



Low rank mechanisms underlying flexible visual representations

Douglas A. Ruff^{a,1}, Cheng Xue^{a,1}, Lily E. Kramer^a, Faisal Baqai^b, and Marlene R. Cohen^{a,b,2}

^aDepartment of Neuroscience, University of Pittsburgh, Pittsburgh, PA 15260; and ^bProgram in Neural Computation, Carnegie Mellon University, Pittsburgh, PA 15260

Edited by Wilson S. Geisler, The University of Texas at Austin, Austin, TX, and approved September 10, 2020 (received for review April 9, 2020)

Neuronal population responses to sensory stimuli are remarkably flexible. The responses of neurons in visual cortex have heterogeneous dependence on stimulus properties (e.g., contrast), processes that affect all stages of visual processing (e.g., adaptation), and cognitive processes (e.g., attention or task switching). Understanding whether these processes affect similar neuronal populations and whether they have similar effects on entire populations can provide insight into whether they utilize analogous mechanisms. In particular, it has recently been demonstrated that attention has low rank effects on the covariability of populations of visual neurons, which impacts perception and strongly constrains mechanistic models. We hypothesized that measuring changes in population covariability associated with other sensory and cognitive processes could clarify whether they utilize similar mechanisms or computations. Our experimental design included measurements in multiple visual areas using four distinct sensory and cognitive processes. We found that contrast, adaptation, attention, and task switching affect the variability of responses of populations of neurons in primate visual cortex in a similarly low rank way. These results suggest that a given circuit may use similar mechanisms to perform many forms of modulation and likely reflects a general principle that applies to a wide range of brain areas and sensory, cognitive, and motor processes.

variability | dimensionality | visual cortex | normalization

Understanding the biological mechanisms by which sensory, cognitive, and motor processes modulate sensory responses will likely be critical for understanding how our brains convert information about the sensory world into action. Many such processes modulate sensory responses in qualitatively similar ways, multiplicatively scaling trial-averaged responses (1) and changing the mean noise correlation in a population, defined as the correlation between the responses of a pair of neurons to repeated presentations of the same stimulus (2). In the visual system, these population-averaged effects describe modulation involving changes to the visual stimulus (e.g., stimulus contrast or surround suppression) (1, 3–10), modulation originating from the earliest stages of visual processing in the retina (e.g., adaptation) (1, 11–14), and modulation originating from cognitive processes internal to the nervous system (e.g., attention, task switching, learning, arousal, or multisensory integration) (15–26).

Although these sensory and cognitive processes often have qualitatively similar effects on neuron-averaged sensory responses, they are notable for their heterogeneity. Even among simultaneously recorded neurons with similar tuning properties, these sensory and cognitive processes have diverse effects on rates, noise correlations, and other metrics of neuronal activity (for example, attention has heterogeneous effects on simultaneously recorded visual neurons) (27). Because these processes are typically studied one at a time and/or one neuron at a time, it is unclear whether they act on different neuronal subpopulations and whether they have similar effects on populations.

In particular, there is an emerging body of work that demonstrates a strong relationship between response covariability and behavior (28). Further, covariability strongly constrains

mechanistic models. We and others showed that response covariability in visual cortex is low rank (29–38). This means that shared variability is well described as a low-dimensional process that affects neurons with different weights rather than higher-order interactions between neurons or subpopulations. Furthermore, we showed that attention has an even lower rank effect on covariability (approximately rank one), as evidenced by the observation that the relationship between noise and signal correlation is largely unchanged by attention (39) and by direct measures of the rank of attention-related modulation of shared variability (30–32).

However, many models of cortical circuits (including balanced excitatory–inhibitory networks with fast inhibition or with slow inhibition and broad connectivity) produce high rank variability (31, 40, 41). This indicates a complexity in the covariance matrix that is absent from the data. Low rank covariability indicates that neuronal population responses to repeated presentations of the same stimulus are oriented along a small number of dimensions in neuronal population response space. In contrast, covariability in most models is spread among many dimensions. We and our collaborators showed recently that requiring models to have realistic timescales and connectivity and constraining parameters by observed effects of attention on covariability place strong constraints on the underlying mechanisms (31, 42).

We therefore hypothesized that measuring how different sensory and cognitive processes affect population covariability may reveal whether they could be mediated by analogous mechanisms. We designed our experiment to simultaneously record the effects of multiple sensory and cognitive processes on the same neuronal population and also to evaluate the generality of our findings by recording from multiple brain areas during different behavioral tasks. We therefore have two datasets. Our first dataset consists of simultaneously recorded effects of contrast, adaptation, and spatial attention on the responses of small populations of neurons in visual area V4. We chose those three

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Brain Produces Mind by Modeling” held May 1–3, 2019 at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. NAS colloquia began in 1991 and have been published in PNAS since 1995. From February 2001 through May 2019, colloquia were supported by a generous gift from The Dame Jillian and Dr. Arthur M. Sackler Foundation for the Arts, Sciences, & Humanities, in memory of Dame Sackler’s husband, Arthur M. Sackler. The complete program and video recordings of most presentations are available on the NAS website at <http://www.nasonline.org/brain-produces-mind-by>.

Author contributions: D.A.R., C.X., and M.R.C. designed research; D.A.R., C.X., and L.E.K. performed research; D.A.R., C.X., and M.R.C. contributed new reagents/analytic tools; D.A.R., C.X., L.E.K., F.B., and M.R.C. analyzed data; and D.A.R., C.X., and M.R.C. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹D.A.R. and C.X. contributed equally to this work.

²To whom correspondence may be addressed. Email: cohenm@pitt.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2005797117/-DCSupplemental>.

First published November 23, 2020.

processes for three reasons. First, their responses on the trial-averaged responses of V4 neurons are well described by the same computation (divisive normalization) (1). Second, they are all known to affect the shared variability between pairs of visual neurons (39, 43–47), which makes it possible to compare their effects on the dimensionality of correlated variability. Third, and most importantly, these processes represent a strong test of the hypothesis that many sensory and cognitive processes involve analogous mechanisms because they originate at different stages of visual processing (contrast is a change in the visual stimulus and affects neuronal responses at all processing stages; adaptation affects responses beginning in the retina, and spatial attention more strongly affects later stages of visual processing). We found that although the way contrast, adaptation, and attention modulate a given neuron's mean response was uncorrelated with modulation by other factors, all three of these processes affect covariability in a low rank way.

Our second dataset shows that these low rank effects are not limited to area V4 or to processes that are associated with average changes in firing rate gains and correlations. We found that even when different visual features of a stimulus are encoded by the same group of primary visual cortical (V1) neurons, changing the behavioral relevance of those features produces low rank effects on neuronal response variability. Together, these data are consistent with the idea that despite the differences between contrast, adaptation, attention, and task switching, they affect the responses of populations of neurons in primate visual cortex in similar ways. More broadly, they suggest that simultaneous recordings from small groups of neurons will be key for understanding the biological bases of a large range of neuronal computations.

