

A Fast and Simple Population Code for Orientation in Primate V1

Philipp Berens,^{1,2,3,4} Alexander S. Ecker,^{1,2,3,4} R. James Cotton,⁴ Wei Ji Ma,⁴ Matthias Bethge,^{1,2,3} and Andreas S. Tolias^{2,4,5,6}

¹Werner Reichardt Centre for Integrative Neuroscience and Institute of Theoretical Physics, University of Tübingen, 72076 Tübingen, Germany, ²Bernstein Center for Computational Neuroscience Tübingen and ³Max Planck Institute for Biological Cybernetics, 72076 Tübingen, Germany, ⁴Department of Neuroscience, Baylor College of Medicine, and ⁵Michael E. DeBakey Veterans Affairs Medical Center, Houston, Texas 77030, and ⁶Department of Computational and Applied Mathematics, Rice University, Houston, Texas 77005

Orientation tuning has been a classic model for understanding single-neuron computation in the neocortex. However, little is known about how orientation can be read out from the activity of neural populations, in particular in alert animals. Our study is a first step toward that goal. We recorded from up to 20 well isolated single neurons in the primary visual cortex of alert macaques simultaneously and applied a simple, neurally plausible decoder to read out the population code. We focus on two questions: First, what are the time course and the timescale at which orientation can be read out from the population response? Second, how complex does the decoding mechanism in a downstream neuron have to be to reliably discriminate between visual stimuli with different orientations? We show that the neural ensembles in primary visual cortex of awake macaques represent orientation in a way that facilitates a fast and simple readout mechanism: With an average latency of 30–80 ms, the population code can be read out instantaneously with a short integration time of only tens of milliseconds, and neither stimulus contrast nor correlations need to be taken into account to compute the optimal synaptic weight pattern. Our study shows that—similar to the case of single-neuron computation—the representation of orientation in the spike patterns of neural populations can serve as an exemplary case for understanding the computations performed by neural ensembles underlying visual processing during behavior.

Introduction

How do populations of neurons represent information in their joint firing patterns and how can this information be read out by downstream neurons (Averbeck et al., 2006)? In the motor system, for instance, numerous studies have applied decoding methods to obtain a detailed understanding of the population code for movement direction (Georgopoulos et al., 1986; Hatsopoulos and Donoghue, 2009). In the visual system, however, the situation is different: Ever since Hubel and Wiesel discovered that neurons in the primary visual cortex (V1) of cats and monkeys are tuned to the orientation of a visual stimulus (Hubel and Wi-

esel, 1959, 1968), orientation selectivity of single neurons and the underlying mechanisms have been studied extensively (De Valois et al., 1982; Ferster and Miller, 2000; Ringach et al., 2002). Despite this intense attention, little is known about how populations of V1 neurons represent orientation, in particular in alert primates. Only recently, decoding methods have been applied to population recordings from primate V1 to infer optimal spatial and temporal pooling rules for stimulus detection (Chen et al., 2006, 2008) and to investigate the importance of correlations for decoding from the spiking responses of V1 population in anesthetized monkeys (Graf et al., 2011). In addition, a similar approach has been used to investigate the representation of dynamic stimulus sequences in the anesthetized cat (Benucci et al., 2009).

Here, we study the population code for orientation in V1 by simultaneously recording spiking activity from populations of up to 20 well isolated V1 neurons in alert macaque monkeys, while the animals are viewing static gratings with different orientations and contrasts. We apply a simple neurally plausible decoder to discriminate between two orientations based on the population response patterns in short time windows. We focus on two questions: First, what are the time course and the timescale at which orientation can be read out from the population response? Second, how complex does the decoding mechanism in a downstream neuron have to be to reliably discriminate between visual stimuli with different orientations? The answers to these questions place constraints on the perceptual performance the popu-

Received March 19, 2012; revised May 3, 2012; accepted May 24, 2012.

Author contributions: P.B. and A.S.T. designed research; P.B., A.S.E., and A.S.T. performed research; R.J.C., W.J.M., and M.B. contributed unpublished reagents/analytic tools; P.B. analyzed data; P.B., A.S.E., and A.S.T. wrote the paper.

This work was partially supported by the German National Academic Foundation (P.B.); by the German Ministry of Education, Science, Research and Technology through the Bernstein Award (FKZ 01GQ0601) (M.B.) and the Bernstein Centre for Computational Neuroscience (FKZ 01GQ1002); the German Excellency Initiative through the Centre for Integrative Neuroscience Tübingen (EXC307); the Max Planck Society; National Eye Institute—National Institutes of Health Grant R01 EY018847 (A.S.T.); The Arnold and Mable Beckman Foundation Young Investigator Award (A.S.T.); and The McKnight Endowment Fund for Neuroscience Scholar Award (A.S.T.). We thank M. Subramanian, T. Shinn, and D. Murray for technical assistance.

This article is freely available online through the *JNeurosci* Open Choice option.

Correspondence should be addressed to either of the following: Philipp Berens, Werner Reichardt Centre for Integrative Neuroscience, Ottfried-Müller-Straße 25, 72076 Tübingen, Germany, E-mail: philipp@bethgelab.org; or Andreas S. Tolias, Baylor College of Medicine, Department of Neuroscience, One Baylor Plaza, S553, Houston, TX 77030, E-mail: astolias@bcm.edu.

DOI:10.1523/JNEUROSCI.1335-12.2012

Copyright © 2012 the authors 0270-6474/12/3210618-09\$15.00/0

lation code can support and constrain the neural machinery necessary to perform further processing downstream.

We find that the population code for orientation in the primary visual cortex of alert macaques is remarkably fast and simple: The decoder can discriminate between two orientations already after 30–80 ms, using the spike count in a time window as small as 30 ms, and its performance is best during the initial transient phase of the trial. Surprisingly, a fixed decoder can perform well throughout the whole trial and even across different contrasts, leading to a largely contrast-invariant population code. Finally, we show that—at the population size studied here—the decoder does not have to take correlations between neurons into account, substantially simplifying the process of learning the optimal readout weights.

Materials and Methods

Electrophysiological recordings. Recordings were made from two adult male rhesus monkeys (*Macaca mulatta*) using chronically implanted arrays of 12 and 24 custom-build tetrodes (monkeys D and H, respectively). The general methods, including implant design and surgical procedures, have been described previously (Tolias et al., 2007; Ecker et al., 2010). All procedures were conducted in accordance with the guidelines of the local authorities (Regierungspräsidium) and the European Community (EUVD 6/609/EEC) for the care and use of laboratory animals (monkey D) and the Baylor College of Medicine Institutional Animal Care and Use Committee-approved protocols and the NIH *Guide for the Care and Use of Laboratory Animals* (monkey H).

Spike sorting was done by fitting a mixture of Gaussians model to the detected and aligned spikes after reducing dimensionality by principal component analysis. The mixture model obtained from the spike sorting algorithm allows assessing single-unit isolation quantitatively. Details of spike sorting and quantification of single-unit isolation have been described previously (Tolias et al., 2007). We used only very well isolated cells for which the sum of estimated false positives and misses did not exceed 5% of their total number of spikes. We used only neurons whose average firing rate during stimulation across orientation and contrast exceeded 0.1 Hz.

