences are shown in bins of 0.25°). The orientation preference of neurons did not change significantly when the level of response heterogeneity was increased.

each session, we generated a new recurrent network by randomizing the values of the synaptic conductances while keeping intracortical connections fixed (see *Materials and Methods*). Our key finding is that the network orientation discrimination threshold is lower when the heterogeneity of orientation-selective responses is increased. Thus, we found (Fig. 3*A*) that increasing the degree of response heterogeneity by increasing the synaptic variability from 0% to 35% causes an increase in information, I_{LD}, relative to the homogeneous case ($P < 0.05$ for each comparison at each level of synaptic variability; bootstrap method). Correspondingly, there was a decrease in the network orientation discrimination threshold as response heterogeneity was increased (Fig. 3*B*; $P < 0.05$ for each comparison; bootstrap method). These results hold when noisy oriented bars are used instead of pure gratings (Fig. S1). Thus, despite the overall decrease in the decoder information because of an increase in the level of visual noise, the performance of heterogeneous networks remained greater than that of the homogeneous network for a broad range of synaptic variability.

What could explain the improvement in population coding efficiency when the degree of response heterogeneity is increased? FI is known to depend on the shape of the orientation tuning curves of the neurons in the population and on the noise correlations between cells, i.e., the covariance matrix (11–13, 16, 28). To understand whether the increase in FI in the heterogeneous neuronal population is primarily because of the changes in the shape of the orientation tuning curves (Fig. 1 *E–G*) or the changes in the structure of neuronal correlations, we computed the shuffled information (for both the homogeneous and heterogeneous networks) by randomly rearranging neuronal responses across trials (see *Materials and Methods*). This was performed by independently shuffling the trial-by-trial responses for each pair of cells in the network (28). Thus, shuffling the trials preserves the profile of the neurons' tuning curves while significantly reducing the correlations between cells.

Fig. 4 *A* and *B* shows that although shuffled information is higher than the unshuffled information, it is relatively insensitive to the changes in the degree of response heterogeneity. In other words, shuffling trials, and thus removing correlations among cells, causes a statistically significant decrease in the improvement in information because of response heterogeneity relative to the unshuffled case ($P < 0.05$, for each comparison between shuffled and unshuffled ΔI_{LD} for each level of synaptic variability; bootstrap method). This result demonstrates that increasing the diversity of orientation-selective responses alters correlations between neurons to account for the improvement in population coding.

We directly tested the relationship between response heterogeneity and the structure of correlations across the network by examining whether the increase in the diversity of cortical responses is accompanied by a decrease in noise correlations. Clearly, the covariance matrices associated with the homogeneous and heterogeneous networks show pronounced differences (Fig. 4 *C* and *D*). Thus, for homogeneous networks, we found that the covariance matrix of the responses exhibits both large negative and positive values for a broad range of orientations (Fig. 4*C*). Consistent with experimental data (18, 29), we found that the correlations between neurons decayed exponentially with the difference in the cells' preferred orientation (Fig. S2). However, when response heterogeneity was increased, the correlation coefficients decreased in value and range and were significant only for the cells preferring similar orientations (Fig. 4*D*; see also Figs. S2 and S3). Nonetheless, we did not notice any changes in the spatial organization of noise correlations as the correlations associated with the heterogeneous networks simply decreased exponentially as a function of the orientation difference between cells (Fig. S2). Because the efficiency of population coding has been shown to depend critically on the strength of correlations between neurons (a reduction in correlations is typically associated with an increase in FI, e.g., refs. 13 and 29, but see refs. 11 and 16), these results indicate that increasing the diversity of orientation-selective responses (Figs. 1 and 2) decreases neuronal correlations (Fig. 4) to improve the efficiency of population coding.