Results

Simultaneous Psychophysical and Neurophysiological Measurements of Contrast, Adaptation, and Attention. To simultaneously measure the effects of contrast, adaptation, and attention on neuronal populations, two monkeys (*Macaca mulatta*) performed the cued detection task illustrated in Fig. 1. The animals maintained fixation while waiting for the onset of a target (*Methods*). During this time, we presented flashing pairs of static gratings at a range of orientations and contrasts (200-ms duration per stimulus, one stimulus per hemifield). The animal was rewarded for making a saccade to a small red bar within 450 ms of its onset or for maintaining fixation if no bar was presented after 8,000 ms. The animal was cued using unanalyzed instruction trials as to which hemifield the bar was likely to appear in. This cue accurately predicted the side where the bar would appear on subsequent trials 85% of the time. The animals' task performance shows that they respected this attention cue. The animals detected a greater proportion of targets when they appeared at one of the locations in the cued (attended) hemifield compared to in the uncued hemifield. Overall, the animals detected 84% of cued targets and 78% of the uncued targets (paired *t* test, $P < 0.05$).

While the monkeys performed the detection task, we measured the effects of contrast, adaptation, and attention on the spiking of neuronal populations in visual area V4 using chronically implanted microelectrode arrays (*Methods* and Fig. 1*B*). Importantly, we measured contrast, adaptation, and attention simultaneously. We measured the effects of contrast by comparing responses to stimuli with high or low contrast, adaptation by comparing responses to the first or second presentation of a given orientation, and attention by comparing responses when attention was directed toward or away from the units' receptive fields.

Contrast, Adaptation, and Attention Have Diverse and Largely Separable Effects on Individual Units and Pairs of Units. Consistent with previous results, we found that contrast (6, 8, 48, 49), adaptation (11, 13, 50), and spatial attention (51–53) were all

associated with changes in the mean firing rates of V4 units. On average, contrast and attention increased mean rates (Fig. 2*A*); indexes were calculated for each unit as the difference between responses in two conditions divided by their sum (*Methods*). The mean index value comparing high and low contrast = 0.17 and the mean index value comparing attention toward and away from the hemifield containing the unit's receptive field for attention = 0.01 (two-tailed Wilcoxon signed rank test, both $P < 0.001$). Adaptation decreased mean rates (mean index value comparing the first and second presentation of the same orientation = 0.01, two-tailed Wilcoxon signed rank test, $P < 0.001$).

We note that the magnitude of these mean effects, particularly for adaptation, is smaller than that in published results from single-neuron studies (for example, see ref. 54). There are two important differences between our study and traditional studies of the effect of different sensory and cognitive processes in general and in adaptation, in particular. First, unlike in single-neuron studies in which the visual stimulus can be optimized for the tuning properties of the neuron under study, our multi-neuron recordings require us to use a single set of stimuli for all neurons. Therefore, the observed response heterogeneity may be caused by true heterogeneity in the extent to which the sensory and cognitive processes affect each neuron, differences in the extent to which stimuli matched the tuning properties of each neuron, or a combination. Second, to simultaneously measure changes in population responses to contrast, adaptation, and attention, we made several experimental compromises. Typically, adaptation is studied using long (>1 s) adapting stimuli (54). However, to reduce nonstationarities in attention, we used shorter stimulus presentations (200 ms). In addition, each trial consisted of many stimuli, meaning that even the first presentation of a given stimulus was typically preceded by many non-identical stimuli. These factors likely diminished the magnitude of the effects of adaptation compared to that in single-neuron studies.

Our data suggest that contrast, adaptation, and attention affect the firing rates of individual units in a largely separable way (15, 54–62). The extent to which each unit's mean rate was modulated by one of the three processes was weakly, but significantly correlated with modulation by the other processes (Fig. 2*B*; Pearson's correlation coefficients between contrast and adaptation indexes = 0.13, between contrast and attention indexes = -0.1, and between adaptation and attention = -0.12, all $P < 0.001$). These results suggest that contrast, adaptation, and attention are associated with robust changes to mean rates, but do not provide strong evidence that these changes equally target the same subsets of neurons.

Previous studies have demonstrated that a variety of processes including anesthesia (37, 63, 64), attention (39, 47, 65), arousal, alertness, or task difficulty (66–68), and locomotion (69) modulate the strength of covariance or pairwise correlations (2) in sensory areas. Consistent with these reports, we found that contrast, adaptation, and attention were associated with changes in covariability (39, 43–47). On average, increased stimulus contrast increased covariance between pairs of simultaneously recorded neurons (Fig. 3*A*; mean difference comparing high and low contrast = 2.43, two-tailed Wilcoxon signed rank test, $P < 0.001$). Covariance was lower as a result of adaptation (mean difference comparing the first and second presentation of the same orientation = 0.19, two-tailed Wilcoxon signed rank test, $P < 0.001$). Attention directed into a pair of neurons' receptive fields decreased covariance (mean difference comparing attention toward and away from the hemifield containing the unit's receptive field = -0.30, two-tailed Wilcoxon signed rank test, $P < 0.001$). The effects of contrast, adaptation, and attention on covariance were not consistently strongly interrelated (Fig. 3*C*; Pearson's correlation between each unit's mean change in covariance, across all stimulus conditions, with all simultaneously