For most of the analysis in the paper, we used a dataset of 17 sessions (monkey D, 12; monkey H, 5) in which static gratings were shown (see below for details on visual stimulation). According to our isolation and firing rate criteria, those sessions contained data from 6 to 20 neurons each (average population size, 10). For the control analysis of Figure 6, we used an additional dataset with 27 sessions (monkey D, 14; monkey H, 13) in which drifting gratings were shown. This dataset contained 5–19 neurons per session (average population size, 10.7). These datasets have also been analyzed in the study by Ecker et al. (2010).

Behavioral paradigm and visual stimulation. For the main dataset, static sine wave gratings were displayed on CRT monitors. Trials were initiated by a sound and the appearance of a colored fixation target ($\sim 0.15^\circ$) at the center of the screen. Three hundred milliseconds after the monkey acquired fixation, the stimulus was presented for 500 ms. The monkey had to maintain fixation for another 300 ms until the fixation target disappeared. Upon successful completion of a trial, it was rewarded by a drop of juice. The animals were implanted with a scleral search coil and required to fixate within a radius of $0.5\text{--}1^\circ$.

The sine wave gratings always covered and extended beyond the receptive fields of all neurons recorded from the array (typical size, 4.5° ; spatial frequency, 4 cycles/ $^\circ$; fixed phase across trials). They were presented at eight different orientations and two different contrasts in each session (1, 2, 3, 10, 20, 30, 50, and 100% Michelson contrast were used). For each condition, we collected between 10 and 85 trials per stimulus condition. In four sessions, only seven of eight orientations were shown.

We labeled contrast levels up to 3% as “low,” between 10 and 30% as “medium,” and $>50\%$ as “high.” According to this scheme, 8 datasets were recorded at low contrast, 16 datasets at medium contrast, and 10 datasets at high contrast, with each session resulting in two datasets with two different contrasts (see above).

For the dataset used in the control analysis of Figure 6, drifting sine wave gratings were used (contrast, 100%; drift speed, 3.4 cycles/s). All other stimulation procedures were identical to the main dataset.

Preprocessing. We binned the spike trains at a resolution of 10 ms. Binned spike trains were convolved with a boxcar filter, typically 50 ms wide (except when otherwise noted). Times were aligned to the end of the boxcar window (not the center) such that the bin with the label “100 ms” contains the number of spikes fired by the neuron between 50 and 100 ms after stimulus onset.

Decoding analysis. We probe the population representation of orientation using a classification task. We used regularized logistic regression (Bishop, 2007) to decode the stimulus orientation from the population response $X[t] = (X_1[t], \dots, X_N[t])^T$ on a single trial, where $X_j[t]$ is the number of spikes fired by neuron j in the t th time bin and N is the number of neurons in the population. The task of the decoder was to decide whether this population response vector occurred in a trial with stimulus orientation θ_1 or θ_2 . Such a classification task is commonly used in psychophysical studies in humans and monkeys (Vázquez et al., 2000) and has also proven useful to theoretically assess the quality of different population coding schemes (Berens et al., 2011).

Logistic regression is a generalized linear model for two-class classification, where the posterior over the class is a log-linear function of the population response (Bishop, 2007) as follows:

$$P(\theta = \theta_1) = \sigma(w^T X[t]). \quad (1)$$

Here, $\sigma(a) = 1/(1 + \exp(-a))$ is the logistic function and w is a vector of weights. We trained the logistic regression model with L_2 regularization using the glmnet toolbox (Friedman et al., 2010) in MATLAB (Mathworks). The instantaneous decoder is a decoder that was trained separately for each time bin t and each combination of stimulus orientations. We deviated from this scheme to test certain properties of the population code as described below.

We performed cross-validation using the repeated random subsample technique (80% training data, 20% test data, balanced between both classes, 100 repetitions) for the whole regularization path (100 regularization parameters logarithmically spaced between 10^{-10} and 10). We apply the “1 SE” rule (Friedman et al., 2010) and report the decoding error/percentage correct of the most strongly regularized decoder within 1 SE of the decoder with the best decoding performance. This is a conservative estimate of the true decoding performance.

In several cases, we made modifications to the decoder to test whether the readout quality would be affected by those changes. (1) The cumulative decoder was trained and tested on the spike count across the whole trial. (2) The constant decoder was trained on the spike count across the whole trial (like to the cumulative one) and tested on the spike count for each specific time bin individually (in contrast to the cumulative one). (3) The decoding weights assuming that the spike counts of the population follow an independent Poisson distribution were computed as follows:

$$w_i = \frac{f_i(\theta) - f_i(\theta + \Delta\theta)}{0.5(f_i(\theta) + f_i(\theta + \Delta\theta))} \quad (2)$$

(for derivation, see Ma, 2010). (4) Contrast invariance was tested in two ways: First, by training the decoder on the combined trials of both high and low contrast and testing it on either one of them (contrast-independent decoder). For training, we randomly selected 50% of the trials from each contrast to ensure that the contrast-independent decoder had the same amount of information available for learning as the specialized one. Second, by training a decoder on the high-contrast trials and evaluating its performance on the low-contrast trials of the same population and vice versa (cross-contrast decoder). (5) The weight vector of the linear logistic regression decoder is influenced by the average spike count and the covariance matrix of the neurons. To test for the importance of covariations in the firing rates for the decoding performance, we trained the logistic regression model on trial-shuffled data. The performance of the decoder was then evaluated on the original da-

taset (Latham and Nirenberg, 2005). In addition, we tested whether the variances and covariances of the firing rates contained extra information not contained in the spike counts by incorporating the quadratic features $\sqrt{X_i X_j}$ as additional features into the logistic regression decoder as one specific example of a nonlinear readout.

Population model. We use a standard population model consisting of 200 neurons with homogenous orientation tuning functions modeled by cosine-like tuning functions (Berens et al., 2011). The model population has two subpopulations, differing in the width of their tuning curve and the semisaturation contrast of their contrast response functions. Contrast response functions were modeled as:

$$f(c) = \frac{c^n}{c_{50}^n + c^n} \quad (3)$$

where n is a constant determining the steepness of the curve. The semisaturation contrast c_{50} was set to 5% for subpopulation 1 and to 50% for subpopulation 2. Thus, at low contrast, only the broadly tuned subpopulation was activated, while at high contrasts >50% of both subpopulations were activated. For illustration, the steepness parameter n was set to 20, a value higher than what is commonly observed experimentally ($\sim 1-5$) (Albrecht and Hamilton, 1982). This was done to achieve a fast transition between contrasts in which only the broad subpopulation was activated and contrasts in which both subpopulations were activated. We calculated the minimum discrimination error for two gratings 10° apart (Berens et al., 2011). We used a short time window (50 ms) and set the peak firing rate of the neurons to 50 Hz, the baseline rate to 5 Hz.

Other statistical methods. For an estimate of the decoding performance achieved by single neurons, we computed:

$$d'_i = \frac{\mu_i(\theta_1) - \mu_i(\theta_2)}{\sqrt{(\lambda_i^2(\theta_1) + \lambda_i^2(\theta_2))/2}} \quad (4)$$

where $\mu_i(\theta)$ and $\lambda_i^2(\theta)$ are the average spike count and the variance of neuron i in response to stimulus θ and the dependence on t has been omitted for clarity. From d' we compute the classification error (Duda et al., 2000) by $1 - \Phi(d'/2)$, where $\Phi(x)$ is the cumulative normal distribution function. The population d' was computed by inverting this formula and applying it to the discrimination error of the population decoder.