Discussion

Although heterogeneous response properties have been observed in many cortical areas and species, the functional role of neuronal response diversity has been elusive. We have shown that the spatial heterogeneity of orientation selectivity improves coding efficiency in V1 networks, despite the high variability in the tuning curve profiles of individual neurons. These results challenge the commonly held notion that response heterogeneity is consistent with noisy, unreliable, population codes. Indeed, the heterogeneity of

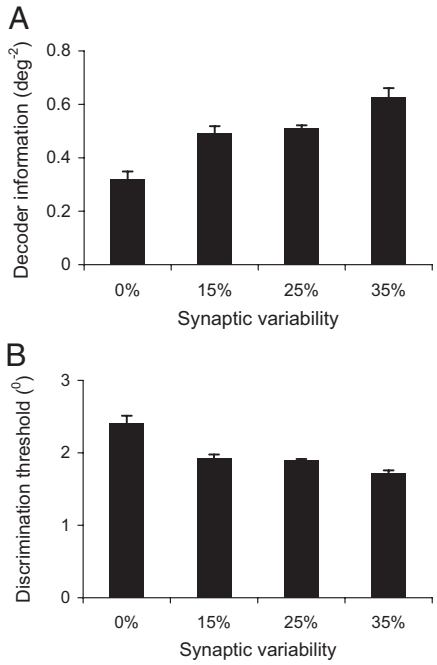

Fig. 3. Relationship between network performance and response heterogeneity. (*A*) The network-orientation-discrimination performance (decoder information) increases with the degree of induced synaptic variability ($P < 0.05$ for each comparison at each level of synaptic variability; bootstrap method). The stimulus was decoded from the population response by using a linear decoder optimized such that information, computed as the lower limit of the FI, was maximized (see *SI Text*). Information was calculated for stimuli differing by 2° in orientation for 20 experimental sessions (5 sessions for each level of synaptic variability). In each session, we generated a new recurrent network that had unchanged intracortical connections but different synaptic conductances (randomly generated as described in *SI Text*). (*B*) The orientation discrimination threshold, which is inversely related to the decoder information, decreases with the increase in the induced synaptic variability ($P < 0.05$; bootstrap method). Error bars represent SEM.

neuronal responses is often considered noise and is typically disregarded in studies of cortical function. In principle, a broad distribution of response amplitudes and tuning curve shapes for nearby neurons within the same functional column could result in highly variable neuronal responses to small differences between nearby input stimuli. Therefore, the downstream neurons decoding these responses would have to rely on unreliable reports about differences between stimuli, which could deteriorate network performance. However, we have shown that this is not the case; response heterogeneity is in fact beneficial when stimuli are decoded from the population response. This result suggests that response diversity should no longer be ignored in models of cortical function because heterogeneous networks offer superior decoding compared with homogeneous networks.

Our results are consistent with previous studies (30) that have suggested that, in certain conditions, the diversity of the tuning curve shapes of individual neurons could influence the information encoded in population activity. However, previous studies have not examined whether the heterogeneity of neuronal responses influences the relationship between the changes in neurons' average response properties and the structure of trial-by-trial neuronal correlations, as well as the relationship with population coding. We found that increasing the degree of response heterogeneity improves the efficiency of the population code, a result that is due primarily to a change in the structure of correlations. Indeed, when neuronal responses were decorrelated by shuffling the trials, the increase in neuronal diversity caused significantly smaller changes in coding efficiency. Nonetheless, our results do not imply that a reduction in noise correlations should always be associated with an increase in network efficiency. Indeed, it has been shown that correlations are beneficial or detrimental for coding depending on the distribution of noise correlations across the entire network (11) and the relationship to signal correlations (16). However, in the context of the network examined here, correlated noise tends to occur along the same dimension as the signals that distinguish between the input stimuli (27, 29), and is therefore detrimental to decoding performance.