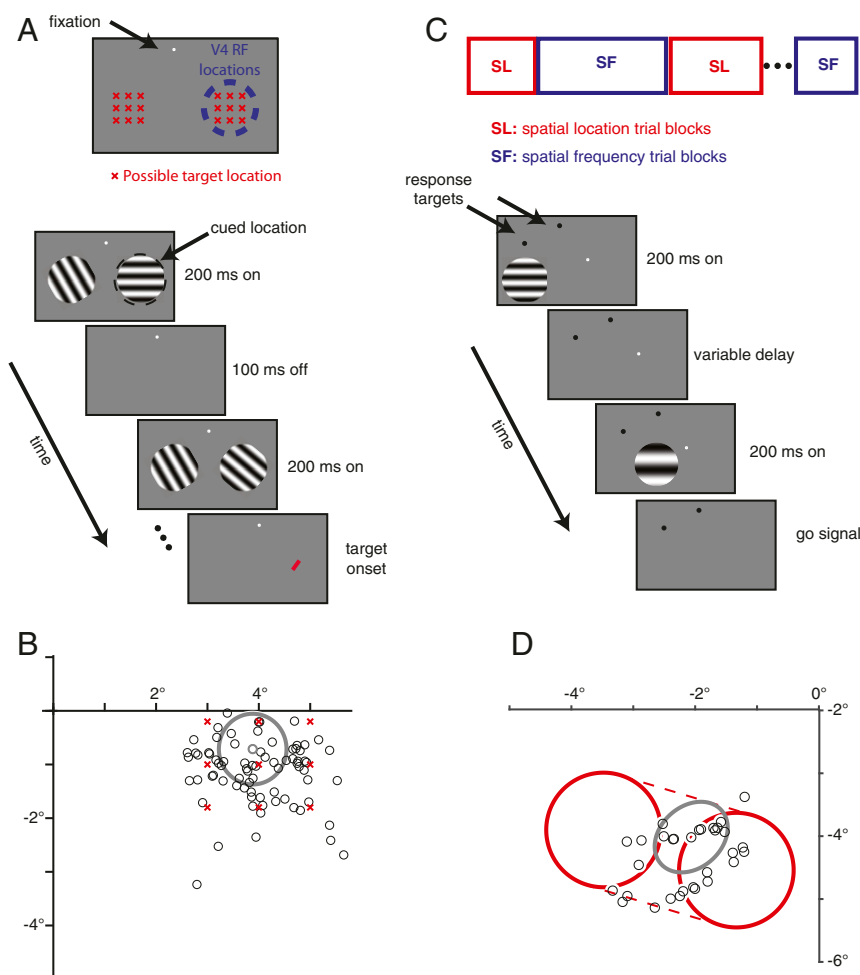


Fig. 1. Methods. (A) Possible target locations in the bar detection task and task structure. Possible target locations are marked with red Xs in *Top*. The animal maintained fixation while static grating stimuli flashed on and off. The animal was rewarded for directing a saccade to the target stimulus (red bar) within 450 ms of its appearance. (B) Estimates of receptive field centers from an example V4 recording session (black circles). An estimate of the size of one example unit's receptive field is drawn in gray. (C) Schematic of the task-switching task. Two subsequent briefly displayed drifting Gabors differ in both their spatial location and spatial frequency, only one of which is behaviorally relevant in each trial. The monkey was rewarded for discriminating the change in the relevant feature while ignoring the change in the irrelevant feature. (D) Estimates of receptive field centers from an example V1 recording session (black circles). An estimate of the size of one example unit's receptive field is drawn in gray. The stimulus size and range of stimulus locations are shown in red.

recorded units = 0.22 between contrast and adaptation, 0.001 between contrast and attention, and -0.07 between adaptation and attention; $P < 0.001$, $P = 0.29$, and $P < 0.001$, respectively).

Contrast, Adaptation, and Attention Are Associated with Low Rank Modulation of Response Covariability. We reasoned that if the effects of contrast, adaptation, and attention on visual responses are mediated by analogous mechanisms, they might all affect covariability in similarly low rank ways. We used factor analysis (36, 70, 71) on the z-scored responses from each condition to evaluate the rank of the changes in response variability associated with each sensory and cognitive process. The majority of the population variance is concentrated in the first eigenmode (62.7% variance explained for the attended, first, high contrast stimulus; 65.2% when that stimulus was unattended; 60.5% when that stimulus was the second stimulus; and 49.3% when that stimulus was low contrast; Fig. 4A). As in previous studies (30–32), we found that attention primarily affects variability in the first eigenmode (Fig. 4A, compare black and red lines). Consistent with the idea that different sensory and cognitive processes affect a small number of dimensions, we also found that contrast and adaptation affected variability predominantly

in this first mode (Fig. 4A, compare black to blue and green lines, respectively).

These changes in covariability can be completely explained by changes in a single dimension of population response space. The changes in covariance associated with each process are essentially abolished when the first eigenmode is removed (31) (Fig. 4B; mean proportion of the total change in covariance that is abolished by removing the first eigenmode for V4 = 0.92 for contrast, = 0.87 for adaptation, and = 0.84 for attention). To assess whether these changes are consistent with changes in a single dimension, given the size of our dataset, we performed a numerical simulation where we estimated the proportion of change in covariance in the first eigenmode as a function of the rank of a simulated covariance change. We constructed simulated datasets using the mean number of units and trials as our real data and generated changes in covariance of different ranks using a procedure described in *Methods*. We then analyzed the simulated data the same way as the real data in Fig. 4B and plotted 1 minus the ratio of the covariance change in the residual and raw data. The mean values for each of the datasets in Fig. 4B and for the V1 data (which are discussed in the next section) are plotted as horizontal lines and the results of the simulation are

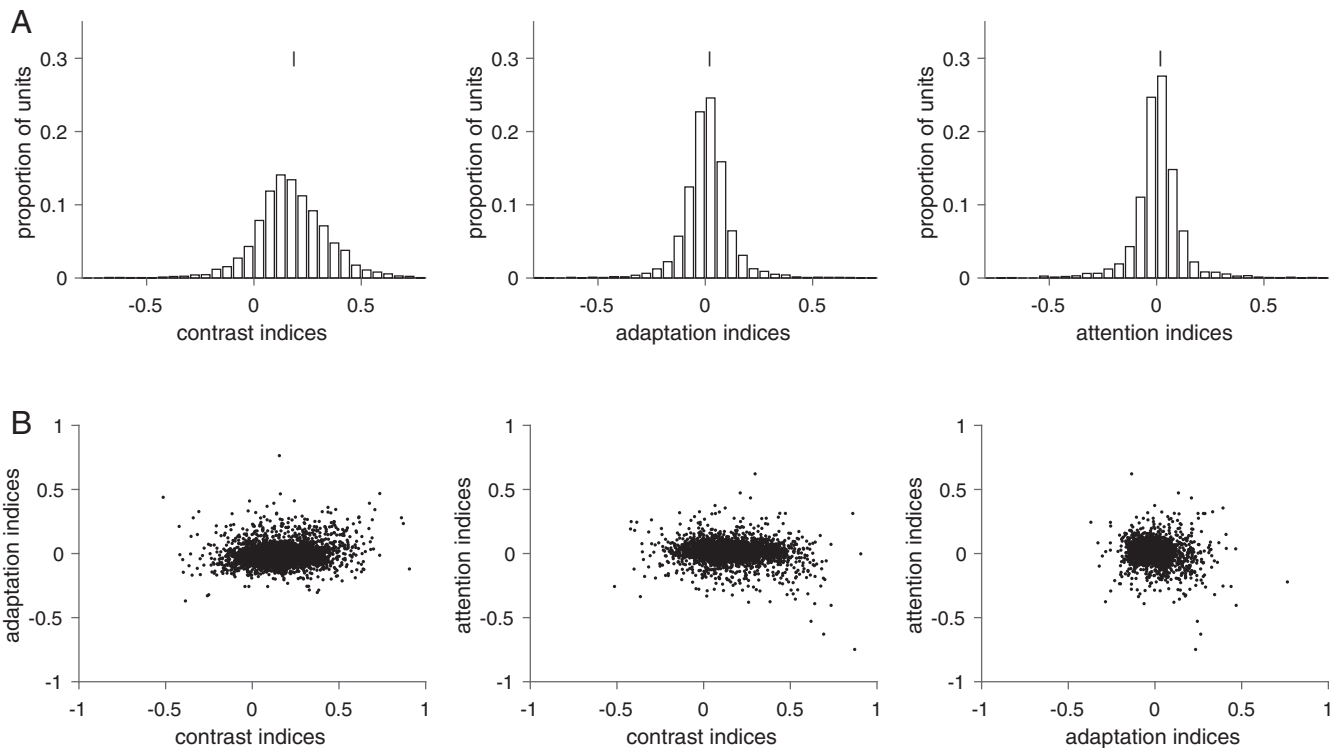


Fig. 2. Contrast, adaptation, and attention affect firing rates in a largely separable manner. (A) Indexes (difference divided by the sum of responses for two conditions) for each unit for contrast, adaptation, and attention. Distribution means are indicated by a vertical line. A total of 5,328 units are shown from 10 sessions. (B) The effects of contrast, adaptation, and attention on individual units are weakly related.

plotted as a solid black line. Given the size of our dataset, the magnitude of the changes in covariance that we measured is consistent with a change in a single dimension (Fig. 4C; black dots signify the prediction from the simulation for rank 1 and rank 2 changes, respectively, from left to right). In principle, if we recorded from much larger populations, we might find modulations that span more than a single dimension. However, these analyses show that while contrast, adaptation, and attention are associated with diverse (and different signed) modulations of mean responses and covariance in visual cortex, they all modulate responses in a small number of dimensions. These results are therefore consistent with the hypothesis that many distinct sensory and cognitive processes affect populations of V4 neurons via similar mechanisms.