The Fano factor was computed as

$$F = \frac{\lambda_i^2[t]}{\mu_i[t]} \quad (5)$$

for each neuron and time window. Time bins with no spikes were omitted from further analysis.

Latency was computed as the time at which the decoding performance reached 75% of its maximal value.

Relative performance was computed as the ratio of the performance gains over chance level, $(P_c(A) - 0.5)/(P_c(B) - 0.5)$, where $P_c(A)$ is the performance of the decoder (in fraction correct) in condition A. Therefore, a decoder that achieves 55% correct in condition A and 60% in condition B has a relative performance of 50%.

The two-way ANOVA for testing the effect of (1) the difference of the preferred orientation of a neuron to the decision boundary and (2) the difference in the orientation of the stimulus, $\Delta\theta$, was computed on the average weights between 50 and 250 ms after stimulus onset. We computed the average weight per session for a given combination of the two factors and treated the different sessions as replicates ($N = 34$).

Noise correlations for pairs of neurons were measured by the correlation coefficient

$$r_{sc} = \frac{\text{Cov}[x, y | \theta]}{\sqrt{\text{Var}[x | \theta]\text{Var}[y | \theta]}} \quad (6)$$

of the spike counts x, y of the two neurons. Spike counts were computed in the whole trial of 500 ms duration.

If not otherwise stated, all reported average values are medians with SEs. The SE of the median was computed using the MATLAB function bootstrap with 1000 bootstrap samples.

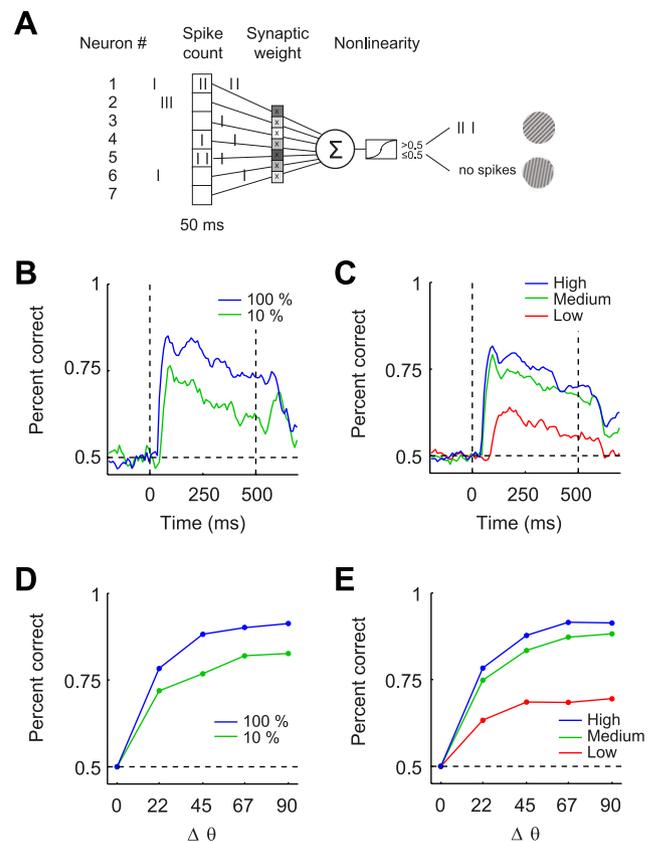


Figure 1. *A*, Illustration of the decoder as a linear–nonlinear model neuron. It adds the spikes of all presynaptic neurons in a small time window and multiplies them with a synaptic weight. The result is summed and passed through a nonlinearity, and a threshold is applied. This indicates which of two gratings was presented. *B*, Time course of the decoding performance for an example session with 20 neurons recorded at 10 and 100% contrast (green and blue lines, respectively) averaged over $\Delta\theta$. *C*, Average time course for all 17 sessions sorted by contrast (low contrast: 1–3%, red; medium contrast: 5–30%, green; high contrast: 50–100%, blue) averaged over $\Delta\theta$. *D*, Neurometric functions for the example session in *B* with decoding performance as a function of $\Delta\theta$. The color code is as in *B*. *E*, Average neurometric functions for all 17 sessions sorted by contrast. The color code is as in *D*.

Results

Our goal was to study how orientation is represented in the joint activity patterns of neural ensembles in primary visual cortex of awake primates. To this end, we recorded simultaneously from populations of up to 20 well isolated V1 neurons in two awake macaques using chronically implanted tetrode arrays (Tolias et al., 2007; Ecker et al., 2010). The animals were viewing static gratings with different orientations and contrasts (17 sessions, 6–20 neurons; two contrasts per session; for details, see Materials and Methods). We took a decoding approach to study how fast and on what timescale orientation could be read out from the neural activity and how this readout depended on contrast or correlations. We used logistic regression (Bishop, 2007) to discriminate between two gratings of different orientation based on the instantaneous population response patterns in short time windows of 50 ms (see Materials and Methods), which may be biophysically plausible for dendritic integration (London and Häusser, 2005). Reminiscent of a simplified linear–nonlinear neuron model (Schwartz et al., 2006), the logistic regression decoder adds the weighted spike counts of its inputs, applies a sigmoidal nonlinearity, and decides whether the output is larger or smaller than a threshold to indicate whether the population pattern is a response to one or the other of the two orientations

(Fig. 1A; see Materials and Methods, Eq. 1). We trained a specific decoder on each time bin and studied how the classification performance of this decoder evolved as a function of time and changed with the difference between the orientations of the gratings and their contrast (Fig. 1).

We found that the readout of orientation is most accurate during the transient phase of the neural response. The decoding performance rose sharply after trial onset with a latency of 77.5 ± 4.1 ms (75% of the peak performance; Figs. 1B,C, 2A). This means that the decoder achieved 75% of the peak performance using the spikes in the time bin between 30 and 80 ms after stimulus onset. On average, the peak performance was reached after 120.0 ± 19.0 ms (Fig. 2B). After this, the performance typically decayed until the end of the trial (Fig. 1B,C): In 30 of 34 datasets, the decoder achieved the best discrimination performance during the initial transient phase of the neural response (90–130 ms). During this time window, the relative performance compared with the sustained response (300–400 ms) was $125.9 \pm 9.2\%$ (Fig. 2C). An additional peak in decoding performance during the off-response at the end of each trial was present in 3 of 17 sessions (data not shown).

Both the difference between the orientations of the gratings ($\Delta\theta$) and the contrast had an effect on the latency of the population readout and its overall performance: The latency decreased with increasing $\Delta\theta$ (Fig. 2D; two-way ANOVA on latencies with factors $\Delta\theta$ and contrast; main effect of $\Delta\theta$, $p = 0.021$) and with increasing contrasts (Fig. 2A; Spearman's rank correlation: -0.63 , $p < 0.001$). Similarly, the decoder achieved higher performance at higher contrasts, both on average (Fig. 1C,E) and in 15 of 17 sessions individually (Fig. 2E, binomial test, $p < 0.001$), and classification performance was highest for orthogonal gratings (Fig. 1D,E).