One important issue is whether other types of heterogeneities in neuronal responses can influence population coding efficiency. Thus, it is well known that the local network environment, not only the responses of individual neurons, is heterogeneous in the way it represents specific aspects of the sensory input. For instance, within primary visual cortex, neurons are clustered according to their orientation preference in orientation domains that converge at singularities or pinwheel centers (31–33). The structure of the orientation map in V1 implies that the orientation distribution of local connections would vary with a neuron's position within the map: Neurons in pinwheel centers are likely to be connected to neurons of a broader range of orientations than neurons in orientation domains (9) to possibly influence information coding. However, despite the fact that the structure of the stimulus preference map is heterogeneous, it is unlikely to influence population coding. Indeed, coding efficiency depends on the shape of the neurons' tuning curves and on the structure of neuronal correlations. Importantly, despite the difference in local inputs to neurons in pinwheel centers and orientation domains, both these classes of cells have been described to have similar orientation tuning characteristics (9, 34, 35). In addition, optical imaging and electrophysiological studies have previously reported (36, 37) that correlations between neurons depend only on the relative separation between cells and are independent on the cells' location on the cortical surface. Altogether, these studies suggest that the heterogeneous structure of the orientation preference maps is unlikely to have an impact on population coding.

It could be argued that our demonstration that the heterogeneity of neuronal responses improves coding accuracy relies critically on the fact that orientation tuning emerges from recurrent, intracortical, interactions between V1 neurons. However, a competing model of orientation tuning, i.e., the feedforward model, which assumes that orientation selectivity arises from the pattern of thalamocortical inputs, raises the issue of whether our conclusions would remain valid if tuning in V1 would arise from feedforward rather than recurrent processing. To address this issue, we implemented a feedforward model of orientation selectivity in which excitatory cortical neurons receive only excitatory inputs from LGN and orientation nonspecific inhibitory inputs from local inhibitory neurons (27, 38, 39). To directly compare the impact of response heterogeneity on information coding in recurrent and feedforward networks, we ensured that the orientation-selective responses of individual neurons in both models matched each other in terms of tuning strength and response amplitude (Fig. S4 *A–C*). Confirming previous results (27), we found that despite the fact that both models show similar orientation tuning properties, the feedforward model conveys approximately three times more information about the input stimuli than the recurrent model (comparing Fig. 3*A* and Fig. S4*D*). This increase in information is mainly caused by a significant reduction in noise correlations in feedforward networks in the absence of recurrent excitatory connections (compare Fig. 4*C* and Fig. S4*E*). Importantly, however, heterogeneous feedforward networks tend to encode more information in the population response than homogeneous networks tend to encode (Fig. S4*D*). Although this result is consistent with the improvement in coding accuracy in heterogeneous recurrent networks (Fig. 3), the degree of improvement dropped from 76% to 23% (comparing the mean change of information in the heterogeneous vs. homogeneous network in the 25% and 35% synaptic variability conditions for the