Low Rank Modulation of Response Variability Is Not Limited to Homogeneous Processes. One possibility is that the modulation we observed is low rank because contrast, adaptation, and attention are each associated with relatively homogeneous effects on the groups of neurons we recorded. For example, contrast typically increases the responses of V4 neurons (60), making it possible that contrast has a low rank, monolithic effect on population responses. A related possibility is that trial-to-trial variability in these processes (such as, for example, the animal's internal attention state) is the source of this low rank variability (32, 72). The different sign relationships between the strength of the low rank modulation and performance [e.g., better behavioral performance is associated with lower variability as a result of attention (28) but stimuli are more readily detected when they are higher contrast, which is associated with higher variability (73)] make it unlikely that the explanation is that simple, but measuring trial-to-trial variability in internal states is notoriously difficult.

While contrast, adaptation, and attention are maximally different in terms of the stage of visual processing in which they originate, they have something important in common in the context of our experiment. All three processes involve a non-specific drive to the population of neurons we recorded. We reasoned that if we were ever going to observe higher rank changes in population response variability, it might be in a situation in which the stimulus or task manipulation had qualitatively different effects on different subsets of neurons.

To test the generality of our results, we recorded the responses of neurons in primary visual cortex (V1) while two different monkeys discriminated changes in one of two different stimulus features. Changing the behavioral relevance of stimulus features will presumably induce heterogeneous effects on neurons in a way that depends on their tuning preferences for these stimulus features. Specifically, the monkeys discriminated either the spatial frequency or the spatial location of two consecutively presented Gabor patches. In alternating blocks of trials, they indicated whether the spatial frequency of the Gabor stimuli increased or decreased or whether the location of the stimuli moved to the left or to the right (Fig. 1 C and D). The change amounts were chosen so that the difficulties of the two tasks would be comparable (monkey 3 average performance, 74% for spatial location task, 72% for spatial frequency task; monkey 4 average performance, 79% for spatial location task, 80% for spatial frequency task; these differences were not significant across sessions: $P = 0.12$ for monkey 3, $P = 0.37$ for monkey 4, paired t test). We expected that switching between task relevant features would induce heterogeneous effects on the recorded neuronal population because the two features were encoded with very different combinations of V1 neurons. On average, the best linear decoder for one feature correctly predicted 77% of the changes in the same feature in held-out test trials, but only 58% of the changes in the other feature (significantly different across

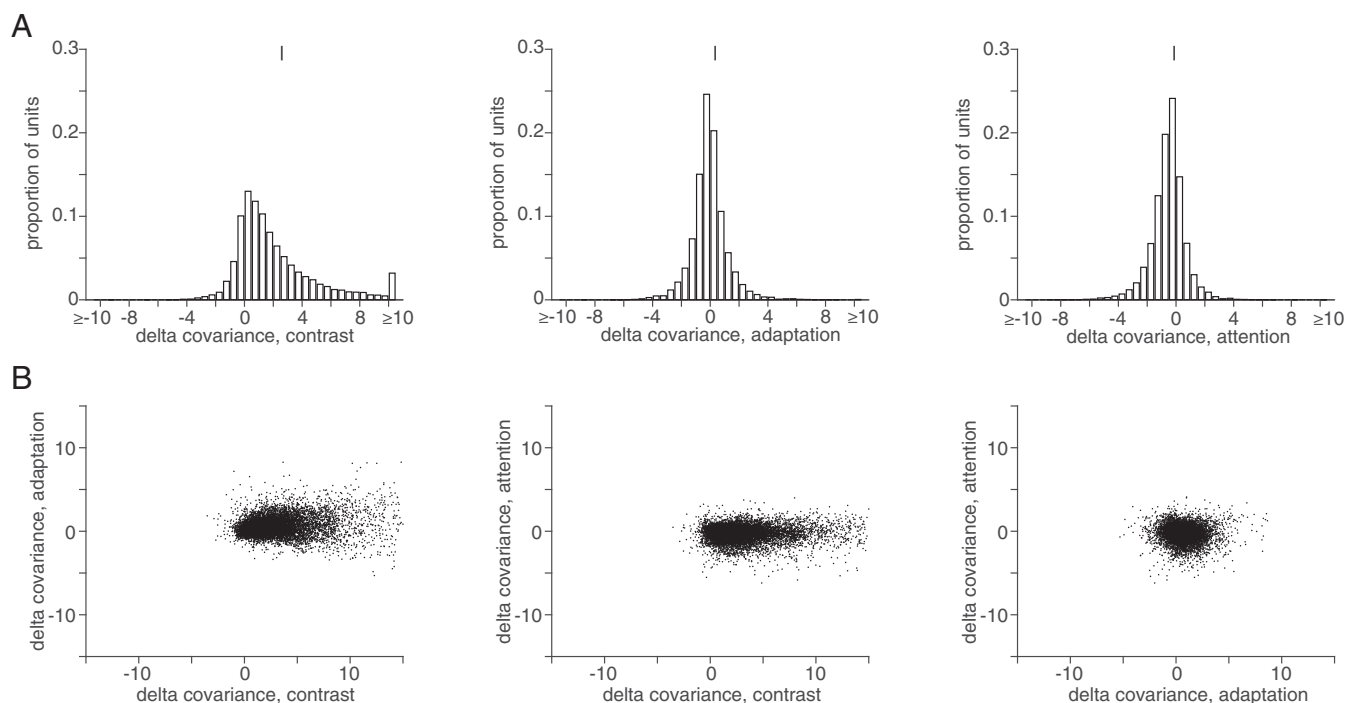


Fig. 3. Contrast, adaptation, and attention affect covariance between pairs of units. (A) Histograms of covariance changes for contrast, adaptation, and attention. Distribution means are indicated by a vertical line. A total of 20,774 pairs are shown from 10 sessions. (B) Covariance changes associated with contrast, adaptation, and attention on individual units are weakly related.

sessions, $P < 0.001$, two-tailed Wilcoxon signed rank test), with chance performance at 50%. Overall, 54.2% of units had higher mean rates in the spatial frequency than the spatial location task, and 57.0% of unit pairs had higher mean noise correlations in the spatial frequency than the spatial location task. Switching between spatial frequency and location discrimination tasks did not lead to an overall change in mean firing rates or for pairwise covariance ($P = 0.8$ and $P = 0.4$, Wilcoxon signed rank test, respectively). Therefore, by comparing the covariability in blocks where different encoding dimensions are behaviorally relevant, we can test whether higher rank changes result from task manipulations that induce heterogeneous modulation.