We next changed the size of the time window used to count the input spikes for the decoder to study the timescale necessary to reliably discriminate two orientations based on the population response. In addition, we asked whether a decoder that can count spikes over hundreds of milliseconds outperforms one that can count spikes only in a small time window of tens of milliseconds. We found that an instantaneous readout based on the spikes in a small time window of tens of milliseconds is sufficient to achieve good decoding performance. Specifically, the relative performance was hardly changed if the duration of the integration window was shortened to 30 ms or extended to 70 ms (Fig. 3A; relative performance, 92 and 105%, respectively). Using an integration window of only 10 ms, the decoder still achieved a relative performance of 67% on average. We explicitly tested whether the readout could be improved by accumulating spikes over longer periods of time (cumulative decoder; see Materials and Methods). We found that the instantaneous decoder with an integration window of 50 ms achieved a relative performance of $78.7 \pm 2.1\%$ compared with the cumulative decoder across the whole trial (Fig. 3B, dark). For comparison, an instantaneous readout of single neurons achieved only $34.6 \pm 0.8\%$ relative performance compared with the cumulative one (Fig. 3B, light).

We next wondered to what extent single-neuron properties such as firing rate or response variability determine the time course of the population readout accuracy. The firing rates of the neurons in our sample showed a transient increase right after stimulus onset, after which the firing rates decayed until the end of the trial (Fig. 4A; average firing rate before stimulus onset, 4.02 Hz), in line with earlier reports (Müller et al., 2001). The response variability as measured by the Fano factor dropped after stimulus onset before returning back to its baseline value (Fig.

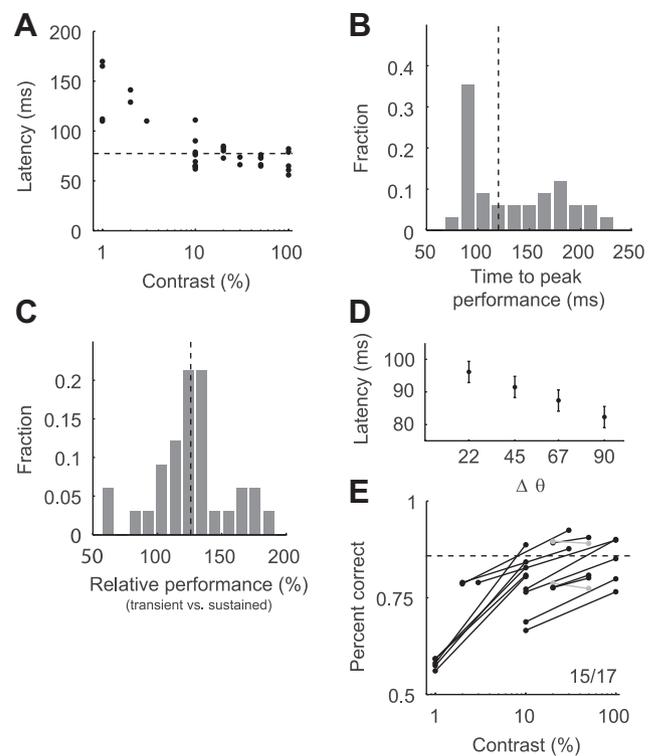


Figure 2. *A*, Latency as a function of contrast. There are 34 data points in total, 2 for each of the 17 sessions. The dashed line indicates the median latency. *B*, Distribution of the time to peak performance. The dashed line indicates the median time to peak performance. *C*, Distribution of the relative performance of the transient period (80–130 ms) versus the sustained period (300–400 ms). The dashed line indicates the median relative performance. *D*, Mean latency as a function of the difference in orientation, $\Delta\theta$. *E*, Peak performance as a function of contrast. The solid lines connect the two data points belonging to the same session. The black dots and lines indicate sessions in which the decoding performance increased with contrast, and the gray dots and lines sessions in which the performance decreased with contrast. The dashed line indicates median peak performance.

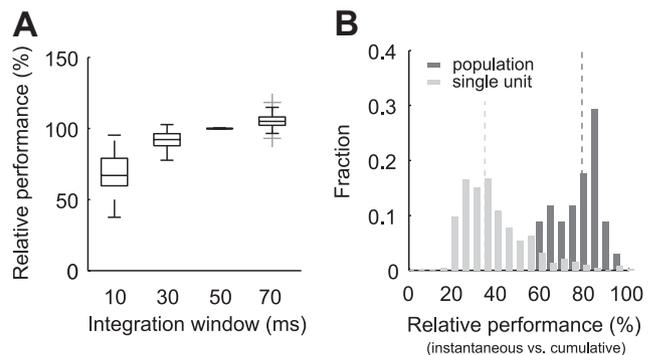


Figure 3. *A*, Relative peak performance using the integration window size denoted on the x-axis compared with the standard integration window size of 50 ms. The peak performance is computed for each length of the integration window separately and was defined as the performance in the time window that gave the best performance with that integration time. *B*, Distribution of the relative peak performance of the instantaneous decoder compared with the cumulative decoder for populations (dark) and single neurons (light).

4B; average Fano factor before stimulus onset, 1.04), akin to what has been observed in V1 and other brain areas before (Müller et al., 2001; Churchland et al., 2010). Interestingly, contrast had slightly different effects on the firing rates and the response variability: During the transient phase of the response, the firing rates hardly differed between the medium and high contrast condition,

while the Fano factor showed a pronounced decrease from medium to high contrast (Fig. 4*A,B*). The time course of the single-unit discriminability reflects both the firing rate and response variability modulations (Fig. 4*C*): It peaks during the transient phase of the response (Müller et al., 2001) and closely resembles the time course of the population readout accuracy (Fig. 4*C,D*) (see also Fig. 1*C*). Thus, the time course of the accuracy of the population representation seems largely determined by the properties of its constituent neurons (see also below for an analysis of the importance of correlations).

After having established that an instantaneous decoder can provide a fast and reliable readout of the population response, we turned to study the weights of the population decoder. We investigated how they evolve during the trial or change with contrast and whether they depend on noise correlations between neurons. We found that, on average, the readout weights did not change over the time of the trial (Fig. 5*A,B*). They depended on the distance of the preferred orientation of a neuron to the “decision boundary” in the middle between the two discriminated orientations (Fig. 5*B,C*; for statistical analysis, see below), with substantial variability around this mean (Fig. 5*D*). This variability is likely partially due to the small populations sampled as well as true tuning curve and firing rate variability in the neural population. Interestingly, the average weight profile did not depend on contrast (Fig. 5*C*), but was affected by the orientation difference between the two stimuli (Fig. 5*E*): If the stimulus orientations were closer together, neurons with preferred orientations closer to the decision boundary were weighed more strongly than if they were farther apart (two-way ANOVA, main effect of $\Delta\theta$, n.s.; main effect of the difference of the preferred orientation relative to the decision boundary, $p < 10^{-10}$; interaction, $p < 10^{-10}$; see Materials and Methods). For $\Delta\theta = 22.5^\circ$, the weight profile closely resembled that of an optimal decoder reading out a population with spike counts following an independent Poisson distribution (Fig. 5*F*).