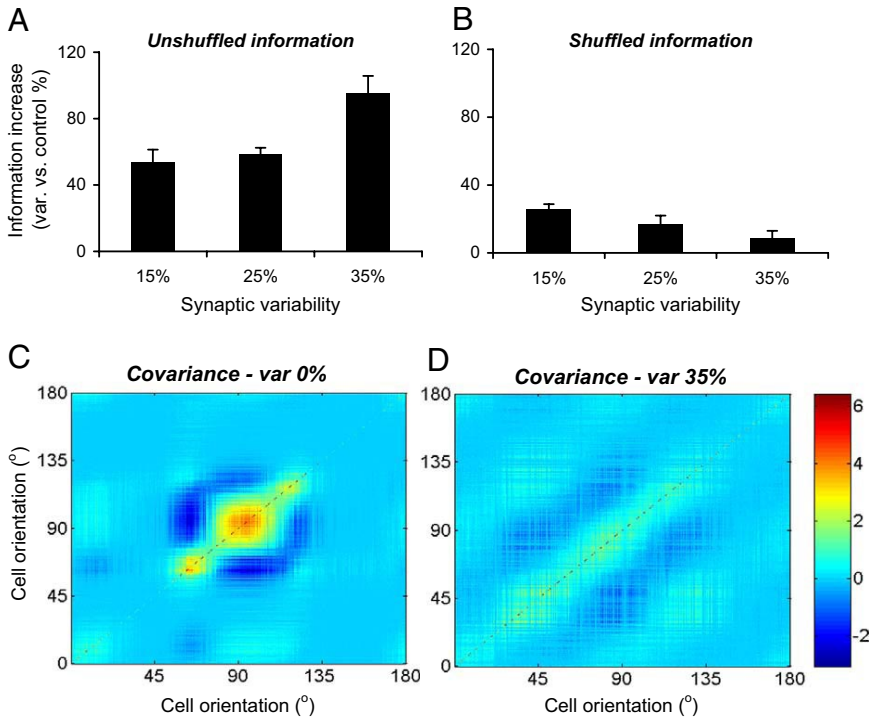

Fig. 4. Response heterogeneity influences the correlations between neurons. (*A* and *B*) Changes in unshuffled (*A*) and shuffled (*B*) information for the heterogeneous vs. homogeneous networks at different levels of synaptic variability. Shuffling trials, and thus removing correlations among cells, causes a decrease in information relative to the unshuffled case ($P < 0.05$; bootstrap method). The shuffled information is insensitive to the changes in the degree of response heterogeneity ($P > 0.05$; bootstrap method). For each experimental session, the information changes were computed relative to the mean over five experimental sessions (see *SI Text*). Error bars represent SEM. (*C* and *D*) Covariance matrices of the V1 neurons for the homogeneous (*C*) and heterogeneous (*D*) networks. In the homogeneous network (induced variability 0%), there are both positive and negative strong correlations between cells preferring a broad range of orientations. In the heterogeneous network (induced variability 35%), correlations are smaller and are limited to the cells preferring similar orientations. Each covariance matrix represents the average of the covariance matrices across five experimental sessions (see *SI Text*).

recurrent and feedforward networks). This difference is because of a pronounced decrease in pairwise noise correlations in the feedforward network (Fig. S4 *E* and *F*), which limits the degree of change in neuronal correlations when the response heterogeneity is increased.

A reduction in neuronal correlations because of an increase in the degree of response heterogeneity is likely to make a recurrent network behave more like a network of independent neurons. This issue is important in light of a recent suggestion (27) that models that rely on recurrent processing encode a smaller amount of information in the population response than networks of independent neurons. However, the fact that recurrent processing is a ubiquitous operating regime of cortical networks *in vivo* raises the issue of whether neuronal networks could retain the computational advantages of recurrent networks while avoiding the decrease in network efficiency because of an excessive increase in correlations. Here, we propose a solution to this problem: response heterogeneity may be the natural way through which correlations between the neurons involved in recurrent processing could be diminished to improve the accuracy of network computations.

Finally, our research captures the role of a key variable, i.e., response heterogeneity, that characterizes the response properties of neurons in many cortical areas. Indeed, although we focus here on orientation selectivity, our analysis applies to any network that relies on feedforward and recurrent processing (e.g., auditory or somatosensory cortex). Hence, our results may have general implications for sensory coding in a wide range of visual and nonvisual cortical areas.

Materials and Methods

Integrate-and-Fire Orientation Selectivity Model. The model has three stages: retina, LGN, and V1. Retina contains two layers of ON-center surround cells, and respectively OFF-center surround cells, driven by image inputs (see *SI Text*). Inputs consist of 4° × 1° oriented bars presented for 500 ms. Retinal cells are modeled by a difference of Gaussian filters. Retinal activity is modulated by a saturating nonlinearity to account for the stimulus contrast sensitivity, and then the output is sent to the LGN cells with a random delay. LGN cells are organized in two ON/OFF layers having one-to-one connections with the retinal cells. LGN cells modulate the Poisson spiking generators that feed the V1 network. In agreement with experimental (40, 41) and modeling studies (19, 38, 39), we ensured that the LGN input to a cortical cell is significantly, but broadly tuned for orientation. Indeed, experimental evidence has indicated that although LGN cells are virtually untuned, their aggregate input to a cortical neuron exhibits significant orientation tuning (because of the spatial organization of the LGN inputs to V1 cells; refs. 42 and 43). Importantly, however, our results do not depend on whether the LGN input is broadly or narrowly orientation tuned (Fig. S5). The layer of cortical cells simulates a hypercolumn of layer 4Cα, consisting of simple cells. The network has 1008 excitatory regular-spiking cells and 252 inhibitory fast-spiking cells. The network exhibits 252 preferred orientations in the interval (0°, 180°), with four excitatory neurons per orientation, and a single inhibitory neuron per orientation. All neurons are modeled as conductance-based, integrate-and-fire neurons. The thalamocortical afferents are established at random from ON and OFF subfields that are defined over the ON and OFF LGN layers by using Gabor functions. We implemented two models of orientation selectivity in V1: a recurrent model and a feedforward model.