We found that task switching affected the covariability of groups of V1 neurons in a low rank manner that was analogous to the effects of contrast, adaptation, and attention on groups of V4 neurons (Fig. 4B). Like in V4, the covariance change was almost exclusively oriented along the primary axis (mean proportion of total change in covariance for V1 resulting from task switching = 0.92). These results suggest that the low rank modulation of population covariability we observed in V4 is not limited to extrastriate cortex or to processes with relatively homogeneous effects on neuronal populations.

Challenges of Determining whether Contrast, Adaptation, and Attention Affect Covariability along Similar Dimensions of Neuronal Population Activity. The analyses in Fig. 4 show that all four sensory and cognitive processes affect covariability in a low rank way, primarily along the first eigenmode of the covariance matrix. Because the covariance matrix in each condition is analyzed separately (including using factor analysis to remove variability that is private to each unit and orthonormalizing the result) (70), these analyses do not make it clear whether they affect the same dimension of neuronal population space or whether the first eigenmode changes between conditions.

Unfortunately, given the size of our datasets and the sensitivity of these measurements to noise, we lack the statistical power to answer the question of whether the first eigenmode changes as a result of these sensory and cognitive processes. Consistent with our previous simulations (74), we found that our data do little to constrain estimates of the dimension affected by each process (SI Appendix, Fig. S1). These results show that experimentally tractable datasets can determine whether covariability changes are low or high rank but serve as a cautionary tale for future efforts to identify dimensions of neural population space using limited data.

Discussion

Constraints on Mechanistic Models. We showed that although contrast, adaptation, attention, and task switching have quantitatively and qualitatively different effects on the responses of visual cortical neurons, they all affect shared response variability in a low rank way. This observation is important because it places strong constraints on models regarding how a variety of sensory and cognitive processes affect cortical circuits.

We and our collaborators showed (31) that although many models produce variability whose magnitude matches that of observed pairwise noise correlations, the only model that produces low rank fluctuations is one in which inhibition is slower than excitation (which is physiologically realistic) (75, 76–78) and in which the connectivity of inhibitory neurons is spatially restricted (which is also realistic) (79–81). In that model, attention could exert low rank changes in shared variability through an input whose effect is to change the balance between excitation and inhibition, increasing the activity or influence of inhibitory neurons relative to excitatory ones (31). This simple mechanism is a good candidate for a general mechanism mediating a broad class of sensory and cognitive processes in primate visual cortex. Our previously published model simulates the responses of a restricted population of neurons (e.g., one hypercolumn). It will

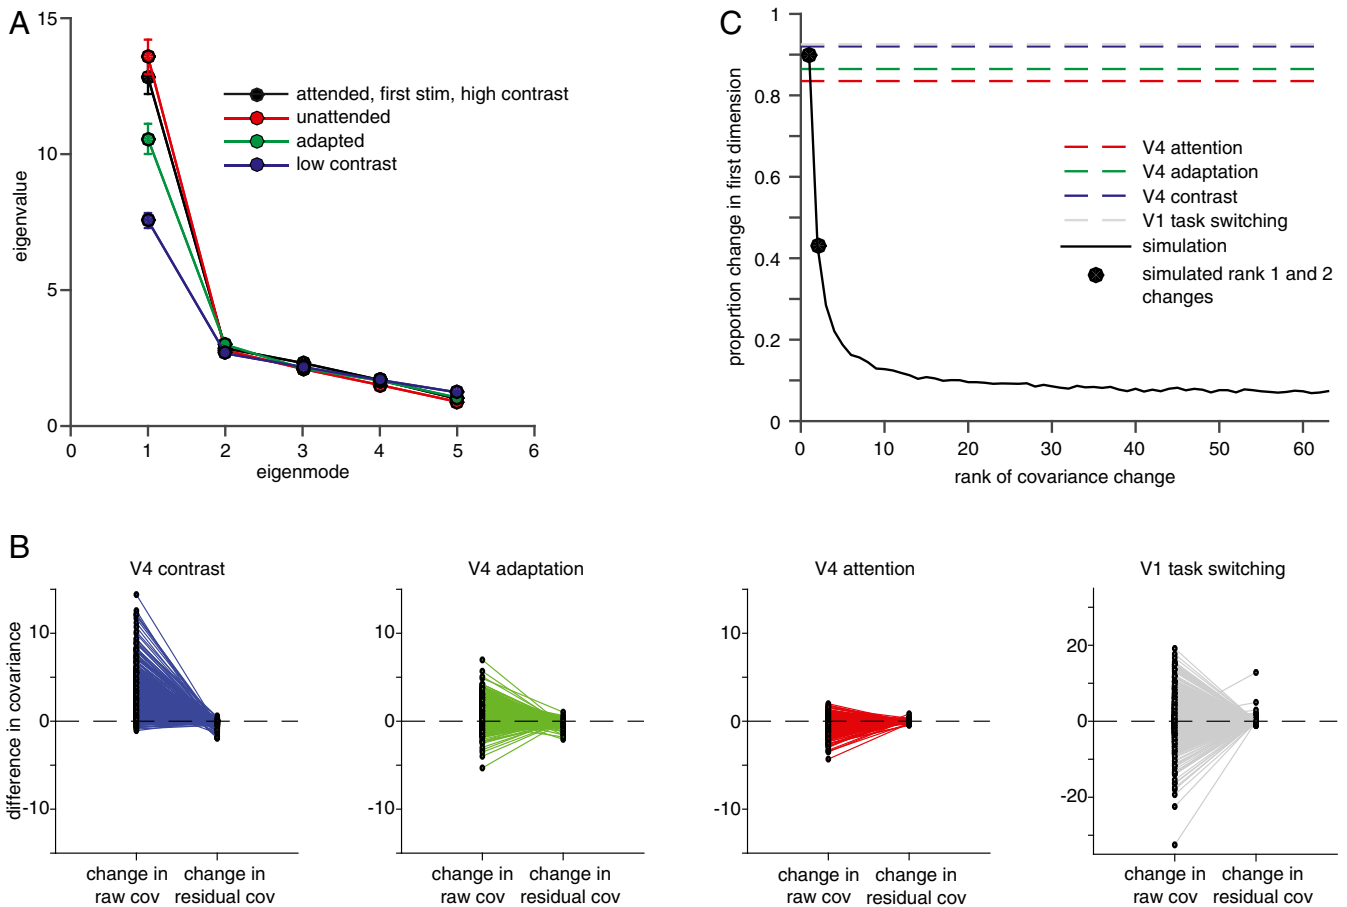


Fig. 4. Contrast, adaptation, and attention affect population responses in a small number of dimensions. (A) Eigenvalues (shared variance) of the population covariance matrix corresponding to the first five modes calculated using factor analysis. (B) Changes in average covariance associated with contrast, adaptation, and attention in V4 or task switching in V1 between each unit's raw (*Left*) or residual (first eigenmode removed) (*Right*) responses and all other simultaneously recorded units. Positive numbers mean that covariance was higher for high contrast, the first stimulus, the attended stimulus, or the location discrimination task, respectively. (C) Proportion change in covariance in the first eigenmode as a function of the rank of the covariance change in simulation. The mean covariance change values for each of the datasets in B are plotted as horizontal lines in the corresponding color. The black line represents the simulated proportion change in the first dimension as a function of the simulated rank of covariance change. Black dots are the simulated prediction for rank 1 and rank 2 changes, respectively, from left to right. This simulation demonstrates that, because of noise, even when a change is rank 1, slightly less than 100% of the change in covariability is in the first dimension, but that the proportion of change in that first dimension falls off dramatically as rank increases.

be exciting to determine whether simple control signals that alter the balance between excitation and inhibition continue to produce low rank changes in a broader swath of modeled (31, 82) or measured cortex.