As the average weights do not depend on time, we expect a constant decoder trained on the spike count in the whole trial to perform as well as the instantaneous decoder trained on each time bin separately. We found that this was indeed the case, in particular during the sustained phase of the trial (Fig. 6*A,B*; relative peak performance, $92.3 \pm 1.3\%$). Note that this decoder is trained on the average spike count during a trial just as the cumulative decoder, but tested on the spike counts in each individual time bin. To rule out that the success of the constant decoder is due to the constant phase of the stimulus, we performed the same analysis in an additional dataset in which moving gratings were used for stimulation (see Materials and Methods). We found that the readout performance of the constant decoder remained close to the bin-based decoder (Fig. 6*C*; relative peak performance, $91.2 \pm 1.8\%$). Also, the drop in peak performance was similar for both static and moving gratings (Fig. 6*D*). Therefore, the population code seems largely phase invariant, consistent with the large fraction of complex cells

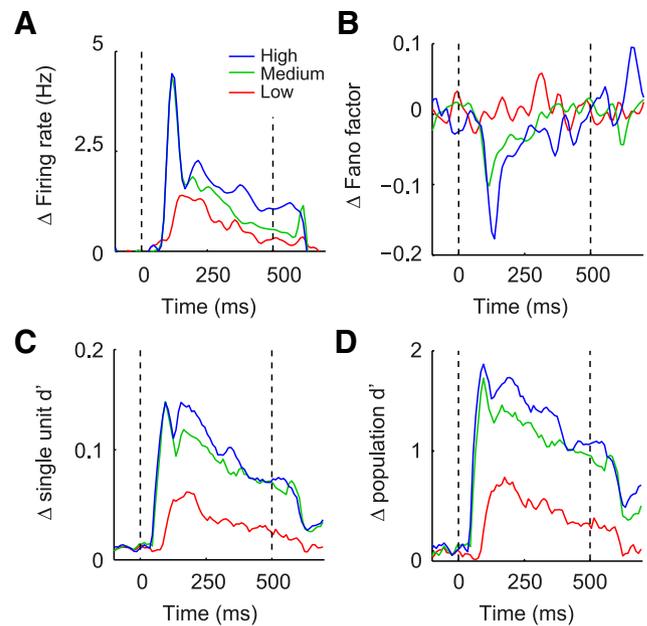


Figure 4. *A*, Average firing rate change after stimulus onset compared with baseline (200 to 0 ms before stimulus onset) for all single neurons sorted by contrast (low, medium, and high contrast in red, green, and blue, respectively) averaged over all orientations. Seventy-seven neurons were recorded at low contrast, 180 neurons at medium contrast, and 119 neurons at high contrast. The firing rate profile was smoothed using a Gaussian kernel with SD of ~ 30 ms. *B*, Average Fano factor change after stimulus onset computed and smoothed as in *A*. *C*, Average single-unit discriminability d' change averaged across all combinations of $\Delta\theta$. *D*, Average population discriminability d' change averaged across all combinations of $\Delta\theta$.

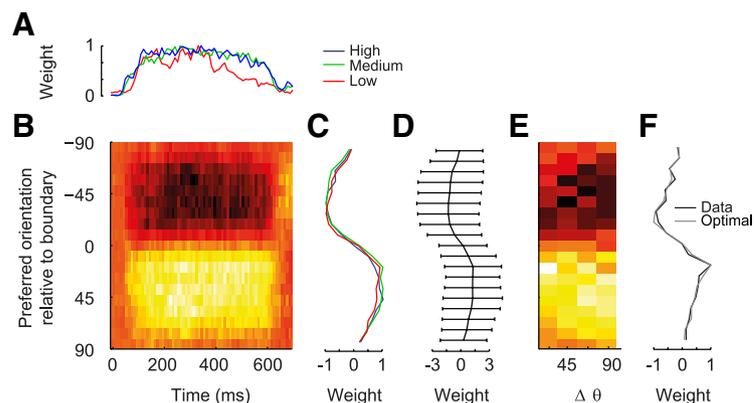


Figure 5. *A*, Average absolute weight as a function of time for the three contrast groups (low, medium, and high contrast in red, green, and blue, respectively), renormalized to a maximum of 1 for each contrast separately. *B*, Average weight as a function of time and the distance of the preferred orientation of a neuron to the decision boundary in the middle of the two orientations that were discriminated. The average is taken across neurons. Weights for medium contrast are shown. The dark color indicates negative weights, and the bright color, positive weights. *C*, Average weight profile as a function of the distance of the preferred orientation of a neuron to the decision boundary for the three contrast groups (low, medium, and high contrast in red, green, and blue, respectively), renormalized to a maximum of 1 for each contrast separately. *D*, As in *C*, but with variability (all contrasts collapsed). Error bars show 1 SD. *E*, Average weight as a function of the distance between the two gratings that were discriminated and the distance of the preferred orientation of a neuron to the decision boundary. The dark color indicates negative weights, and the bright color, positive weights. *F*, The normalized weights of the logistic regression decoder (black) compared with the optimal weights assuming that the spike counts of the population follow an independent Poisson distribution (gray) for $\Delta\theta = 22.5^\circ$.

found in macaque V1 in all layers except layer 4C (Ringach et al., 2002).

In addition, we expect the population readout to be contrast invariant to some degree, as the average weight profile also does not depend on contrast (Fig. 5*C*). In this case, the weights learned at one contrast could be used to decode orientation at another. To

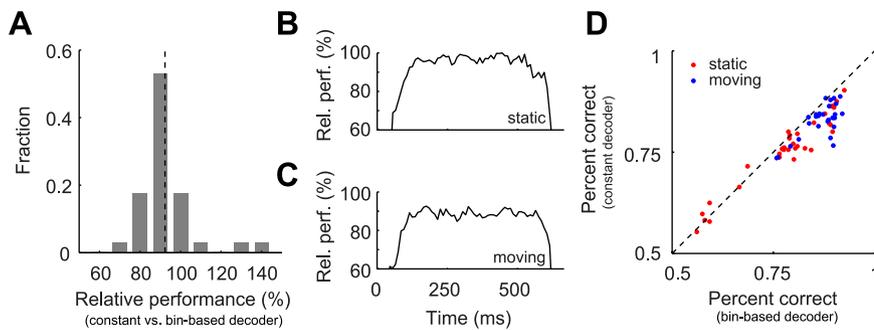


Figure 6. *A*, Relative peak performance of the constant decoder compared with the bin-based one. The dashed line indicates median relative performance. *B*, Time course of the relative performance of the constant decoder compared with the bin-based one. *C*, As in *B* for an additional dataset of sessions in which drifting gratings were shown. Despite the continually changing phase, the relative performance of the constant decoder remains high compared with the bin-based one. *D*, Comparison of the peak performance using a bin-based or constant decoder for both static (red) and moving (blue) gratings.