Recurrent model of orientation selectivity. The recurrent connections between V1 cells are established at random, after a Gaussian probability distribution centered on the cell's preferred orientation (see *SI Text*). The response of LGN afferents (Fig. 1*A*) is broadly tuned to the stimulus orientation, and then intracortical connections within V1 sharpen orientation selectivity (Fig. 1*B*).

Feedforward model of orientation selectivity. The parameters of the Gabor functions were chosen to ensure that the pooled LGN inputs are more selective for orientation than those corresponding to the recurrent network. Cortical excitatory neurons only receive excitatory feedforward inputs from LGN and inhibitory inputs from inhibitory cells (recurrent excitatory connections have been completely removed). Inhibitory neurons receive excitatory inputs from LGN and then send orientation-nonspecific inhibition to excitatory neurons (see *SI Text*). Thus, the feedforward model relies on a spatially uniform thalamocortical inhibition that balances the LGN excitatory inputs at nonpreferred orientations to produce sharp orientation selectivity (Fig. S4*A*). The model implementation is similar to that of other feedforward models of orientation selectivity (27, 38, 39).

Inducing Synaptic Variability. The synaptic conductances used in the integrate-and-fire model are defined as:

$$g(t) = \bar{g} \sum_p [t - t_p]^+ \left(\frac{e}{\tau_{peak}}\right) \exp\left(-\frac{t - t_p}{\tau_{\text{peak}}}\right),$$

where $\bar{g}$ represents the maximum conductance change produced by presynaptic spikes and t_p is the time of the pth spike from the presynaptic cell. The maximum conductance $\bar{g}$ and τ_{peak} have different values depending on the synapse type (see *SI Text*). For both the recurrent and feedforward models, we

controlled the degree of synaptic variability by replacing the fixed values, $\bar{g}$, with Gaussian-distributed random variables of mean $\bar{g}$ and standard deviation $\sigma = c \cdot \bar{g}$, where c is a coefficient that controls the degree of induced variability. We used c values between 0 (control case) and 0.35 (maximum induced variability; because 99% of the random numbers of a Gaussian distribution are within 3σ of their mean, the maximum value for c was 0.35; increasing c beyond 0.35 results in a severe loss of orientation selectivity for >60% of the neurons and a non-Gaussian distribution of response amplitudes, OSIs, and tuning curve slopes). Once the random values for the maximum conductance changes were generated, they were held fixed throughout the trials.

FI Estimator. Stimulus orientation θ was decoded from the firing rates of the excitatory V1 neurons by using a linear decoder:

$$\hat{\theta} = W^T R + b, \qquad [1]$$

where W is a weight vector, b is a scalar, and R is the firing rate of the excitatory neurons. Decoder weights W and b were optimized to minimize the discrimination threshold between two nearby stimulus orientations, θ_1 and θ_2 (we used $\theta_1 = 89°$ and $\theta_2 = 91°$ throughout the study). As discrimination threshold, $\Delta\theta_{TH}$, we used the upper bound of the change in orientation, $\Delta\theta$, which can be detected in 75% of the trials by an ideal observer (44):

$$\Delta\theta_{TH} = \frac{1.35}{\sqrt{I_{LD}}}. \qquad [2]$$

The factor I_{LD} is an estimate of the FI (22, 25):

$$I_{LD} = \frac{((\langle\hat{\theta}_2\rangle - \langle\hat{\theta}_1\rangle)/\Delta\theta)^2}{\sqrt{\sigma^2_{\hat{\theta}_2} + \sigma^2_{\hat{\theta}_1}}} \qquad [3]$$

where $\Delta\theta = \theta_2 - \theta_1$, $\langle\hat{\theta}_i\rangle$ is the mean and $\sigma^2_{\hat{\theta}_i}$ ($i = 1, 2$) is the variance of the linear decoder output when a bar stimulus of orientation θ_i is presented to the retina.