Links with Normalization. Divisive normalization, in which responses are scaled by the total stimulus drive, describes how a wide variety of sensory, cognitive, and motor processes (including the effects of contrast, adaptation, and attention in visual cortex) affect the trial-averaged firing rates of individual neurons (1). The link between normalization-related sensory and cognitive processes and the covariability of population responses is not well understood. However, we recently showed that a simple normalization model can account for the changes in variability observed between neurons in different cortical areas (21), suggesting that, just as with the divisive scaling of average responses, changes in covariability may be a signature of normalization computations. The observation that each of the tested sensory and cognitive processes had a similar effect on covariability is evidence in favor of the tantalizing hypothesis that task switching also involves normalization.

The surface level similarities between the ways such a wide variety of sensory and cognitive processes affect average responses of neurons gave rise to the hypothesis that all normalization-related processes may utilize the same mechanism (1). This idea appears to be contradicted by two recent results showing that in mouse visual cortex, normalization primarily involves changes in excitation (83), but in the *Drosophila* antenna lobe, normalization is mediated through changes in inhibition (84). Our model leaves room for both of these possibilities: Normalization is accomplished by shifting the excitatory and inhibitory (E/I) balance, which could involve changes to excitation or inhibition or both (31, 42). One possibility is that both mouse visual cortex and the *Drosophila* antenna lobe use the same mechanism (changing E/I balance) but employ different means to change that balance. Another possibility is that each circuit (e.g., primate V1 or V4) uses the same mechanism for all forms of normalization-related computations, but that different systems use different mechanisms.

A recent set of studies suggested a mechanism in which each neuron has a set of inputs that perform normalization and that attention and other sensory and cognitive processes act through those inputs (18–22, 85, 86). This idea is supported by the

observation that the extent to which mean responses of individual neurons in visual cortex are modulated by switching attention between two stimuli within their receptive fields is highly correlated with the magnitude of response suppression associated with placing an additional stimulus in the receptive field (18–20, 22, 85). However, this strong correlation does not seem to apply to all attention-related modulation: modulation by feature attention is not correlated with multistimulus suppression (25), and here we showed that modulations by contrast, adaptation, and attention are weakly correlated.

These observations suggest that while all normalization-related computations may involve the same or similar circuit mechanisms, the involvement of particular neurons likely depends on the specific modulatory process or circuit. In some form, the idea that the involvement of specific neurons depends on the specific context is necessary to explain the observations that the extent of firing rate modulation by, for example, feature attention, depends on how closely a neuron's tuning matches the attended stimulus (87–91).

Implications for Information Coding. A curious consequence of the observation that sensory and cognitive processes affect response variability in a low rank way is that those modulations could have little to no effect on the stimulus information that could be read out of a neuronal population by an optimal decoder. An elegant series of theoretical studies showed that shared variability affects the Fisher information encoded in a neuronal population only when it is aligned with the dimensions in population space that are being read out (92–96). Recent experimental studies have begun to characterize the extent to which correlated variability aligns with signal encoding (97–101). In the future, it will be interesting to determine the extent to which variability aligns with the dimensions of population activity that are used by the animal during behavior (28, 74, 102, 103). If contrast, adaptation, and attention affect only a small number of dimensions in population space, their effects could easily be discounted by an optimal decoder.

Why then are contrast, adaptation, attention, and task switching associated with such large changes in behavior in visually guided tasks? Our results are consistent with only two possibilities: 1) that the modulations in shared variability in visual cortex associated with these processes do not affect perception or 2) that the small number of dimensions that are affected by these sensory and cognitive processes are also the ones that are used to guide behavior. While the first possibility has been proposed for attention (in favor of an alternative model in which attention is mediated through changes in the visual

information that is communicated to downstream areas involved in decision making) (74, 103–106), it seems implausible for a low-level process like contrast. Contrast is associated with changes from the earliest stages of visual processing, which are then communicated to visual cortex. The idea that contrast-related changes in visual cortex are not associated with the concomitant changes in perception seems implausible.

On its surface, the second possibility, that visual information is read out in a manner such that low rank changes in variability affect perception, seems equally implausible. It might seem as though readout mechanisms should have access to higher dimensional representations of visual information, which would minimize the impact of low rank changes in response variability. However, we showed recently that during an orientation change detection task, monkeys' choices were much more closely aligned with the axes in V4 population space that correspond to shared variability than would be expected if decoding included more dimensions of population activity (28). One possibility is that, because of either a biological constraint or a need to optimize something other than performance on a specific task (74), responses oriented along the dimensions affected by sensory and cognitive processes carry an outsized influence on perception. Determining the validity of these two possibilities will be an important avenue for further work.

Methods

Using chronically implanted Utah arrays, we recorded population spiking activity from two adult male rhesus monkeys (*M. mulatta*) in visual area V4 while they performed a target detection task. All animal procedures were approved by the Institutional Animal Care and Use Committees of the University of Pittsburgh and Carnegie Mellon University. While monkeys performed this task, we measured the effects of contrast, adaptation, and spatial attention on the responses of individual and simultaneously recorded populations of neurons. In a separate set of experiments, also using Utah arrays, we recorded population spiking activity in primary visual cortex (V1) from two different adult male rhesus monkeys while they performed a task that required them to discriminate either the spatial location or spatial frequency of sequentially presented Gabor stimuli. See *SI Appendix* for detailed methods.

Data Availability. .mat data files and .m functions have been deposited in a Figshare repository, <https://doi.org/10.6084/m9.figshare.13009253>.

ACKNOWLEDGMENTS. We are grateful to Karen McCracken for providing technical assistance and to Amy Ni for comments on an earlier version of this manuscript. Support to M.R.C. is from NIH Grants R00EY020844 and R01EY022930 and Core Grant P30 EY008098s, the McKnight Foundation, the Whitehall Foundation, the Sloan Foundation, and the Simons Foundation.