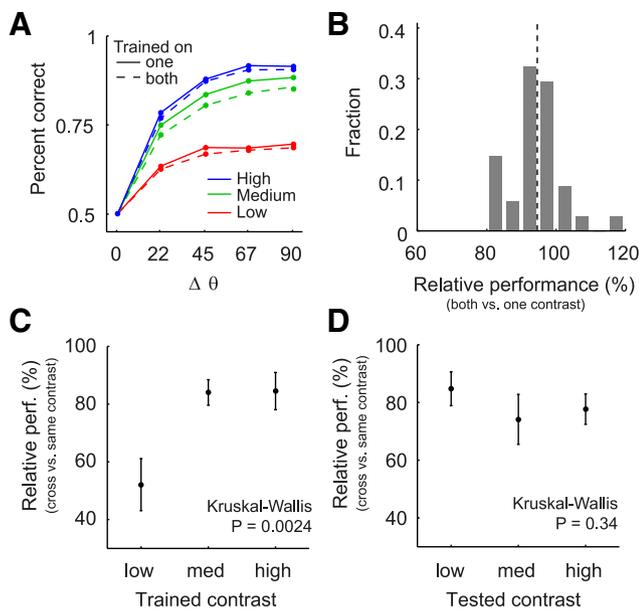


Figure 7. *A*, Average neurometric functions sorted by the contrast the decoder was tested on (low, medium, and high contrast in red, green, and blue, respectively) showing decoding performance as a function of $\Delta\theta$, for both a decoder trained and tested on the same contrast (solid) and a decoder trained on data from both contrasts recorded during a single session and tested on one of them (dashed). *B*, Distribution of the relative peak performance comparing a decoder trained on both contrasts and tested on one of them to one trained and tested on the same contrast. The dashed line indicates median relative performance. *C*, Average relative peak performance comparing a decoder trained on one contrast and tested on the other (“cross-contrast decoder”) to one trained and tested on the same contrast (“same-contrast decoder”) sorted by the contrast level trained on. Error bars denote SEs of the median. *D*, As in *C*, but sorted by the contrast level tested on.

test this property of the code, we first trained the decoder on the data acquired at the two contrasts from one session and evaluated it at one of them (contrast-independent decoder; see Materials and Methods). For this and all following analyses, we returned to the static grating data. The contrast-independent decoder performed almost as well as the contrast-specific decoder (Fig. 7*A, B*; $94.6 \pm 1.1\%$ relative performance). This means that the representations of orientation at the two contrast levels do not interfere with each other—in this case, we would have expected strongly reduced classification performance. We next used a cross-contrast decoder that was trained on one contrast and tested on the other contrast presented during that session. We

compared it with the contrast-specific decoder trained and tested on the same contrast (Fig. 7*C, D*). We first analyzed the relative decoding performance conditioned on whether the decoder was trained on a contrast belonging to the low, medium, or high contrast group (averaging across all test contrasts). If the decoder was trained on medium or high contrast, the relative performance of the cross-contrast decoder compared with the specialized decoder was high (84.0 and 84.5% at medium and high contrast, respectively). When trained on low contrast, the relative performance was significantly lower (Fig. 7*C*; 52.1%; Kruskal–Wallis test, $p = 0.0028$). Next, we analyzed the relative performance depending on which

contrast level the decoder was tested on (averaging across all training contrasts). We found that the relative performance did not change depending on the contrast level the decoder was tested on (Fig. 7*D*; 84.8, 74.1, and 77.7%, respectively; Kruskal–Wallis test, $p = 0.25$). These results indicate that a contrast-invariant decoder can indeed achieve almost as good performance as a contrast-specialized one. Why is the relative performance of the decoder significantly impaired when trained on low contrast? At low contrasts, not all necessary information is present to determine the decoding weights for all neurons accurately. This can occur, for example, if a neuron does not fire at all, because the contrast is far below its semisaturation contrast. If the weights of the decoder are determined in a way such that the correct weights are assigned to each neuron at high enough contrast, contrast-invariant orientation readout is possible.

Finally, noise correlations, shared trial-to-trial fluctuations in the population response, are believed to affect the population code (Averbeck et al., 2006). In the anesthetized macaque, they are typically found to be on the order of ~ 0.1 – 0.2 (Kohn and Smith, 2005; Smith and Kohn, 2008). Here, they contribute $\sim 10\%$ to the total information about the orientation of stimulus contained in pairs of neurons (Montani et al., 2007). Furthermore, ignoring these noise correlations found in anesthetized animals in a population readout can lead to a significant loss in decoding performance (Graf et al., 2011).

In contrast, we found that the decoding weights in the awake macaque can be learned without knowledge of the correlation structure of the population. In our data, noise correlations were small (Ecker et al., 2010) to begin with (0.005 ± 0.002 , mean \pm SEM; Fig. 8*A*; average across pairs from all tetrodes), but the average correlation level was slightly higher at low contrasts compared with high ones (one-way ANOVA, $p = 0.0044$; Fig. 8*A*, inset). We studied whether the decoder would have to be aware of these noise correlations for adjusting its readout weights. We found that a decoder trained on trial-shuffled data (and tested on the original data; see Materials and Methods) performed as well as a decoder with access to the correlation structure of the population response (Fig. 8*B*; relative performance, $100 \pm 0.2\%$), independent of the contrast level (Fig. 8*B*, inset; one-way ANOVA, $p = 0.74$). This confirms our analysis above that the weights of the logistic regression decoder closely resemble the weights derived from assuming that the population has an independent Poisson spike count distribution (Fig. 5*F*).

We also tested whether allowing the decoder to explicitly take correlations into account improves decoding performance. To

this end, we added quadratic features to the logistic regression decoder, resulting in a nonlinear readout. This is different from our analysis above, in which the decoder was linear but once trained on the original data and once on the trial-shuffled data. Since we had only few trials for many sessions, this resulted in a slight drop in performance overall (relative performance quadratic vs linear decoder, 96.6%), as has been reported for the anesthetized macaque (Graf et al., 2011). Thus, quadratic features—and thus more complex nonlinear decoders—are unlikely to improve classification performance substantially.

Discussion

Orientation tuning in single neurons has been a classic model for understanding single-neuron computation. In the same way, the representation of orientation in the spike patterns of neural populations can serve as an exemplary case for the understanding of the computations performed by neural ensembles underlying visual processing during behavior. Our study is a first step toward that goal. We showed that the neural ensembles in primary visual cortex of awake macaques represent orientation in a way that facilitates a fast and simple readout mechanism: With an average latency of 30–80 ms, the population code can be read out instantaneously with a short integration time of only tens of milliseconds and neither stimulus contrast nor correlations need to be taken into account to compute the synaptic weight pattern.

Temporal aspects of the population readout

The accuracy of the population readout was highest during the transient phase of the trial and followed the temporal modulation of the firing rate and variability of the single neurons. Thereby, our findings extend earlier work by Müller et al. (2001)

to neural populations. Interestingly, the weights of the population readout remain largely constant over the course of the trial, even if the phase of the grating is continually changing. This finding may seem surprising given the complex temporal tuning dynamics previously reported for V1 neurons (Ringach et al., 1997). However, Ringach et al. used briefly flashed gratings to map the tuning functions measuring the “impulse response” of the neurons. In contrast, we measured the “step response” by presenting the stimulus continuously for several hundred milliseconds possibly averaging over tuning dynamics reported in the study by Ringach et al. Using briefly flashed gratings, Benucci et al. (2009) studied the population representation of orientation sequences in V1 of anesthetized cats using a decoding approach. The study focused on the transitions between stimuli of different orientation and found that a simple instantaneous readout can perform well despite the interactions between the representations of successive stimuli. This suggests that our finding—a constant instantaneous decoder is sufficient to read out the population representation—may extend even to complex stimulus sequences.