The discrimination threshold $\Delta\theta_{TH}$ was minimized by computing the decoder weights W and b so that the information I_{LD} was maximized. First, we generated 1,008 trials in which a bar of orientation θ_1 was presented, followed by 1,008 trials in which a bar of orientation θ_2 was presented (each stimulus was presented for 500 ms). Then, we divided the trials for each stimulus orientation into equal sets: a training set and a test set. We used the data of the training set (504 trials with orientation θ_1 and 504 trials with orientation θ_2) to compute the linear decoder weights by using the conjugate gradient algorithm (27, 45). For each of the algorithm iterations, we computed information I_{LD} from Eq. **3** by using the test set and the decoder weights corresponding to that iteration. We used 1,008 iterations, which is equal to the dimension of the linear decoder (defined by the number of excitatory neurons), because the minimum of the surface error for a linear decoder trained with the conjugate gradient algorithm is reached at most after a number of iterations equal to the dimension of the decoder. We kept the weights of the decoder for which I_{LD} was maximum. FI was calculated for 20 experimental sessions by using 5 sessions for each value of synaptic variability (0%, 15%, 25%, and 35%). In each session, we generated a new recurrent network that had similar intracortical connections but different synaptic conductances that were randomly generated. This allowed us to compute the information, I_{LD}, and the changes in information (for the heterogeneous vs. homogeneous networks) over 5 experimental sessions for each level of synaptic variability.

Shuffled Information. We computed shuffled information with Eq. **3** by shuffling the responses of neurons across trials in both the training and test sets (the noise correlations among cells were completely removed by shuffling). This was performed by independently shuffling the trial-by-trial responses of each cell in the network (28), thus destroying the temporal structure of the responses.

OSI. The OSI of an orientation tuning curve was computed for 16 oriented stimuli θ_i, $i = 1, 2, \ldots, 16$, uniformly distributed over the entire orientation range (0°, 180°). OSI was computed as OSI = c/M, where M was the mean spike count for all of the orientations and $c = \sqrt{a^2 + b^2}$, where a and b are the second-order Fourier components of the responses $R(\theta_i)$ to orientation θ_i,: $a = \sum_{i=1}^{16} R(\theta_i)\cos(2\theta_i)$, $b = \sum_{i=1}^{16} R(\theta_i)\sin(2\theta_i)$ (46). Responses $R(\theta_i)$ were computed as mean spike counts for 128 trials of 500 ms each.

ACKNOWLEDGMENTS. We thank Diego Gutnisky and Kreso Josic for comments on the manuscript. This work was supported by the Pew Scholars Program and the James S. McDonnell Foundation (V.D.).

