

Network Anisotropy Trumps Noise for Efficient Object Coding in Macaque Inferior Temporal Cortex

Yueh-peng Chen (陳嶽鵬),¹ Chia-pei Lin (林佳霏),^{1,2} Yu-chun Hsu (徐裕鈞),¹ and Chou P. Hung (洪洲伯)^{1,3}

¹National Yang-Ming University, Institute of Neuroscience and Brain Research Center, Taipei 112, Taiwan, ²National Taiwan University, Psychology Department, Taipei 106, Taiwan, and ³Georgetown University Medical Center, Department of Neuroscience, Washington, DC 20007

How neuronal ensembles compute information is actively studied in early visual cortex. Much less is known about how local ensembles function in inferior temporal (IT) cortex, the last stage of the ventral visual pathway that supports visual recognition. Previous reports suggested that nearby neurons carry information mostly independently, supporting efficient processing (Barlow, 1961). However, others postulate that noise covariation effects may depend on network anisotropy/homogeneity and on how the covariation relates to representation. Do slow trial-by-trial noise covariations increase or decrease IT's object coding capability, how does encoding capability relate to correlational structure (i.e., the spatial pattern of signal and noise redundancy/homogeneity across neurons), and does knowledge of correlational structure matter for decoding? We recorded simultaneously from ~80 spiking neurons in ~1 mm³ of macaque IT under light neurolept anesthesia. Noise correlations were stronger for neurons with correlated tuning, and noise covariations reduced object encoding capability, including generalization across object pose and illumination. Knowledge of noise covariations did not lead to better decoding performance. However, knowledge of anisotropy/homogeneity improved encoding and decoding efficiency by reducing the number of neurons needed to reach a given performance level. Such correlated neurons were found mostly in supragranular and infragranular layers, supporting theories that link recurrent circuitry to manifold representation. These results suggest that redundancy benefits manifold learning of complex high-dimensional information and that subsets of neurons may be more immune to noise covariation than others.

Key words: efficient coding; macaque; multielectrode array; noise correlation; object recognition; vision

Significance Statement

How noise affects neuronal population coding is poorly understood. By sampling densely from local populations supporting visual object recognition, we show that recurrent circuitry supports useful representations and that subsets of neurons may be more immune to noise covariation than others.

Introduction

How correlations affect neural coding efficiency is important for theories and models of brain processing. It is actively debated whether and how signal correlations (“ r_{signal} ,” similarity in stimulus tuning) and noise correlations (“ r_{SC} ,” trial-by-trial fluctuations in spike count in simultaneously recorded activity) affect coding efficiency. Previous reports suggested that nearby neu-

rons are mostly independent (i.e., low r_{SC}) (Gawne and Richmond, 1993; Zohary et al., 1994; Vinje and Gallant, 2000; Kohn and Smith, 2005; Mitchell et al., 2007), and it has been postulated that decorrelation (sparseness) supports efficient processing by maximizing the information encoded by a population of neurons (Barlow, 1961; Olshausen and Field, 1996; Bell and Sejnowski, 1997). However, it is conjectured that noise correlations may have more complex effects on population coding, depending on the specific relationship between signal and noise correlation and on population homogeneity and anisotropy (Abbott and Dayan, 1999; Wu et al., 2002; Averbeck et al., 2006; Renart et al., 2010; Cohen and Kohn, 2011; Ganmor et al., 2011; Berens et al., 2012; Shamir, 2014; Okun et al., 2015; Panas et al., 2015). By anisotropy, we mean the variation in functional excitability and functional interactions across a population, as described previously (Hung et al., 2014; Lin et al., 2014; Okun et al., 2015) and as modeled by Panas et al. (2015). In theory, both positive and negative associations between signal and noise correlation may in-

Received Nov. 6, 2014; revised May 22, 2015; accepted May 29, 2015.

Author contributions: C.P.H. designed research; Y.-p.C., C.-p.L., Y.-c.H., and C.P.H. performed research; Y.-p.C. and C.P.H. analyzed data; Y.-p.C. and C.P.H. wrote the paper.

This work was supported by Taiwan Ministry of Education Aim for the Top University Plan NSC-96-2811-B-010-501 and NSC-97-2321-B-010-007 and Georgetown University Medical Center. We thank Cheston Tan, Gabriel Kreiman, Thomas Serre, and Tomaso Poggio for stimuli and helpful discussions.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Chou P. Hung, Georgetown University Medical Center, Department of Neuroscience, 3970 Reservoir Road NW, NRB EP-04, Washington, DC 20007. E-mail: ch486@georgetown.edu.

DOI:10.1523/JNEUROSCI.4595-14.2015

Copyright © 2015 the authors 0270-6474/15/359889-11\$15.00/0

crease or decrease coding efficiency, with possibly different results for encoding versus decoding, depending on how the correlations affect redundancy and depending on the population's size, homogeneity, correlational structure, and how these relate to the representation. Addressing this debate requires dense sampling of many neurons per cortical column in response to naturalistic stimuli, to measure how coding efficiency depends on local correlational structure.

We previously showed that inferior temporal (IT) populations carry generalizable object information that can be read out by downstream neurons via pooling of weighted synaptic inputs, as measured by a linear classifier (Hung et al., 2005). We recently boosted the homogeneity and density of the sampled population, to record cells that have overlapping tuning, by inserting 64-contact multidepth arrays in neighboring cortical columns (~80 neurons within 1 mm³). We reported that, contrary to the current view that decorrelation supports efficient coding, correlated (similarly tuned) neurons have better object coding capability (Hung et al., 2014; Lin et al., 2014). Here, we reanalyzed these data on a trial-by-trial basis. We hypothesized that noise covariation reduces object encoding and decoding capability in larger and more homogeneous populations in IT.

Materials and Methods

Methods. All procedures adhered to the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the Institutional Animal Care and Use Committee of National Yang-Ming University. The methods and recordings were previously described (Lin et al., 2014). Briefly, we recorded spiking activity from right lateral surface of anterior IT (AP16) of 3 *Macaca cyclopis* monkeys (1 male and 2 females) under light neurolept anesthesia (0.9 μg/kg/h i.v. fentanyl, 70%/30% N₂O/O₂, 0.3%–0.5% isoflurane, 0.25 mg/kg i.m. droperidol) (Fujita et al., 1992; Wang et al., 1996; Brown et al., 2011; Sato et al., 2013) and muscle relaxation (rocuronium bromide).

We inserted dense arrays (8 shanks, 8 contacts per shank, spanning 1.4 × 1.4 mm at 0.2 × 0.2 mm spacing horizontally and in depth, A8x8–5mm–200–200–413, Neuronexus Technologies) across 5 recording sessions, one array insertion per session. Arrays 2, 4, and 5 were recorded across 3 sessions from the same monkey at locations ~3 mm apart (these were mislabeled in our previous report). Spikes (400–5000 Hz) were filtered (48 dB/octave) and continuously digitized at 24.4 kHz (RZ2, Tucker-Davis Technologies). Spikes were detected as voltage threshold crossings (“hash MUA”), or isolated (WaveClus_2.0) as single-unit activity (SUA) and then grouped as multiunit activity (MUA, ~1–2 SUAs each) (Lin et al., 2014). Most analyses were based on SUA, but results were similar for SUA, MUA, and hash MUA (see Figs. 4, 5, 6). Although spike sorting has been postulated to bias toward weaker r_{SC} , it avoids a possible bias of hash MUA toward more active and correlated neurons (Manning et al., 2009). We excluded units with <2000 spikes (<4 Hz) and rare units with slow (<0.1 Hz) rate fluctuations. We excluded multiple detection of the same unit across contacts via coincidence detection. Even-versus-odd trial consistency and consistency of tuning selectivity (“sparseness”) across stimulus sets were previously reported (Lin et al., 2014). All data will be uploaded to www.crcns.org.

We measured cortical depth by visually tracking individual contacts as they disappeared into the brain during insertion and, because of the small footprint of the shanks (15 × 33 μm), tracking individual units as they transitioned from the deepest to the most superficial contacts. Recording directly from the lateral surface avoids distortions from electrode bending, cortical compression, uneven sampling, and fine blood vessel damage. We estimate that the deviation of the array from vertical was <8 deg (<0.2 mm horizontal offset at the deepest site) because we could see and hear spikes at the beginning of array insertion on all 8 shanks (contacts 8, 16, ...64). Deviation in the orthogonal plane would have resulted in depth-specific correlation, whereas we observed correlation between the most superficial and deepest depths (Lin et al., 2014). Sectioning confirmed that the arrays were at a flat part of lateral cortex (i.e., depth was

not distorted by cortical curvature). Histological confirmation was impossible due to damage from later sessions. Layers were estimated from cortical depth, based on our previous results (cytochrome oxidase staining, current source density, and temporal frequency analysis) in V1 with the same arrays (Chu et al., 2014; their Fig. S1).