How contrast-invariant is the population code?

The computation of invariant representations is difficult (Rust and Stocker, 2010). We found that the population code in pri-

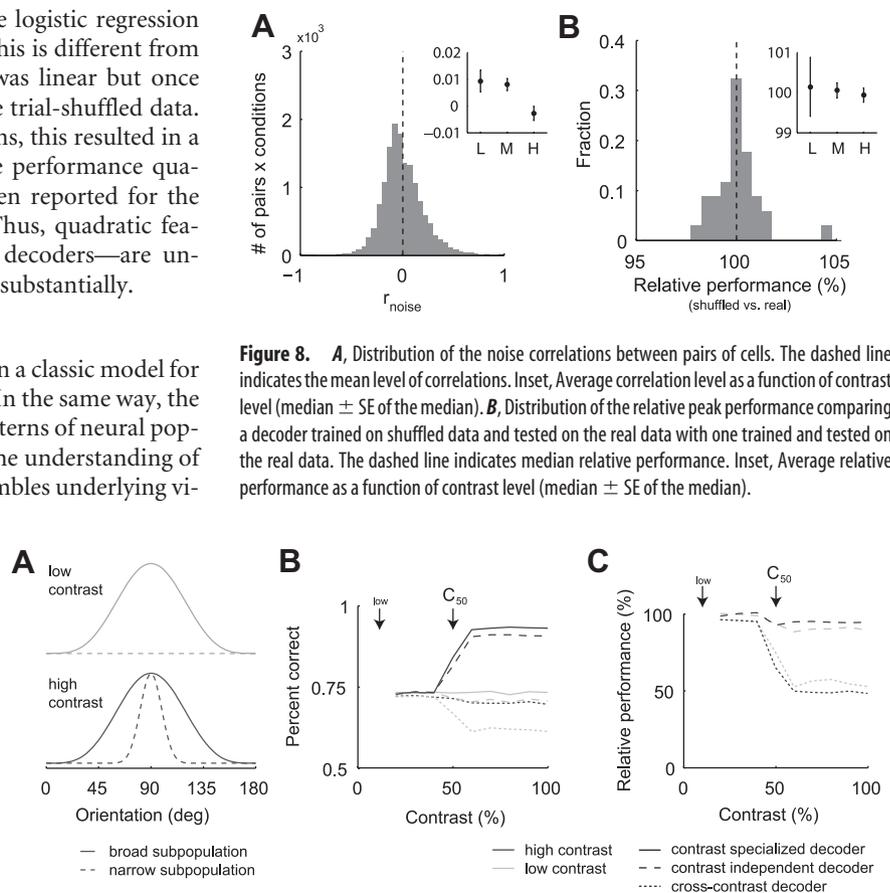


Figure 8. *A*, Distribution of the noise correlations between pairs of cells. The dashed line indicates the mean level of correlations. Inset, Average correlation level as a function of contrast level (median \pm SE of the median). *B*, Distribution of the relative peak performance comparing a decoder trained on shuffled data and tested on the real data with one trained and tested on the real data. The dashed line indicates median relative performance. Inset, Average relative performance as a function of contrast level (median \pm SE of the median).

Figure 9. *A*, Illustration of the model with two subpopulations. The broad subpopulation has a semisaturation contrast of 5% and the narrow subpopulation of 50%. At low contrast (<50%), only the broad subpopulation is activated (top). At high contrasts (>50%), both populations are active (bottom). *B*, Performance of a decoder discriminating between two gratings 10° apart at low and high contrast using a contrast-specialized decoder (trained and tested on the same contrast; solid), a contrast-independent decoder (trained on both contrasts, tested on one; dashed), and a cross-contrast decoder (trained on one contrast, tested on the other; dotted). The lower contrast was 10% in all cases, while the higher contrast varied between 15 and 100%. *C*, Relative performance of the contrast-independent and cross-contrast decoders compared with the specialized decoder.

mary visual cortex exhibits a high degree of contrast invariance; that is, a readout neuron does not have to adjust its synaptic weights to a specific contrast. In contrast, a recent study reports that the population code is less contrast invariant in the anesthetized monkey [Graf et al., (2011), their supplemental material]. It is possible that the differences in the correlation structure of the population activity between the anesthetized and awake state (see below) are responsible for the different degrees of contrast invariance, as the magnitude of correlations in the anesthetized state strongly depends on contrast (Kohn and Smith, 2005).

One might assume that contrast invariance of the population code follows from the well known contrast invariance of single neurons (Sclar and Freeman, 1982) in a straightforward manner: because single neurons show contrast-invariant orientation tuning (i.e., the orientation tuning curve and the contrast response function factorize), the population activity will be contrast invariant as well. However, since neurons in V1 have differing tuning and contrast response functions (Albrecht and Hamilton, 1982; Ringach et al., 2002), this does not have to be the case. To illustrate this point, consider a model with two subpopulations, one broadly tuned and one narrowly tuned, in which the narrowly tuned neurons have a higher semisaturation contrast than the broadly tuned neurons (Ganmor et al., 2009). The activity profile of such a population is not contrast invariant, despite

contrast invariance at the single-neuron level (see Materials and Methods) (Fig. 9A). For this model, the contrast-independent decoder trained on data collected at both contrasts performs well (Fig. 9B,C), while the relative performance of a cross-contrast decoder is only $\sim 50\%$ for both high and low contrasts. When trained at low contrast, the decoder has no way to assign the proper weights to the narrowly tuned subpopulation as these neurons are not activated by the stimuli, similar to what we find in our data (Fig. 7C). When trained at high contrast, the decoder assigns higher weights to the narrowly tuned subpopulation than to the broadly tuned neurons as they are more informative about the stimulus. If this decoder is used at low contrast, the narrowly tuned neurons only add noise to the decoding process, leading to lower performance.

In our data, the population code achieves $\sim 80\%$ relative performance for the cross-contrast decoder (Fig. 7C,D); this is considerably more than could be achieved in the model, but direct comparisons between model and data are difficult as the performance of the cross-contrast decoder depends on the exact setting of the model parameters. The simple model illustrates, however, that a high degree of contrast invariance is not a trivial consequence of contrast invariance at the single-neuron level. In cats, semisaturation contrast and tuning width are independent (Busse et al., 2009), such that the resulting population code is indeed contrast invariant. Our study is the first to show that the population code in V1 in alert macaques shares this property. It will be interesting to explore how invariant the orientation representation is to changes in other parameters (e.g., spatial frequency or speed).

Are correlations important for decoding?

The magnitude and functional consequences of noise correlations have been subject to intense debate in recent years, both in the theoretical and the experimental literature (Averbeck et al., 2006; Smith and Kohn, 2008; Ecker et al., 2010, 2011; Cohen and Kohn, 2011). Extending our previous work (Ecker et al., 2010), we report here that the noise correlations measured in alert monkeys (~ 0.01) do not have to be taken into account when discriminating between different orientations, for both high and low contrasts. In the anesthetized monkey, ignoring correlations leads to significantly impaired readout performance even at a population size similar to ours (Graf et al., 2011). The reason for this discrepancy likely is that correlated variability is higher under anesthesia (~ 0.16) (Smith and Kohn, 2008) than in the awake state (~ 0.01) (Ecker et al., 2010), despite similar firing rates (average geometric mean firing rate, 3.4 vs 5 spikes/s, respectively) [see also Greenberg et al. (2008) and Renart et al. (2010) for direct comparisons]. This effect is possibly a result of ongoing brain state modulations that occur under anesthesia (Kohn et al., 2009; Ecker et al., 2010; Renart et al., 2010).