1. Hubel D, Wiesel T (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154.
2. Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: Diversity and laminar dependence. *J Neurosci* 22:5639–5651.
3. Kilgard MP, Merzenich MM (1999) Distributed representation of spectral and temporal information in rat primary auditory cortex. *Hear Res* 134:16–28.
4. Staiger JF, *et al*. (2004) Functional diversity of layer IV spiny neurons in rat somatosensory cortex: Quantitative morphology of electrophysiologically characterized and biocytin labeled cells. *Cereb Cortex* 14:690–701.
5. Lee D, Port NL, Kruse W, Georgopoulos AP (1988) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18:1161–1170.
6. Katz LC, Shatz CJ (1996) Synaptic activity and the construction of cortical circuits. *Science* 274:1133–1138.
7. Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433:597–603.
8. Ohki K, *et al*. (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925–928.
9. Dragoi V, Rivadulla C, Sur M (2001) Foci of orientation plasticity in visual cortex. *Nature* 411:80–86.
10. Dragoi V, Turcu CM, Sur M (2001) Stability of cortical responses and the statistics of natural scenes. *Neuron* 32:1181–1192.
11. Abbott L, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
12. Sompolinsky H, Yoon H, Kang KJ, Shamir M (2001) Population coding in neuronal systems with correlated noise. *Phys Rev E* 6405.
13. Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. *Nat Rev Neurosci* 1:125–132.
14. Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349.
15. Miller JP, Jacobs GA, Theunissen F (1991) Representation of sensory information in the cricket cercal sensory system. I. Response properties of the primary interneurons. *J Neurophysiol* 66:1680–1689.
16. Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
17. Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.
18. Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
19. Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *J Neurosci* 15:5448–5465.
20. McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* 54:782–806.
21. Ringach DL, Hawken MJ, Shapley R (1997) Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387:281–284.
22. Douglas R, Martin K (2004) Neuronal circuits of the neocortex. *Annu Rev Neurosci* 27:419–451.
23. Chelaru MI, Dragoi V (2008) Asymmetric synaptic depression in cortical networks. *Cereb Cortex* 18:771–788.
24. Dragoi V, Sharma J, Miller EK, Sur M (2002) Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nat Neurosci* 5:883–891.
25. Kay SM (1993) *Fundamentals of Statistical Signal Processing*. Prentice–Hall Signal Processing Series (Prentice–Hall, Englewood Cliffs, NJ).
26. Paradiso MA (1988) A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biol Cybern* 58:35–49.
27. Seriès P, Latham PE, Pouget A (2004) Tuning curve sharpening for orientation selectivity: Coding efficiency and the impact of correlations. *Nat Neurosci* 7:1129–1135.
28. Dayan P, Abbott LF (2001) *Theoretical Neuroscience* (MIT Press, Cambridge, MA).
29. Gutnisky D, Dragoi V (2008) Adaptive coding of visual information in neural populations. *Nature* 452:220–224.
30. Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. *Neural Comput* 18:1951–1986.
31. Hubel DH, Wiesel TN (1974) Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J Comp Neurol* 158:267–293.
32. Bonhoeffer T, Grinvald A (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353:429–431.
33. Blasdel GG (1992) Differential imaging of ocular dominance and orientation selectivity in monkey striate cortex. *J Neurosci* 12:3115–3138.
34. Maldonado PE, Gödecke I, Gray CM, Bonhoeffer T (1997) Orientation selectivity in pinwheel centers in cat striate cortex. *Science* 276:1551–1555.
35. Mariño J, *et al*. (2005) Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat Neurosci* 8:194–201.
36. Das A, Gilbert CD (1999) Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature* 399:655–661.
37. Schummers J, Sur M (2000) Rules of functional connectivity in primary visual cortex. *Soc Neurosci Abstr* 26:53.12.
38. Troyer TW, Krukowski AE, Priebe NJ, Miller KD (1998) Contrast-invariant orientation tuning in cat visual cortex: Thalamocortical input tuning and correlation-based intracortical connectivity. *J Neurosci* 18:5908–5927.
39. Ferster D, Miller KD (2000) Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci* 23:441–471.
40. Ferster D, Chung S, Wheat H (1996) Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380:249–252.
41. Chung S, Ferster D (1998) Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron* 20:1177–1189.
42. Tanaka K (1983) Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J Neurophysiol* 49:1303–1318.
43. Reid RC, Alonso JM (2005) Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–284.
44. Green DM, Swets JA (1966) *Signal Detection Theory and Psychophysics* (Wiley, New York).
45. Hagan MT, Demuth HB, Beale MH (1996) *Neural Network Design* (PWS Publishing, Boston).
46. Worgotter F, Eysel UT (1991) Correlations between directional and orientational tuning of cells in cat striate cortex. *Exp Brain Res* 83:665–669.