Visual stimuli. Objects ~10° wide were shown monocularly to the left eye, which was focused via contact lens upon a CRT monitor 57 cm away. Objects were positioned foveally via alignment of the optic disc. We used rapid serial visual presentation at 5 Hz (94 ms ON, 106 ms OFF), interleaved with the gray background. Stimuli were shown in pseudorandom order (10 repetitions, i.e., all objects in random order, followed by the same objects in a different order, etc.) and consisted of a block of 240 grayscale rendered 3D objects (no color/texture) at center pose and illumination, then a block to test generalization of 10 preferred objects across 25 variations in pose and illumination (see Fig. 3A). Stimuli belonged to a broad variety of categories, including but not limited to animals, faces, plants, foods, tools, vehicles, appliances, and furniture. The monkeys had never seen these images.

Classifier analysis. We tested classifier readout via a linear support vector machine classifier (MATLAB bioinformatics toolbox, The Math-Works). The effect of noise correlation was similar across a wide range of soft margins ($c = 1$ to 10^{-8}), and results are based on $c = 10^{-8}$ to avoid overfitting. For all analyses except the binned analyses, the classifier input was based on the z -normalized spike count between 100 and 300 ms after stimulus onset, where z -normalization was across stimuli for each unit.

For within-category generalization (see Fig. 1B), we trained 8 one-versus-all binary classifiers to choose among 8 categories via a winner-take-all strategy. The classifiers were trained on 8 objects per category (randomly chosen from 13 to 21 objects per category) and then tested on 5 other objects per category. Chance is 12.5% for 8 categories. The input matrix for classifier training was 640 trials (8 categories × 8 objects × 10 trials) × up to 306 units and was either unshuffled (simultaneously recorded) or trial-shuffled within each object and unit, to mimic noise-independent units. Classifier testing was based on 5 other objects per category, and the input matrix for the classifier was 400 trials (8 categories × 5 objects × 10 trials) × up to 306 units, also shuffled or unshuffled.

For object identification (see Fig. 2), we trained 240 binary classifiers on responses in 5 trials to 240 objects and tested their ability to identify these objects based on responses in the other 5 trials. Chance is 0.42% for 240 objects. The input matrix for classifier training was based on the binned spike count in the 100:300 ms window after stimulus onset (i.e. (240 objects × 5 trials) × 306 units, and the input for classifier testing was the same matrix for the remaining 5 trials. Figure 2 shows the performance of the best bin for each bin size (e.g., for 150:175 ms at 25 ms bin size).

Generalization across object pose and illumination was analyzed separately for arrays 2 and 4 (see Fig. 3). Object pose/illumination generalization was not tested for other arrays because their recording blocks for pose/illumination generalization contained slow (<0.1 Hz) firing rate fluctuations, likely from anesthetic accumulation in these later blocks. Such rate fluctuations were absent in arrays 2 and 4 and in the initial 240-object recording block in all arrays. The classifier was trained on random subsets of 4 variations in pose or illumination (all illuminations of each pose, or all poses of each illumination) and then tested on generalization to the remaining pose/illumination (all illuminations of one pose, or all poses of one illumination). Chance was 10% for choosing among 10 preferred objects per array. Thus, the matrix for classifier training was 2000 trials (4 variations × 5 subvariations × 10 objects × 10 trials) × N units and the matrix for classifier testing was 500 trials (1 variation × 5 subvariations × 10 objects × 10 trials) × N units. N was all SUAs in each array (75 units for array 2, 25 units for array 4).

For categorization (see Fig. 4), we trained 8 classifiers as before, except that stimulus repetitions were randomly divided into 5 “training” trials and 5 “test” trials. The input for classifier training was a matrix of 520 training trials (8 categories × 13 objects per category × 5 trials) × up to 306 SUA units and was either unshuffled or trial-shuffled within each object and unit. Classifier testing was based on the same matrix using the remaining 5 trials, also shuffled or unshuffled. For time course analysis,

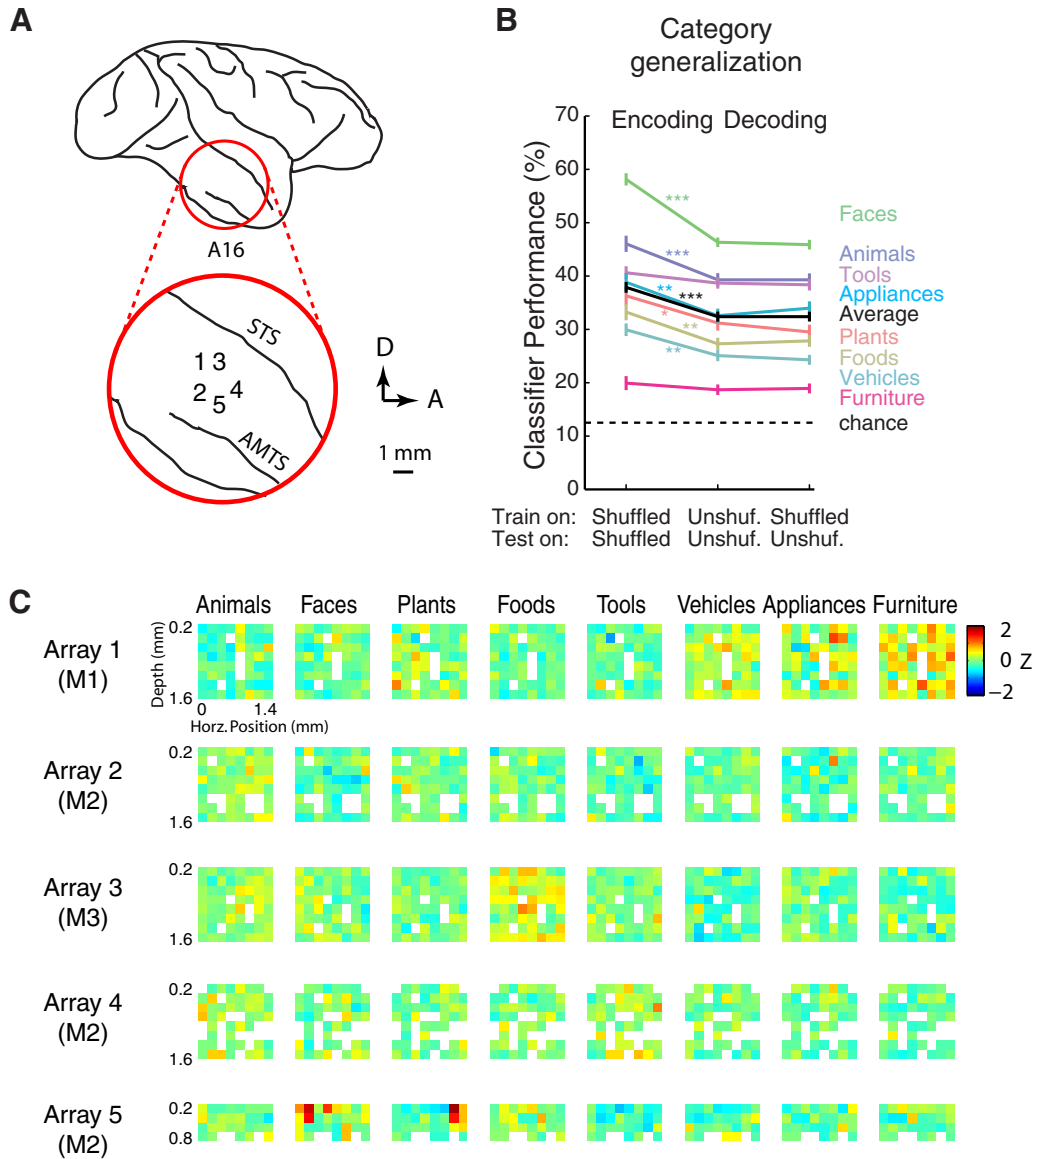


Figure 1. Noise covariation reduces object category generalization encoding in IT. **A**, Array locations 1–5, centered at A16 of right lateral inferior temporal cortex, recorded in separate sessions across 3 macaque monkeys under light neurolept anesthesia. Arrays had 64 contacts (8 shanks, 8 contacts/shank) spaced 0.2 mm apart, spanning 1.4 × 1.4 mm horizontally and in depth. STS, Superior temporal sulcus; AMTS, anterior medial temporal sulcus. **B**, Effect of noise correlation on classifier performance for within-category generalization. Linear classifiers were trained on 8 objects per category and tested on 5 other objects per category. Chance is 12.5% for 8 categories. For encoding, classifiers were trained and tested on unshuffled (actual) data versus trained and tested on trial-shuffled data, shown for individual categories (colors) and for the average across categories (black). For decoding, classifiers were trained and tested on unshuffled data, versus trained on shuffled data and tested on unshuffled data. Classifier input was spike count in the 100:300 ms period, pooled across 5 arrays and all trials. Data are mean ± SEM (50 permutations of objects and trials). **p* < 0.05 (two-tailed *t* test, uncorrected). ***p* < 0.01 (two-tailed *t* test, uncorrected). ****p* < 0.005 (two-tailed *t* test, uncorrected). **C**, Average category responses for the 5 arrays. Spike counts 100:300 ms after stimulus onset were averaged across 10 trials, z-normalized across all 240 objects for each contact, and then averaged within each category. z-scores are low because different objects in a category tend to activate different sets of contacts. White represents broken/inactive contacts.

we binned the spike counts and trained/tested the classifiers separately for each bin.