Our finding, however, is subject to some limitations: First, in populations of several hundred neurons, even very small correlations may become important for reading out the population code, in particular as they might imply stimulus specific higher-order correlation structures (Schneidman et al., 2006; Macke et al., 2011). Second, the strength of noise correlations may vary as a function of cortical layer (Hansen and Dragoi, 2011) or correlations in certain pairs of neurons may be stronger than the average correlations in the populations studied here (Ecker et al., 2010; Ko et al., 2011). In this case, a detailed characterization of the precise shape of the correlation structure is needed to determine the impact of noise correlations on neural coding. Third, we evaluated the performance of the population code in a fairly easy

task, when discriminating between coarsely spaced orientations. If the task is more difficult (e.g., when discriminating between finely spaced orientations), correlations may become important (Samonds and Bonds, 2004).

Implications for models of cortical computation

We have shown that the readout weights are independent of contrast and that potential improvements in decoding accuracy through quadratic features are small at best. These results provide evidence that the population variability in V1 is Poisson-like (i.e., it belongs to the exponential family with linear sufficient statistics) (Ma et al., 2006; Ma, 2010). This type of variability is a crucial assumption made in theoretical studies of probabilistic population codes. In such codes, uncertainty about the stimulus is represented in the neural activity in addition to information about the stimulus itself. Poisson-like variability is particularly suitable to facilitate near-optimal Bayesian computation in tasks such as cue combination (Ma et al., 2006), decision making (Beck et al., 2008), and visual search (Ma et al., 2011). However, experimental evidence that cortical population activity is Poisson-like has been lacking so far. It would be interesting to work out the details of how uncertainty is represented in a real V1 population and how it is used and transformed during perceptual decision making.

Conclusions

This study demonstrates that decoding techniques are a useful tool for investigating the properties of neural population codes, in particular for testing invariance properties (for a similar approach in V4/IT, see Rust and Dicarlo, 2010). As a next step, it will be important to compare neural decoding performance to psychophysical performance. This will allow judging how similar the behavioral readout is to the simple readout mechanism suggested by our results. Testing the contrast invariance of perceptual decisions might be an interesting starting point.

Notes

Supplemental material for this article is available at <http://bethgelab.org/datasets/v1gratings>. This webpage contains all neurophysiological data used in this article and example Matlab code. This material has not been peer reviewed.

References

- Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. *J Neurophysiol* 48:217–237.
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
- Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60:1142–1152.
- Benucci A, Ringach DL, Carandini M (2009) Coding of stimulus sequences by population responses in visual cortex. *Nat Neurosci* 12:1317–1324.
- Berens P, Ecker AS, Gerwinn S, Tolias AS, Bethge M (2011) Reassessing optimal neural population codes with neurometric functions. *Proc Natl Acad Sci U S A* 108:4423–4428.
- Bishop CM (2007) Pattern recognition and machine learning. Berlin: Springer.
- Busse L, Wade AR, Carandini M (2009) Representation of concurrent stimuli by population activity in visual cortex. *Neuron* 64:931–942.
- Chen Y, Geisler WS, Seidemann E (2006) Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat Neurosci* 9:1412–1420.
- Chen Y, Geisler WS, Seidemann E (2008) Optimal temporal decoding of neural population responses in a reaction-time visual detection task. *J Neurophysiol* 99:1366–1379.
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith

- MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378.
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819.
- De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544.
- Duda RO, Hart PE, Stork DG (2000) Pattern classification: pattern classification, Part 1. 2. Auflage. New York: Wiley.
- Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–587.
- Ecker AS, Berens P, Tolias AS, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. *J Neurosci* 31:14272–14283.
- Ferster D, Miller KD (2000) Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci* 23:441–471.
- Friedman J, Hastie T, Tibshirani R (2010) Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* 33:1–22.
- Ganmor E, Okun M, Lampl I (2009) The V1 population gains normalization. *Neuron* 64:778–780.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
- Graf AB, Kohn A, Jazayeri M, Movshon JA (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat Neurosci* 14:239–245.
- Greenberg DS, Houweling AR, Kerr JN (2008) Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nat Neurosci* 11:749–751.
- Hansen BJ, Dragoi V (2011) Correlated variability in laminar cortical circuits. *Soc Neurosci Abstr* 37:175.03.
- Hatsopoulos NG, Donoghue JP (2009) The science of neural interface systems. *Annu Rev Neurosci* 32:249–266.
- Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148:574–591.
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195:215–243.
- Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrsic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473:87–91.
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
- Kohn A, Zandvakili A, Smith MA (2009) Correlations and brain states: from electrophysiology to functional imaging. *Curr Opin Neurobiol* 19:434–438.
- Latham PE, Nirenberg S (2005) Synergy, redundancy, and independence in population codes, revisited. *J Neurosci* 25:5195–5206.
- London M, Häusser M (2005) Dendritic computation. *Annu Rev Neurosci* 28:503–532.
- Ma WJ (2010) Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Res* 50:2308–2319.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Ma WJ, Navalpakkam V, Beck JM, Berg R, Pouget A (2011) Behavior and neural basis of near-optimal visual search. *Nat Neurosci* 14:783–790.
- Macke JH, Opper M, Bethge M (2011) Common input explains higher-order correlations and entropy in a simple model of neural population activity. *Phys Rev Lett* 106:208102.
- Montani F, Kohn A, Smith MA, Schultz SR (2007) The role of correlations in direction and contrast coding in the primary visual cortex. *J Neurosci* 27:2338–2348.
- Müller JR, Metha AB, Krauskopf J, Lennie P (2001) Information conveyed by onset transients in responses of striate cortical neurons. *J Neurosci* 21:6978–6990.
- Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. *Science* 327:587–590.
- Ringach DL, Hawken MJ, Shapley R (1997) Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387:281–284.
- Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22:5639–5651.
- Rust NC, Dicarlo JJ (2010) Selectivity and tolerance (“invariance”) both increase as visual information propagates from cortical area V4 to IT. *J Neurosci* 30:12978–12995.
- Rust NC, Stocker AA (2010) Ambiguity and invariance: two fundamental challenges for visual processing. *Curr Opin Neurobiol* 20:382–388.
- Samonds JM, Bonds AB (2004) From another angle: differences in cortical coding between fine and coarse discrimination of orientation. *J Neurophysiol* 91:1193–1202.
- Schneidman E, Berry MJ 2nd, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6:484–507.
- Sclar G, Freeman RD (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp Brain Res* 46:457–461.
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28:12591–12603.
- Tolias AS, Ecker AS, Siapas AG, Hoenselaar A, Keliris GA, Logothetis NK (2007) Recording chronically from the same neurons in awake, behaving primates. *J Neurophysiol* 98:3780–3790.
- Vázquez P, Cano M, Acuña C (2000) Discrimination of line orientation in humans and monkeys. *J Neurophysiol* 83:2639–2648.