1. M. Carandini, D. J. Heeger, Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2011).
2. M. R. Cohen, A. Kohn, Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
3. A. Ayaz, A. B. Saleem, M. L. Schölvinc, M. Carandini, Locomotion controls spatial integration in mouse visual cortex. *Curr. Biol.* **23**, 890–894 (2013).
4. J. J. Nassi, C. Gómez-Laberge, G. Kreiman, R. T. Born, Corticocortical feedback increases the spatial extent of normalization. *Front. Syst. Neurosci.* **8**, 105 (2014).
5. M. Carandini, D. J. Heeger, Summation and division by neurons in primate visual cortex. *Science* **264**, 1333–1336 (1994).
6. M. Carandini, D. J. Heeger, J. A. Movshon, Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* **17**, 8621–8644 (1997).
7. E. P. Simoncelli, O. Schwartz, "Modeling surround suppression in V1 neurons with a statistically-derived normalization model" in *Advances in Neural Information Processing Systems*, M. S. Kearns, S. A. Solla, D. A. Cohn, Eds. (MIT Press, Cambridge, MA, 1999), Vol. vol. 11, pp. 153–159.
8. D. J. Heeger, Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992).
9. L. Busse, A. R. Wade, M. Carandini, Representation of concurrent stimuli by population activity in visual cortex. *Neuron* **64**, 931–942 (2009).
10. K. H. Britten, H. W. Heuer, Spatial summation in the receptive fields of MT neurons. *J. Neurosci.* **19**, 5074–5084 (1999).
11. S. G. Solomon, A. Kohn, Moving sensory adaptation beyond suppressive effects in single neurons. *Curr. Biol.* **24**, R1012–R1022 (2014).
12. R. Shapley, C. Enroth-Cugell, Visual adaptation and retinal gain controls. *Prog. Retinal Res.* **3**, 263–346 (1984).
13. A. Kohn, Visual adaptation: Physiology, mechanisms, and functional benefits. *J. Neurophysiol.* **97**, 3155–3164 (2007).
14. B. P. Olveczky, S. A. Baccus, M. Meister, Retinal adaptation to object motion. *Neuron* **56**, 689–700 (2007).
15. M. Sanayei, J. L. Herrero, C. Distler, A. Thiele, Attention and normalization circuits in macaque V1. *Eur. J. Neurosci.* **41**, 949–964 (2015).
16. T. Ohshiro, D. E. Angelaki, G. C. DeAngelis, A normalization model of multisensory integration. *Nat. Neurosci.* **14**, 775–782 (2011).
17. T. Ohshiro, D. E. Angelaki, G. C. DeAngelis, A neural signature of divisive normalization at the level of multisensory integration in primate cortex. *Neuron* **95**, 399–411.e8 (2017).
18. J. Lee, J. H. R. Maunsell, A normalization model of attentional modulation of single unit responses. *PLoS One* **4**, e4651 (2009).
19. A. M. Ni, J. H. R. Maunsell, Spatially tuned normalization explains attention modulation variance within neurons. *J. Neurophysiol.* **118**, 1903–1913 (2017).