Signal and noise correlation analysis. Signal correlation (r_{signal}) and noise correlation (r_{SC}) were measured for pairs of hash MUA or pairs of SUAs recorded on separate channels, not between SUAs from the same channel, to avoid bias from spike collisions. Signal correlation was measured as the Pearson correlation of trial-averaged responses to different stimuli, based on the z-normalized spike count between 100 and 300 ms after stimulus onset, where z-normalization was across stimuli for each unit. Noise correlation was measured as the Pearson correlation of trial-by-trial spike count between 100 and 300 ms, after subtraction of mean response to that stimulus.

Network anisotropy (homogeneity) analysis. To directly assess the effect of network anisotropy, we sorted the units within each array by their

average r_{signal} with other units in the same array. We then defined Group 1 as units that had the strongest average r_{signal} in each array and defined Group 2 units as units that had median average r_{signal} (termed “choristers” and “soloists” in Hung et al., 2014; Lin et al., 2014), excluding training and test objects. We previously showed that neurons with higher average r_{signal} also tend to have spontaneous coincident firing (Lin et al., 2014; their Figs. 4 and 6; Tamura et al., 2014; their Fig. 10). Our measure of anisotropy is similar to another measure in which “choristers” and “soloists” were characterized by population coupling and synaptic coupling, which were partially linked to the strength of visual drive: compare Okun et al. (2015; their Fig. 3e) with Hung et al. (2014; their Fig. 2B). Our measure is also supported by a model of network stability in which anisotropy was characterized by a combination of functional excitability and functional interactions (Panas et al., 2015).

For Figure 7A, 2 units per array means the 2 Group 1 units with the strongest average r_{signal} and 2 Group 2 units with average r_{signal} closest to the 50th percentile in each array. For Figure 7B, Group 1 units were the 30% (solid lines) of the units that had the strongest average r_{signal} and Group 2 units were the 30% that had median average r_{signal} (35–65th percentile) in each array (i.e., 14 Group 1 and 14 Group 2 units for array 2 and 7 Group 1 and 7 Group 2 units for array 4). For generalization across pose and illumination, we reduced this to 20% (dashed lines, same Group 1 threshold as in Fig. 6A) to highlight the better efficiency of correlated neurons for smaller populations (i.e., 8 Group 1 and 8 Group 2 units for array 2 and 4 Group 1 and 4 Group 2 units for array 4), defined without training and test objects.

Results

Noise covariation reduces category generalization encoding and decoding

We measured spiking responses to 240 object stimuli across 5 sessions (one array insertion per session) in 3 monkeys (Fig. 1A) (Lin et al., 2014). Object category readout was based on 104 of the 240 objects, comprising 8 categories and 13 objects per category (categories with <13 objects were not tested).

A key requirement of cortical computations for object recognition is generalization across stimulus variations (e.g., variations within a category and in object pose and illumination) (Poggio and Bizzi, 2004; DiCarlo et al., 2012). To test the effect of noise covariation on category generalization, we trained one-versus-all binary classifiers on 8 of 13 objects per category and tested them on the 5 remaining “unseen” objects (chance is 12.5% for 8 categories).

Noise covariation significantly reduced, by ~5%–10%, within-category generalization performance for most categories (Fig. 1B; chance is 12.5% for 8 categories, $p < 0.005$ for faces and animals, $p < 0.01$ or 0.05 for other categories, two-tailed t test). Because noise covariation may have different effects on encoding versus decoding (Averbeck et al., 2006), we measured both types of effects. For encoding, we compared the performance when classifiers were trained and tested on unshuffled (simultaneously recorded) data, versus when they were trained and tested on trial-shuffled data (“shuffled,” to mimic independently sampled neurons). For decoding, we compared the performance when classifiers were trained and tested on unshuffled data, versus when they were trained on shuffled data and tested on unshuffled data. Noise covariation significantly reduced category generalization performance for encoding. Knowledge of noise covariation had no significant effect on decoding, suggesting that the correlations are orthogonal to the decision boundary (Eyherrabide and Samengo, 2013; their Fig. 7C).

The effect of noise covariation was uneven across categories. For example, “tools” had no significant effect, despite high overall performance, possibly due to a more distributed representation for that category. However, no category showed an increase in encoding performance from noise covariation. Such stimulus dependence of spike time correlations was previously reported for IT responses to face feature configuration (Hirabayashi and Miyashita, 2005) and is postulated to be an additional code in addition to rate coding (Reichert and Serre, 2014). We also note that stronger average responses to a category (e.g., to furniture for array 1) do not necessarily guarantee stronger generalization performance or stronger noise covariation effect, because response patterns vary across objects in a category.

Also, “faces” and “animals” had higher performance and stronger noise covariation effect, consistent with stronger tuning to these categories in some penetrations (Fig. 1C). The stronger noise covariation effect for “faces” and “animals” is consistent with previous reports of spatially clustered representation for these categories in IT (Wang et al., 1996; Freiwald and Tsao, 2010;

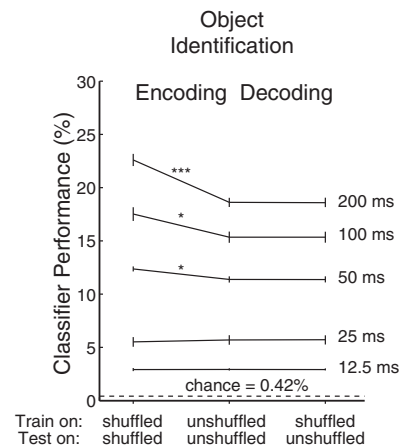


Figure 2. Noise covariation reduces object identification encoding. Effect of noise covariation on classifier performance for object identification, for encoding and decoding across different bin sizes within the 100:300 ms window. Classifier training was based on 5 trials, and classifier testing was based on 5 other trials. Performance is shown for the best bin of each bin size (e.g., the 150:175 ms bin for the 25 ms bin size). Chance is 0.42% for 240 objects. Data are mean \pm SEM (50 permutations). * $p < 0.05$ (two-tailed t test, uncorrected). *** $p < 0.005$ (two-tailed t test, uncorrected).

Ku et al., 2011; Sato et al., 2013). Array 5’s tuning and location are suggestive of a previously reported “AL” face patch, but we only sampled two other array locations in this monkey.

The effect of noise covariation on encoding and decoding was not due to collapsing across responses to different objects in each category (i.e., misinterpreting signal as noise) because we also observed this effect for object identification (of 240 objects; Fig. 2). The effect on encoding was significant across a range of bin sizes down to 50 ms. The absence of the effect at smaller bin sizes is likely due to less training data (only 1 object per class) for identification compared with categorization.

Noise covariation reduces pose and illumination generalization

Next, we measured how noise covariation affects generalization across changes in object pose and illumination. For each array, we selected 10 objects that strongly drove a few neurons and presented 25 variations of each object (5 poses \times 5 illuminations, -45° to 45° at 22.5° steps; Fig. 3A). We trained the classifier on responses to four poses or illuminations per object (all illuminations at each pose or all poses at each illumination), then tested readout of object identity at the remaining pose or illumination. In all cases, noise covariation reduced generalization performance for encoding by ~5%–10% (Fig. 3B, $p < 0.005$; Fig. 3C, $p < 0.05$), and knowledge of noise covariation had no significant effect on decoding performance.

Noise covariation effect is fast and increases with ensemble size

We tested how the effect of noise covariation on category encoding and decoding depends on ensemble size. Category readout was based on one-versus-all binary classifiers with cross-validation across trials, randomly assigned to 5 training trials and 5 test trials per object.

Consistent with theoretical predictions (Sompolinsky et al., 2001; Averbeck et al., 2006), the effect on encoding increased with population size (Fig. 4A), reaching 10%–15% relative reduction in classifier performance ($\Delta P/P_{\text{corrected}}$) at 32–64 randomly selected units per array (Fig. 4C). With fewer simultaneously re-

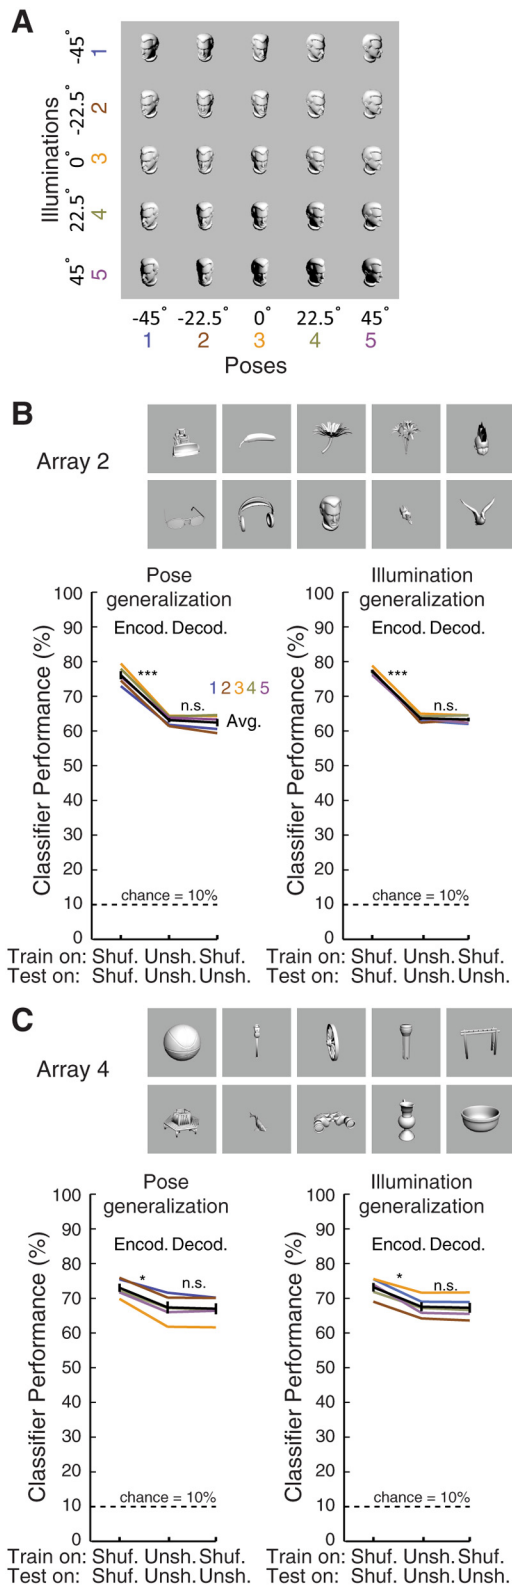


Figure 3. Noise covariation reduces pose and illumination generalization. **A**, Pose and illumination variations for one object. **B, C**, Classifier performance for generalization across pose and illumination (colors) for encoding and decoding. Based on arrays 2 and 4, recorded in separate sessions. For each array, we tested 25 variations of 10 preferred objects. Black represents average across poses/illuminations. Classifiers were trained on 4 poses or 4 illuminations (20 variations per object) and tested on object identification at the unseen pose or illumination. Data are mean \pm SD (50 trial permutations). Chance is 10% for 10 objects. * $p < 0.05$. *** $p < 0.005$. n.s., Not significant.

recorded units, although performance was above chance, noise covariation had negligible effect on encoding performance, consistent with Aggelopoulos et al. (2005) and Anderson et al. (2007). However, this was not simply due to flooring of the performance with fewer units because array 3 (orange) also showed an increased effect with population size, despite low overall performance. For decoding (Fig. 4B,C), the effect on $\Delta P/P_{corrected}$ was nonsignificant for all population sizes and reached up to 0.8% for 64 units per array. $\Delta P/P_{corrected}$ was slightly negative for 2–8 units per array, possibly because shuffling avoids errors from overtraining with few units. For both encoding and decoding, the increased effect with population size was similar for multiunit and single-unit activity, including unsorted hash, and it was consistent across the 5 arrays.

The time course of the noise covariation effect was fast for encoding (Fig. 4D) and was similar for categorization and identification (data not shown). The effect was similar for single-unit and multiunit activity and across different object categories and bin sizes down to 25 ms (0–1 spike per neuron). The effect of noise covariation on encoding was significant even from the earliest response (~ 100 ms) at 25 ms bin size, and its time course was similar to that of the classifier readout performance, supporting that its dynamics coincide with rapid feedforward (“core”) object recognition (DiCarlo et al., 2012). For decoding, time course analysis did not reveal a significant effect of noise covariation.

Signal correlation and noise correlation are strongly linked

Previous reports conjectured that noise correlation (r_{SC}) may increase or decrease coding efficiency, depending on its relationship to signal correlation (r_{signal}). Noise correlation has been predicted to increase information if signal and noise correlation are negatively related, and to decrease information if they are positively related. Consistent with this, signal and noise correlation were strongly positively linked in all 5 arrays (Fig. 5; Fisher-corrected $r = 0.38$ to 0.79 for SUA, 0.49 to 0.81 for hash MUA, $p < 0.005$ for each array). Mean R was similar for SUA (0.63) and for hash MUA (0.69), indicating that these results are not strongly biased by oversorting or by low spike count (Cohen and Kohn, 2011). Also consistent with previous reports, both signal and noise correlation were weak (mean $r_{signal} = 0.11$, mean $r_{SC} = 0.05$, $N = 10,645$ pairs) despite strong even-versus-odd trial tuning consistency (mean $r = 0.6$) (Lin et al., 2014). These results extend upon our previous report linking tuning correlation to spike time synchrony during spontaneous activity (Lin et al., 2014), to slow trial-by-trial noise correlations.

Signal and noise correlation effects depend on cortical depth

In primary visual cortex, signal and noise correlation are weaker in the granular layer than in supragranular and infragranular layers, and this has been tied to better orientation decoding efficiency for neurons in the granular layer (Hansen et al., 2012). In our IT data, signal and noise correlation also depended on cortical depth. Within each array, we sorted units by average r_{signal} with other units in the same array. Units in the top 20th percentile of average r_{signal} (Group 1 units, Fig. 6A, based on SUA) were disproportionately fewer at 1.0–1.2 mm depth (granular layer, 2 of 55 units). This depth dependence is specific to average r_{signal} (related to “population sparseness”) and is absent for tuning selectivity (Hung et al., 2014; their Fig. 3), a common measure of “sparseness” (Vinje and Gallant, 2000; Zoccolan et al., 2007; Willmore et al., 2011).

Noise correlation (r_{SC}) was weaker in the granular layer (mean $r_{SC} = 0.021$ for SUA, 0.035 for hash MUA) compared with supra-

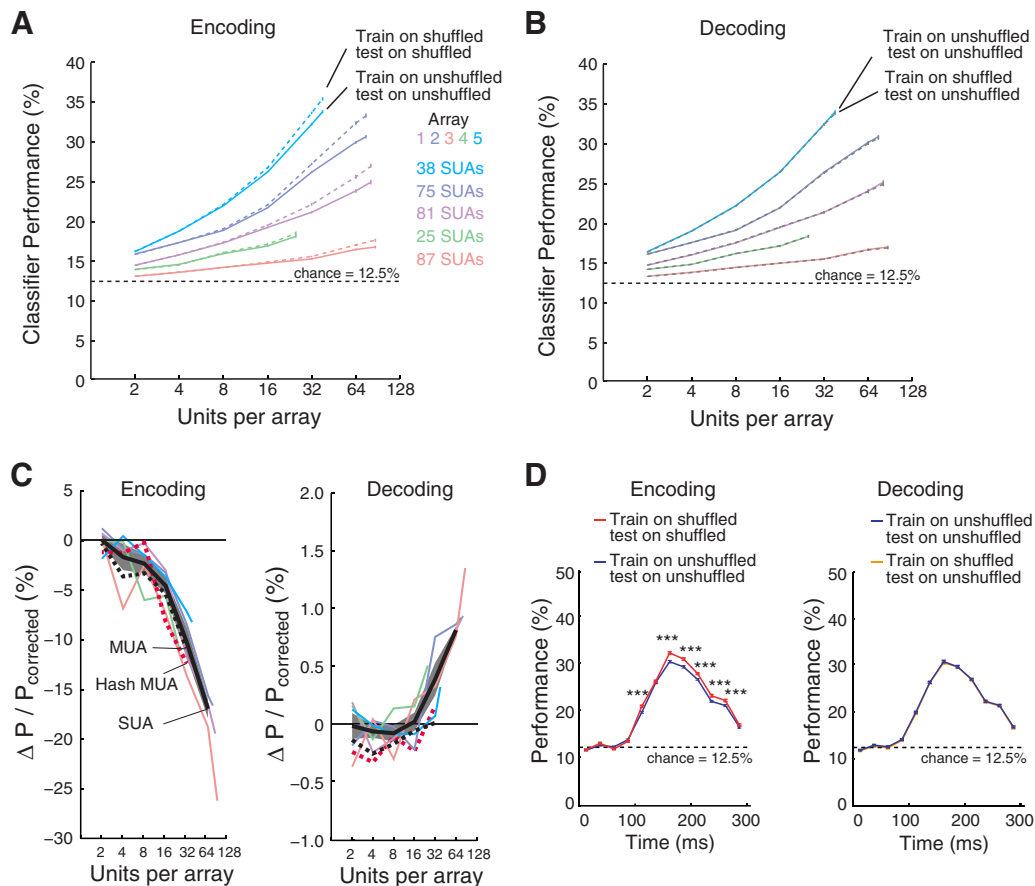


Figure 4. Dependence of noise covariation effect on ensemble size and latency. **A**, Effect of noise covariation on category encoding as a function of number of units in each of 5 arrays (colors). Classifiers were trained and tested on unshuffled data (solid lines) or trained and tested on shuffled data (dashed lines). Classifiers were trained on 5 trials and then tested on the remaining 5 trials, based on spike count in the 100:300 ms window. Data are mean \pm SEM (50 permutations of units and trials). Chance = 12.5% for 8 categories. **B**, Effect of noise correlation on category decoding. Classifiers were trained and tested on unshuffled data (colored solid lines) or trained on shuffled and tested on unshuffled data (gray dashed lines). **C**, Dependence of effect on population size for encoding (left) and decoding (right). Colors show results for individual arrays, based on SUA. $\Delta P/P_{corrected}$ is calculated as follows: $(P_U - P_S)/(P_U - P_C)$, where P_U , P_S , and P_C are performance on unshuffled training trials, shuffled training trials, and chance (12.5% for 8 categories), respectively. Mean $\Delta P/P_{corrected}$ is shown as thick black line for SUA, dashed black line for MUA (pooled from SUA), and dashed red line for hash MUA. Shaded region is SEM (50 permutations). **D**, Time course of object categorization performance for encoding (left) and decoding (right). Classifier input was the average spike rate for one 25 ms bin in the 100:300 ms window, pooled across 5 arrays. Data are mean \pm SEM (50 permutations, SUA). * $p < 0.05$; ** $p < 0.01$; *** $p < 0.005$ (two-tailed t test, uncorrected).

granular and infragranular layers (0.066 for SUA, 0.096 for hash MUA; Fig. 6B; $p < 0.001$ for all comparisons, unpaired t test). A recent report also found low noise correlation in anesthetized IT, based on tetrode recordings in the granular layer (mean $r_{SC} = 0.02$ for SUA) (Tamura et al., 2014). These values are lower than previously reported for SUA in awake V1, where mean r_{SC} was 0.04 in the granular layer versus 0.23–0.24 in supragranular and infragranular layers (Hansen et al., 2012; see also Kohn and Smith, 2005; Ecker et al., 2010; Cohen and Kohn, 2011). This decrease in noise correlation from V1 to IT is consistent with the weaker signal correlation and rarer spontaneous coincident spiking in IT compared with V1 (Chu et al., 2014; Lin et al., 2014).

The effect of noise covariation on classifier performance also depended on cortical depth. Categorization, within-category identification, and within-category generalization performances were lower in the granular layer than in supragranular and infragranular layers (Fig. 6C; all 12.5% chance, for 8 categories and 8 objects/category). The similar depth dependence across these tasks supports that the better performance for output layers was not task-specific. For encoding (red, blue, and black lines), noise correlation significantly reduced classifier performance in supragranular and infragranular layers ($p < 0.05$ to $p < 0.005$) and had

no significant effect in the granular layer. However, supragranular and infragranular layers still outperformed the granular layer for unshuffled data. For decoding (gray lines), knowledge of noise covariation had no significant effect on performance in any layer. The better performance of output layers in IT for object coding is surprising compared with the better performance of the input layer in V1 for orientation decoding (Hansen et al., 2012), and this difference is discussed below.

Noise covariation effect depends on network anisotropy

The dependence of the noise covariation effect on ensemble size and on cortical depth indicates that it should also depend on network anisotropy (variations in homogeneity, measured as average r_{signal}). We recently reported that, in IT, within-category generalization performance is better for correlated neurons (units with the highest average r_{signal} and spontaneous coincident spiking), consistent with a link between correlated activity and manifold theories of representation (DiCarlo et al., 2012; Hung et al., 2014; Lin et al., 2014). Here, we examine how noise covariation's effect on encoding and decoding depends on network anisotropy.

We compared the effect of noise covariation on two subsets of units. Group 1 was units with the highest average r_{signal} of each

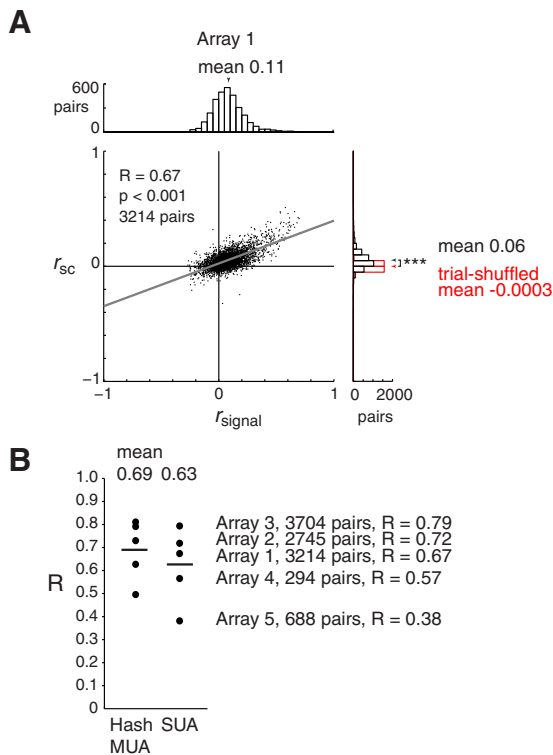


Figure 5. Signal correlation and noise correlation are strongly linked. **A**, Signal correlation (r_{signal}) versus noise correlation (r_{sc}) of neuronal pairs in array 1, based on 81 SUAs (3214 pairs, recorded from separate contacts). r_{signal} and r_{sc} are measured as Pearson correlation of trial-averaged stimulus responses and trial-by-trial response deviations, respectively, based on spike count from 100 to 300 ms after stimulus onset. Stimuli were 240 objects shown for 10 repetitions in pseudorandom order. Gray line indicates model II regression of r_{signal} versus r_{sc} . Red marginal represents r_{sc} after trial shuffling. $***p < 0.005$. **B**, Correlation (R) between r_{signal} and r_{sc} across 5 arrays. Horizontal lines are mean across 5 arrays for hash MUA ($r = 0.69$) and for SUA ($r = 0.63$). $p < 0.005$ for each array.

array, and Group 2 was units with average r_{signal} closest to the median of each array (Fig. 7A; see Materials and Methods). This designation of neurons based on functional interactions directly addresses the issue of network anisotropy raised in our previous reports and in recent reports (Okun et al., 2015; Panas et al., 2015). Object categorization performance was better for Group 1 than for Group 2, and noise covariation had weak but significant effects on classifier performance for encoding (colors) and no significant effect for decoding (gray). The stronger effect for Group 1 is consistent with theoretical predictions for homogeneous populations, and the effect required ensembles of at least 8–16 units to be detected, consistent with Figure 4.

To improve sensitivity, we pooled responses across arrays and included all units in the top 30th percentile as Group 1 and units in the median 30th percentile as Group 2 (Fig. 7B, red and black solid lines, 7–14 units per group per array). The noise covariation effect on encoding was stronger for Group 1 than for Group 2 for categorization, within-category generalization, and pose and illumination generalization. This differential effect was not simply due to the higher performance of Group 1 units than Group 2 units, particularly for smaller ensemble sizes (e.g., 20th percentile for pose and illumination generalization, dashed lines). For pose and illumination generalization at 30% ensemble size (solid lines), although performance was similar for both groups, only the Group 1 noise covariation effect was significant. For decoding, knowledge of noise covariation had no significant effect for both Group 1 and Group 2. Overall, these results show that

Group 1 units tend to be more efficient than Group 2 units (better performance for smaller ensembles) despite the cost of noise covariation.

Simulation of noise correlation effects on classifier performance

A possible concern is whether both positive and negative effects of noise covariation are detectable with linear classifiers, or whether higher-order kernels are necessary (Reichert and Serre, 2014). We simulated the effect of injecting artificial noise correlation to trial-averaged responses (Fig. 8). Simulated negative r_{sc} slightly improved both encoding and decoding performance, and the results for actual (unshuffled) data were approximated by a simulated r_{sc} of 0.2 (Averbeck et al., 2006; compare their Figs. 2 and 4). This approximation to r_{sc} of 0.2, instead of the much weaker average r_{sc} of 0.05 (Fig. 5) (Hung et al., 2014; their Fig. 2), is consistent with the effect of anisotropy. It suggests that encoding and decoding are mainly supported by Group 1 units in output layers, which have higher average r_{sc} (Figs. 5, 6B). However, we note that our results do not preclude the possibility that the brain uses higher-order kernels.

Are there any conditions in which increases in noise covariation might actually improve classifier performance for encoding? Although our sample was dominated by positive r_{sc} , we also observed subpopulations with negative r_{sc} . Neuronal pairs that were selected for $r_{sc} < -0.05$ on even trials (mean $r_{sc} = -0.082$, $N = 434$ of 10,645 pairs) also had negative r_{sc} on odd trials (mean $r_{sc} = -0.040$, $p < 10^{-9}$ for each of 5 arrays, one-sample t test). A plausible biological basis is that inhibition could produce negative r_{sc} between nearby neurons (Tamura et al., 2014; their Fig. 6) or possibly neighboring areas (consistent with interareal balancing) (Ramsden et al., 2001), which could be useful for denoising the effects of traveling waves (Xu et al., 2007). This potential benefit of negative r_{sc} should be explored in future studies.

Discussion

Our results provide a rare glimpse into the effect of noise covariation on local spiking ensembles at the end of the ventral visual pathway, and they suggest possible computational strategies for networks of neurons in higher cortex. As predicted, noise covariation reduced object encoding in IT, including generalization, and this effect increased with ensemble size and with population homogeneity, tied to a positive link between signal and noise correlation. Consistent with previous reports in V1, the effect of noise correlation depended on cortical depth and was weakest in the granular layer. However, performance was higher in output layers of IT, supporting that manifold representation is tied to recurrent circuitry despite stronger noise correlation. This dependence of coding efficiency on homogeneity and cortical depth was captured by network anisotropy, which was directly addressed by our designation of Group 1 and Group 2 units. Finally, our decoding results suggest that it may be possible to extract all the information in a population without knowledge of noise correlations. Overall, these data support that even weak noise correlations reduce coding efficiency but that, contrary to previous reports in the early visual pathway, redundancy is beneficial for manifold representation of complex high-dimensional object information, and subsets of neurons may be more immune to noise covariation than others.

The generality of these results is based on the assumption that neurons downstream of IT read out the object code via linear pooling of weighted synaptic inputs. Different results might be obtained by decoders that know how to exploit these correlations

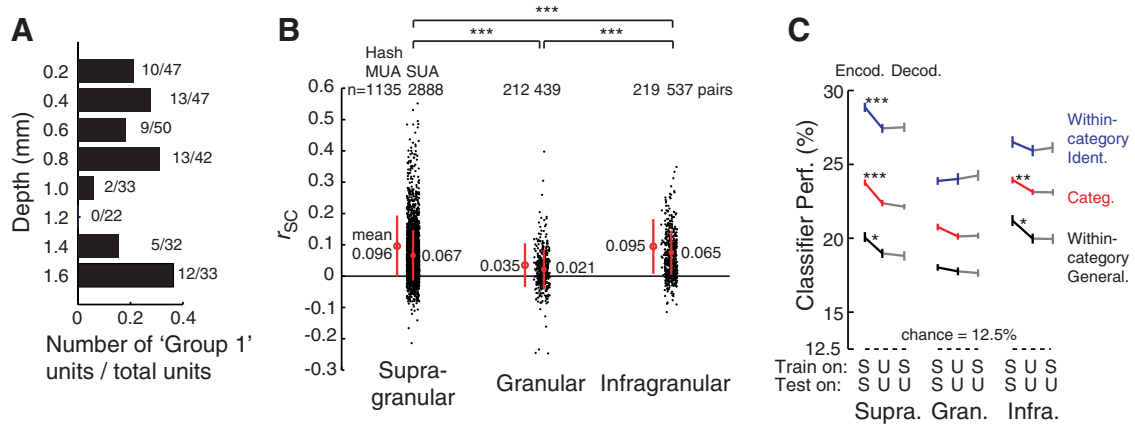


Figure 6. Cortical depth dependency of signal correlation, noise correlation, and coding efficiency. **A**, Cortical depth versus proportion of Group 1 units, defined as the top 20th percentile of units in each array sorted by average r_{signal} with other units in the same array, based on SUA. **B**, Noise correlation (r_{SC}) at supra-granular (0.2–0.8 mm), granular (1.0–1.2 mm), and infragranular (1.4–1.6 mm) depths for arrays 1–3, based on hash MUA and SUA. Layers were estimated from cortical depth. Red lines indicate mean \pm SD. Arrays 4 and 5 were excluded because they sampled mainly supra-granular layers. * $p < 0.05$. ** $p < 0.01$. *** $p < 0.005$. **C**, Effect of noise correlation on encoding (blue, red, and black) and decoding (gray) at supra-granular, granular, and infragranular depths for within-category identification, categorization, and within-category generalization, for arrays 1–3. Based on 8 units with highest average r_{signal} per depth group per array, defined without training/test objects. S, Shuffled; U, unshuffled. Data are mean \pm SEM (50 permutations of stimuli and trials). Chance is 12.5% for all tests, for 8 categories and 8 objects/category.

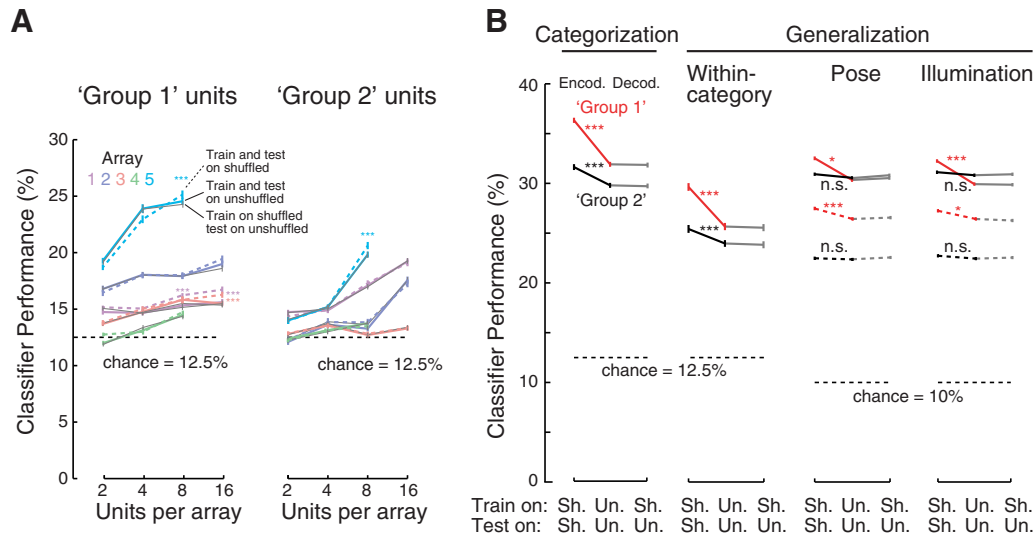


Figure 7. Noise covariation effect on coding efficiency depends on network anisotropy. **A**, Effect of noise covariation on category encoding (colors) and decoding (gray) for Group 1 versus Group 2 units in each array. Group 1 is units with the highest average r_{signal} with other units in the same array, and Group 2 is units with average r_{signal} closest to the median of each array, defined without training/test objects. Classifier training/testing is same as in Figure 4, cross-validated across trials. For both panels: * $p < 0.05$; *** $p < 0.005$. Data are mean \pm SEM (50 permutations). n.s., Not significant. **B**, Effect of noise covariation on encoding (red, black) and decoding (gray) for categorization, within-category generalization, and pose and illumination generalization for Group 1 and Group 2 units pooled across arrays (compare with Figs. 1B, 3B, C). Solid lines are Group 1 and Group 2 units defined as the 30% highest (red) and 30% median (black, 35th–65th percentile) average r_{signal} of each array, defined without training/test objects. Dashed lines are Group 1 and Group 2 units defined as the 20% highest (red) and 20% median (black) average r_{signal} of each array. Data are mean \pm SEM (50 permutations of stimuli and trials).

(e.g., via higher-order kernels) (Reichert and Serre, 2014). This possibility is supported by the shift from negative $\Delta I_{diag}/I$ for small ensembles to positive $\Delta I_{diag}/I$ for larger ensembles (Figs. 4C, 8B) and by differences in the noise correlation effect across categories (Fig. 1B). These conclusions also assume that IT population dynamics are not substantially altered under light neurolept anesthesia compared with awake behaving animals. Because of the difficulty of access, anesthesia and muscle relaxation were necessary to study noise covariation in IT, and they were also used in a previous study of noise correlation in V1 (Kohn and Smith, 2005).

We think that anesthesia did not have a strong effect in our data for four reasons. First, our fentanyl concentration was 10 \times lower than in reports that did not find an effect on neu-

ronal dynamics (Loughnan et al., 1987; Constantinople and Bruno, 2011) and was 100 \times lower than in a recent report that found an effect at frequencies slower than 2 Hz (Ecker et al., 2014). Second, we estimate that our noise correlations are faster than 3.3 Hz and that these results are not due to slow up/down fluctuations because injection of slower artificial noise correlations (via spike jittering) (Smith and Kohn, 2008; their Fig. 2) resulted in effects on classifier performance that exceeded that of actual noise correlations. We did observe such slow up/down fluctuations in some of our later recording blocks, possibly due to anesthetic accumulation, and we excluded those blocks from analysis. Third, a report from the same group found that noise correlations are more similar in active awake and anesthetized animals than during quiet

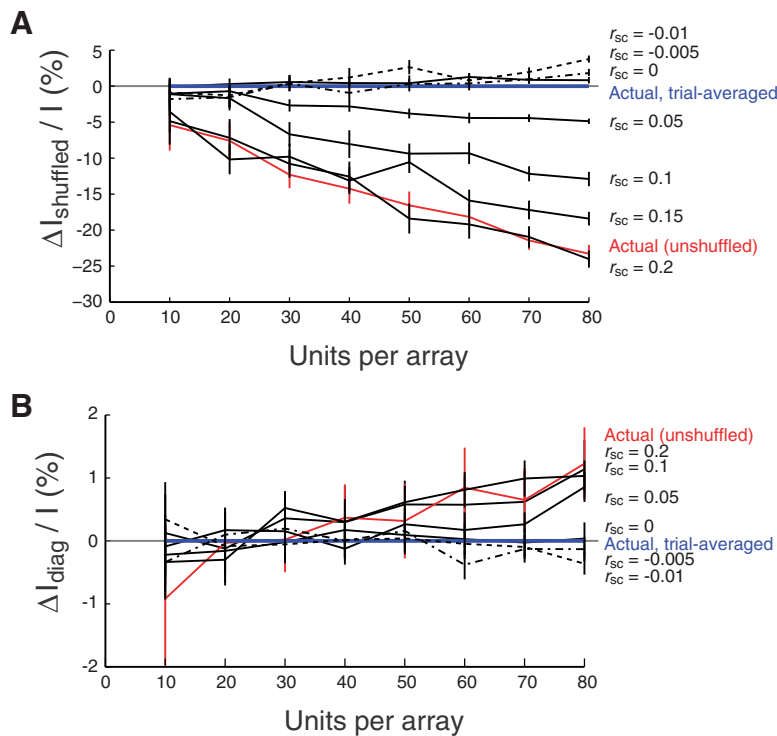


Figure 8. Effect of simulated noise correlation on category encoding and decoding. **A**, Effect of simulated noise correlation on category encoding. $\Delta I_{\text{shuffled}}/I$ versus population size for simulated noise correlation r_{sc} of -0.01 to 0.2 (black). I , Information; I_{shuffled} , trial-shuffled information; $\Delta I_{\text{shuffled}}/I = I_{\text{shuffled}}/I - I$. Noise correlation was simulated by injecting noise to trial-averaged data (blue). Results for unshuffled data (actual, red) approximate those of simulated r_{sc} of 0.2 . $\Delta I_{\text{shuffled}}/I$ is positive for negative r_{sc} . Compare with Averbeck et al. (2006; their Fig. 2). **B**, $\Delta I_{\text{diag}}/I$ versus population size. I_{diag} , Information that would be extracted by a decoder trained on trial-shuffled data but tested on actual correlated (unshuffled) data; $\Delta I_{\text{diag}}/I = I_{\text{diag}}/I - I$. Compare with Averbeck et al. (2006; their Fig. 4).

wakefulness (Froudarakis et al., 2014). Fourth, the noise correlations we measured were very low (mean $r_{\text{SC}} = 0.021$ for SUA, 0.035 for hash MUA in the granular layer), lower than in most reports of recordings in awake V1 (Kohn and Smith, 2005; Ecker et al., 2010; Hansen et al., 2012) and awake IT (Gawne and Richmond, 1993). The weakness of the noise correlation in the granular layer here and in a recent report also in anesthetized IT (Tamura et al., 2014) shows that these results cannot be due to strong thalamocortical fluctuations under anesthesia. In both awake V1 and in our lightly anesthetized IT results, noise correlations are tied to recurrent activity in supragranular and infragranular layers, consistent with an intermediate level of synchronization that is optimal for information flow between sparsely connected areas, rather than with a state of global synchronization that is exhibited during relaxed and sleeping states (Mark and Tsodyks, 2012). A strength of our study is that our use of muscle relaxation rules out correlations due to top-down attention and eye movements (Rajkai et al., 2008; Ito et al., 2011). However, a goal of future studies should be to understand how correlations impact actual behavior.

Why was coding efficiency better for correlated neurons despite stronger noise correlation?

The better coding efficiency of correlated neurons (Group 1) is surprising because sparseness is thought to support efficient coding, and this has been tied to weaker signal and noise correlation in layer 4 of V1 (Cohen and Kohn, 2011; Hansen et al., 2012). Although we also observed weaker signal and noise correlation

in layer 4 (Fig. 6), beyond the already weak correlation in IT (Gawne and Richmond, 1993; Sato et al., 2009; Lin et al., 2014; Tamura et al., 2014) compared with V1 (Chu et al., 2014), our finding of more efficient coding by correlated neurons, mostly in output layers, is counterintuitive.

Several factors may have contributed to this difference. Previous studies of coding efficiency were based on orientation discrimination by a few neurons per penetration in V1 or MT, whereas we measured generalization across naturalistic objects by a larger, denser, and more homogeneous population in IT. We suggest the hypothesis that the higher stimulus dimensionality, representational complexity, and sampling density in our study highlight the benefit of manifold representation and redundancy for noisy spiking populations. The better performance of Group 1 units is unrelated to tuning width (i.e., average r_{signal} and tuning selectivity, commonly reported as “sparseness”, are unrelated) (Vinje and Gallant, 2000; Zoccolan et al., 2007). The surprisingly better performance of Group 1 is consistent with recent reports that tolerance (Rust and Dicarlo, 2010), but not sparseness (Willmore et al., 2011), increases along the ventral visual pathway. The effect cannot be due to feedback from attention (Takeuchi et al., 2011; Zhang et al., 2011; Bansal et al., 2014; Scholl et al., 2014) because its rapid time course (Fig. 4D) matched that of feedforward recognition.

Previous reports implicated statistical learning of the natural environment by spiking neurons (Luczak et al., 2009; Rust and Dicarlo, 2010; Berkes et al., 2011; Zylberberg and DeWeese, 2013), and it has been postulated that IT supports invariant recognition by encoding a manifold representation of complex features (DiCarlo et al., 2012). Our results support this by showing that a low-dimensional correlational structure, which we describe as a “pipe cleaner” model (Hung et al., 2014; Lin et al., 2014), is more efficient for coding naturalistic shapes in IT (Baldassi et al., 2013). In our model, uncorrelated units (the bristles, Group 2) in the input layer, possibly decorrelated and orthogonalized to the local manifold (Bengio, 2009) by balanced excitation and local inhibition (Renart et al., 2010; Graupner and Reyes, 2013), act as tensors to fine-tune the manifold representation in output layers (the spine, Group 1). This is consistent with reports that that columnar-scale organization supports efficient coding (Tsodyks et al., 1999; Tanaka, 2003; Chen et al., 2006) and behavior (Nienborg and Cumming, 2014). This is also consistent with a recent report that continuous remodeling of neural networks depends on network anisotropy, with faster remodeling at the functional periphery (Group 2) and slower remodeling at the core (Group 1) (Panas et al., 2015). Overall, these results imply that, in dense populations that encode complex, high-dimensional information, coding is efficiently concentrated in a low-dimensional correlational structure, namely, synchronized neurons in output layers.

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101. [CrossRef Medline](#)
- Aggelopoulos NC, Franco L, Rolls ET (2005) Object perception in natural scenes: encoding by inferior temporal cortex simultaneously recorded neurons. *J Neurophysiol* 93:1342–1357. [CrossRef Medline](#)
- Anderson B, Sanderson MI, Sheinberg DL (2007) Joint decoding of visual stimuli by IT neurons' spike counts is not improved by simultaneous recording. *Exp Brain Res* 176:1–11. [CrossRef Medline](#)
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366. [CrossRef Medline](#)
- Baldassi C, Alemi-Neissi A, Pagan M, Dicarlo JJ, Zecchina R, Zoccolan D (2013) Shape similarity, better than semantic membership, accounts for the structure of visual object representations in a population of monkey inferotemporal neurons. *PLoS Comput Biol* 9:e1003167. [CrossRef Medline](#)
- Bansal AK, Madhavan R, Agam Y, Golby A, Madsen JR, Kreiman G (2014) Neural dynamics underlying target detection in the human brain. *J Neurosci* 34:3042–3055. [CrossRef Medline](#)
- Barlow HB (1961) Possible principles underlying the transformation of sensory messages. In: *Sensory communication* (Rosenblith W, ed), pp 217–234. Cambridge, MA: Massachusetts Institute of Technology.
- Bell AJ, Sejnowski TJ (1997) The “independent components” of natural scenes are edge filters. *Vision Res* 37:3327–3338. [CrossRef Medline](#)
- Bengio Y (2009) Learning deep architectures for AI. *Found Trends Machine Learning* 2:1–127. [CrossRef](#)
- Berens P, Ecker AS, Cotton RJ, Ma WJ, Bethge M, Tolias AS (2012) A fast and simple population code for orientation in primate V1. *J Neurosci* 32:10618–10626. [CrossRef Medline](#)
- Berkes P, Orbán G, Lengyel M, Fiser J (2011) Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331:83–87. [CrossRef Medline](#)
- Brown EN, Purdon PL, Van Dort CJ (2011) General anesthesia and altered states of arousal: a systems neuroscience analysis. *Annu Rev Neurosci* 34:601–628. [CrossRef Medline](#)
- Chen Y, Geisler WS, Seidemann E (2006) Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat Neurosci* 9:1412–1420. [CrossRef Medline](#)
- Chu CC, Chien PF, Hung CP (2014) Tuning dissimilarity explains short distance decline of spontaneous spike correlation in macaque V1. *Vision Res* 96:113–132. [CrossRef Medline](#)
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819. [CrossRef Medline](#)
- Constantinople CM, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical networks. *Neuron* 69:1061–1068. [CrossRef Medline](#)
- DiCarlo JJ, Zoccolan D, Rust NC (2012) How does the brain solve visual object recognition? *Neuron* 73:415–434. [CrossRef Medline](#)
- Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–587. [CrossRef Medline](#)
- Ecker AS, Berens P, Cotton RJ, Subramaniyan M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82:235–248. [CrossRef Medline](#)
- Eyherabide HG, Samengo I (2013) When and why noise correlations are important in neural decoding. *J Neurosci* 33:17921–17936. [CrossRef Medline](#)
- Freiwald WA, Tsao DY (2010) Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330:845–851. [CrossRef Medline](#)
- Froudarakis E, Berens P, Ecker AS, Cotton RJ, Sinz FH, Yatsenko D, Saggau P, Bethge M, Tolias AS (2014) Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat Neurosci* 17:851–857. [CrossRef Medline](#)
- Fujita I, Tanaka K, Ito M, Cheng K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360:343–346. [CrossRef Medline](#)
- Ganmor E, Segev R, Schneidman E (2011) Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc Natl Acad Sci U S A* 108:9679–9684. [CrossRef Medline](#)
- Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13:2758–2771. [Medline](#)
- Graupner M, Reyes AD (2013) Synaptic input correlations leading to membrane potential decorrelation of spontaneous activity in cortex. *J Neurosci* 33:15075–15085. [CrossRef Medline](#)
- Hansen BJ, Chelaru MI, Dragoi V (2012) Correlated variability in laminar cortical circuits. *Neuron* 76:590–602. [CrossRef Medline](#)
- Hirabayashi T, Miyashita Y (2005) Dynamically modulated spike correlation in monkey inferior temporal cortex depending on the feature configuration within a whole object. *J Neurosci* 25:10299–10307. [CrossRef Medline](#)
- Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. *Science* 310:863–866. [CrossRef Medline](#)
- Hung CP, Cui D, Chen YP, Lin CP, Levine MR (2014) Correlated activity supports efficient cortical processing. *Front Comput Neurosci* 8:171. [CrossRef Medline](#)
- Ito J, Maldonado P, Singer W, Grün S (2011) Saccade-related modulations of neuronal excitability support synchrony of visually elicited spikes. *Cereb Cortex* 21:2482–2497. [CrossRef Medline](#)
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673. [CrossRef Medline](#)
- Ku SP, Tolias AS, Logothetis NK, Goense J (2011) fMRI of the face-processing network in the ventral temporal lobe of awake and anesthetized macaques. *Neuron* 70:352–362. [CrossRef Medline](#)
- Lin CP, Chen YP, Hung CP (2014) Tuning and spontaneous spike time synchrony share a common structure in macaque inferior temporal cortex. *J Neurophysiol* 112:856–869. [CrossRef Medline](#)
- Loughnan BL, Sebel PS, Thomas D, Rutherford CF, Rogers H (1987) Evoked potentials following diazepam or fentanyl. *Anaesthesia* 42:195–198. [CrossRef Medline](#)
- Luczak A, Barthó P, Harris KD (2009) Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62:413–425. [CrossRef Medline](#)
- Manning JR, Jacobs J, Fried I, Kahana MJ (2009) Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *J Neurosci* 29:13613–13620. [CrossRef Medline](#)
- Mark S, Tsodyks M (2012) Population spikes in cortical networks during different functional states. *Front Comput Neurosci* 6:43. [CrossRef Medline](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55:131–141. [CrossRef Medline](#)
- Nienborg H, Cumming BG (2014) Decision-related activity in sensory neurons may depend on the columnar architecture of cerebral cortex. *J Neurosci* 34:3579–3585. [CrossRef Medline](#)
- Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD (2015) Diverse coupling of neurons to populations in sensory cortex. *Nature* 521:511–515. [CrossRef Medline](#)
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–609. [CrossRef Medline](#)
- Panas D, Amin H, Maccione A, Muthmann O, van Rossum M, Berdondini L, Hennig MH (2015) Sloppiness in spontaneously active neuronal networks. *J Neurosci* 35:8480–8492. [CrossRef Medline](#)
- Poggio T, Bizzi E (2004) Generalization in vision and motor control. *Nature* 431:768–774. [CrossRef Medline](#)
- Rajkai C, Lakatos P, Chen CM, Pincze Z, Karmos G, Schroeder CE (2008) Transient cortical excitation at the onset of visual fixation. *Cereb Cortex* 18:200–209. [CrossRef Medline](#)
- Ramsden BM, Hung CP, Roe AW (2001) Real and illusory contour processing in area V1 of the primate: a cortical balancing act. *Cereb Cortex* 11:648–665. [CrossRef Medline](#)
- Reichert DP, Serre T (2014) Neuronal synchrony in complex-valued deep networks. In: [arXiv:1312.6115v5](http://arxiv.org/abs/1312.6115). <http://arxiv.org/abs/1312.6115>.
- 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. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Sato T, Uchida G, Tanifuji M (2009) Cortical columnar organization is reconsidered in inferior temporal cortex. *Cereb Cortex* 19:1870–1888. [CrossRef Medline](#)
- Sato T, Uchida G, Lescroart MD, Kitazono J, Okada M, Tanifuji M (2013) Object representation in inferior temporal cortex is organized hierarchically in a mosaic-like structure. *J Neurosci* 33:16642–16656. [CrossRef Medline](#)
- Scholl CA, Jiang X, Martin JG, Riesenhuber M (2014) Time course of shape and category selectivity revealed by EEG rapid adaptation. *J Cogn Neurosci* 26:408–421. [CrossRef Medline](#)
- Shamir M (2014) Emerging principles of population coding: in search for the neural code. *Curr Opin Neurobiol* 25:140–148. [CrossRef Medline](#)
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28:12591–12603. [CrossRef Medline](#)
- Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. *Phys Rev E Stat Nonlin Soft Matter Phys* 64:051904. [CrossRef Medline](#)
- Takeuchi D, Hirabayashi T, Tamura K, Miyashita Y (2011) Reversal of interlaminar signal between sensory and memory processing in monkey temporal cortex. *Science* 331:1443–1447. [CrossRef Medline](#)
- Tamura H, Mori Y, Kaneko H (2014) Organization of local horizontal functional interactions between neurons in the inferior temporal cortex of macaque monkeys. *J Neurophysiol* 111:2589–2602. [CrossRef Medline](#)
- Tanaka K (2003) Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. *Cereb Cortex* 13:90–99. [CrossRef Medline](#)
- Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946. [CrossRef Medline](#)
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276. [CrossRef Medline](#)
- Wang G, Tanaka K, Tanifuji M (1996) Optical imaging of functional organization in the monkey inferotemporal cortex. *Science* 272:1665–1668. [CrossRef Medline](#)
- Willmore BD, Mazer JA, Gallant JL (2011) Sparse coding in striate and extrastriate visual cortex. *J Neurophysiol* 105:2907–2919. [CrossRef Medline](#)
- Wu S, Amari S, Nakahara H (2002) Population coding and decoding in a neural field: a computational study. *Neural Comput* 14:999–1026. [CrossRef Medline](#)
- Xu W, Huang X, Takagaki K, Wu JY (2007) Compression and reflection of visually evoked cortical waves. *Neuron* 55:119–129. [CrossRef Medline](#)
- Zhang Y, Meyers EM, Bichot NP, Serre T, Poggio TA, Desimone R (2011) Object decoding with attention in inferior temporal cortex. *Proc Natl Acad Sci U S A* 108:8850–8855. [CrossRef Medline](#)
- Zoccolan D, Kouh M, Poggio T, DiCarlo JJ (2007) Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. *J Neurosci* 27:12292–12307. [CrossRef Medline](#)
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143. [CrossRef Medline](#)
- Zylberberg J, DeWeese MR (2013) Sparse coding models can exhibit decreasing sparseness while learning sparse codes for natural images. *PLoS Comput Biol* 9:e1003182. [CrossRef Medline](#)