20. J. H. Reynolds, D. J. Heeger, The normalization model of attention. *Neuron* **61**, 168–185 (2009).
21. D. A. Ruff, M. R. Cohen, A normalization model suggests that attention changes the weighting of inputs between visual areas. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E4085–E4094 (2017).
22. A. M. Ni, S. Ray, J. H. R. Maunsell, Tuned normalization explains the size of attention modulations. *Neuron* **73**, 803–813 (2012).
23. K. Louie, M. W. Khaw, P. W. Glimcher, Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 6139–6144 (2013).
24. H. Yamada, K. Louie, A. Tymula, P. W. Glimcher, Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nat. Commun.* **9**, 162 (2018).
25. A. M. Ni, J. H. R. Maunsell, Neuronal effects of spatial and feature attention differ due to normalization. *J. Neurosci.* **39**, 5493–5505 (2019).
26. G. M. Boynton, A framework for describing the effects of attention on visual responses. *Vision Res.* **49**, 1129–1143 (2009).
27. M. R. Cohen, J. H. R. Maunsell, A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* **30**, 15241–15253 (2010).
28. A. M. Ni, D. A. Ruff, J. J. Alberts, S. Symmonds, M. R. Cohen, Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018).
29. R. L. T. Goris, J. A. Movshon, E. P. Simoncelli, Partitioning neuronal variability. *Nat. Neurosci.* **17**, 858–865 (2014).
30. N. C. Rabinowitz, R. L. Goris, M. Cohen, E. P. Simoncelli, Attention stabilizes the shared gain of V4 populations. *eLife* **4**, e08998 (2015).
31. C. Huang *et al.*, Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**, 337–348.e4 (2019).
32. A. S. Ecker, G. H. Denfield, M. Bethge, A. S. Tolia, On the structure of neuronal population activity under fluctuations in attentional state. *J. Neurosci.* **36**, 1775–1789 (2016).
33. I. C. Lin, M. Okun, M. Carandini, K. D. Harris, The nature of shared cortical variability. *Neuron* **87**, 644–656 (2015).
34. B. R. Cowley, M. A. Smith, A. Kohn, B. M. Yu, Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Comput. Biol.* **12**, e1005185 (2016).
35. J. D. Smedo, A. Zandvakili, C. K. Machens, B. M. Yu, A. Kohn, Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259.e4 (2019).
36. R. C. Williamson *et al.*, Scaling properties of dimensionality reduction for neural populations and network models. *PLoS Comput. Biol.* **12**, e1005141 (2016).
37. A. S. Ecker *et al.*, State dependence of noise correlations in macaque primary visual cortex. *Neuron* **82**, 235–248 (2014).
38. C. Stringer *et al.*, Inhibitory control of correlated intrinsic variability in cortical networks. *eLife* **5**, e19695 (2016).
39. M. R. Cohen, J. H. R. Maunsell, Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
40. M. N. Shadlen, W. T. Newsome, The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
41. C. van Vreeswijk, H. Sompolinsky, Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
42. T. Kanashiro, G. K. Ocker, M. R. Cohen, B. Doiron, Attentional modulation of neuronal variability in circuit models of cortex. *eLife* **6**, e23978 (2017).
43. E. Zavitz, H. H. Yu, E. G. Rowe, M. G. Rosa, N. S. Price, Rapid adaptation induces persistent biases in population codes for visual motion. *J. Neurosci.* **36**, 4579–4590 (2016).
44. D. A. Gutnitsky, V. Dragoi, Adaptive coding of visual information in neural populations. *Nature* **452**, 220–224 (2008).
45. A. Benucci, A. B. Saleem, M. Carandini, Adaptation maintains population homeostasis in primary visual cortex. *Nat. Neurosci.* **16**, 724–729 (2013).
46. A. Kohn, M. A. Smith, Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
47. J. F. Mitchell, K. A. Sundberg, J. H. Reynolds, Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879–888 (2009).
48. D. J. Tolhurst, J. A. Movshon, I. D. Thompson, The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp. Brain Res.* **41**, 414–419 (1981).
49. A. Dean, The relationship between response amplitude and contrast for cat striate cortical neurones. *J. Physiol.* **318**, 413–427 (1981).
50. M. A. Webster, Visual adaptation. *Annu. Rev. Vis. Sci.* **1**, 547–567 (2015).
51. J. H. Reynolds, L. Chelazzi, Attentional modulation of visual processing. *Annu. Rev. Neurosci.* **27**, 611–647 (2004).
52. J. H. R. Maunsell, Neuronal mechanisms of visual attention. *Annu. Rev. Vis. Sci.* **1**, 373–391 (2015).
53. R. Desimone, J. Duncan, Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (1995).
54. A. Kohn, J. A. Movshon, Neuronal adaptation to visual motion in area MT of the macaque. *Neuron* **39**, 681–691 (2003).
55. A. Pooresmaeili, J. Poort, A. Thiele, P. R. Roelfsema, Separable codes for attention and luminance contrast in the primary visual cortex. *J. Neurosci.* **30**, 12701–12711 (2010).
56. A. Thiele, A. Pooresmaeili, L. S. Delicato, J. L. Herrero, P. R. Roelfsema, Additive effects of attention and stimulus contrast in primary visual cortex. *Cereb. Cortex* **19**, 2970–2981 (2009).
57. K. Anton-Erxleben, V. M. Stephan, S. Treue, Attention reshapes center-surround receptive field structure in macaque cortical area MT. *Cereb. Cortex* **19**, 2466–2478 (2009).
58. J. Martínez-Trujillo, S. Treue, Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* **35**, 365–370 (2002).
59. J. Lee, J. H. R. Maunsell, The effect of attention on neuronal responses to high and low contrast stimuli. *J. Neurophysiol.* **104**, 960–971 (2010).
60. J. H. Reynolds, T. Pasternak, R. Desimone, Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703–714 (2000).
61. A. Kohn, J. A. Movshon, Adaptation changes the direction tuning of macaque MT neurons. *Nat. Neurosci.* **7**, 764–772 (2004).
62. M. Ghodrati, E. Zavitz, M. G. P. Rosa, N. S. C. Price, Contrast and luminance adaptation alter neuronal coding and perception of stimulus orientation. *Nat. Commun.* **10**, 941 (2019).
63. K. D. Harris, A. Thiele, Cortical state and attention. *Nat. Rev. Neurosci.* **12**, 509–523 (2011).
64. M. L. Schölvinck, A. B. Saleem, A. Benucci, K. D. Harris, M. Carandini, Cortical state determines global variability and correlations in visual cortex. *J. Neurosci.* **35**, 170–178 (2015).
65. D. A. Ruff, M. R. Cohen, Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* **17**, 1591–1597 (2014).
66. M. Vinck, R. Batista-Brito, U. Knoblich, J. A. Cardin, Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740–754 (2015).
67. M. J. McGinley, S. V. David, D. A. McCormick, Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron* **87**, 179–192 (2015).
68. D. A. Ruff, M. R. Cohen, Global cognitive factors modulate correlated response variability between V4 neurons. *J. Neurosci.* **34**, 16408–16416 (2014).
69. S. Erisken *et al.*, Effects of locomotion extend throughout the mouse early visual system. *Curr. Biol.* **24**, 2899–2907 (2014).
70. B. M. Yu *et al.*, Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).
71. J. P. Cunningham, B. M. Yu, Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
72. G. H. Denfield, A. S. Ecker, T. J. Shinn, M. Bethge, A. S. Tolia, Attentional fluctuations induce shared variability in macaque primary visual cortex. *Nat. Commun.* **9**, 2654 (2018).
73. Z.-L. Lu, B. A. Doshier, External noise distinguishes attention mechanisms. *Vision Res.* **38**, 1183–1198 (1998).
74. D. A. Ruff, A. M. Ni, M. R. Cohen, Cognition as a window into neuronal population space. *Annu. Rev. Neurosci.* **41**, 77–97 (2018).
75. J. R. Geiger, J. Lübke, A. Roth, M. Frotscher, P. Jonas, Submillisecond AMPA receptor-mediated signaling at a principal neuron-interneuron synapse. *Neuron* **18**, 1009–1023 (1997).
76. Z. Xiang, J. Huguenard, D. Prince, GABAA receptor-mediated currents in interneurons and pyramidal cells of rat visual cortex. *J. Physiol.* **506**, 715–730 (1998).
77. P. Salin, D. Prince, Spontaneous GABAA receptor-mediated inhibitory currents in adult rat somatosensory cortex. *J. Neurophysiol.* **75**, 1573–1588 (1996).
78. M. C. Angulo, J. Rossier, E. Audinat, Postsynaptic glutamate receptors and integrative properties of fast-spiking interneurons in the rat neocortex. *J. Neurophysiol.* **82**, 1295–1302 (1999).
79. R. B. Levy, A. D. Reyes, Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex. *J. Neurosci.* **32**, 5609–5619 (2012).
80. S. Horvát *et al.*, Spatial embedding and wiring cost constrain the functional layout of the cortical network of rodents and primates. *PLoS Biol.* **14**, e1002512 (2016).
81. J. Mariño *et al.*, Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat. Neurosci.* **8**, 194–201 (2005).
82. C. Huang, A. Pouget, B. Doiron, Internally generated population activity in cortical networks hinders information transmission. *bioRxiv*:2020.02.03.932723 (4 February 2020).
83. T. K. Sato, B. Haider, M. Häusser, M. Carandini, An excitatory basis for divisive normalization in visual cortex. *Nat. Neurosci.* **19**, 568–570 (2016).
84. S. R. Olsen, V. Bhandawat, R. I. Wilson, Divisive normalization in olfactory population codes. *Neuron* **66**, 287–299 (2010).
85. B. E. Verhoef, J. H. R. Maunsell, Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat. Neurosci.* **20**, 969–977 (2017).
86. D. A. Ruff, M. R. Cohen, Stimulus dependence of correlated variability across cortical areas. *J. Neurosci.* **36**, 7546–7556 (2016).
87. J. C. Martínez-Trujillo, S. Treue, Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* **14**, 744–751 (2004).
88. J. H. R. Maunsell, S. Treue, Feature-based attention in visual cortex. *Trends Neurosci.* **29**, 317–322 (2006).
89. D. A. Ruff, R. T. Born, Feature attention for binocular disparity in primate area MT depends on tuning strength. *J. Neurophysiol.* **113**, 1545–1555 (2015).
90. S. Treue, J. C. Martínez Trujillo, Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).

91. C. J. McAdams, J. H. R. Maunsell, Attention to both space and feature modulates neuronal responses in macaque area V4. *J. Neurophysiol.* **83**, 1751–1755 (2000).
92. R. Moreno-Bote *et al.*, Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
93. I. Kanitscheider, R. Coen-Cagli, A. Pouget, Origin of information-limiting noise correlations. *Proc. Natl. Acad. Sci. U.S.A.* **112**, E6973–E6982 (2015).
94. A. Kohn, R. Coen-Cagli, I. Kanitscheider, A. Pouget, Correlations and neuronal population information. *Annu. Rev. Neurosci.* **39**, 237–256 (2016).
95. M. Shamir, H. Sompolinsky, Implications of neuronal diversity on population coding. *Neural Comput.* **18**, 1951–1986 (2006).
96. A. S. Ecker, P. Berens, A. S. Tolias, M. Bethge, The effect of noise correlations in populations of diversely tuned neurons. *J. Neurosci.* **31**, 14272–14283 (2011).
97. A. G. Bondy, R. M. Haefner, B. G. Cumming, Feedback determines the structure of correlated variability in primary visual cortex. *Nat. Neurosci.* **21**, 598–606 (2018).
98. J. S. Montijn *et al.*, Strong information-limiting correlations in early visual areas. *bioRxiv*:10.1101/842724 (15 November 2019).
99. C. Stringer *et al.*, Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
100. O. I. Rumyantsev *et al.*, Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).
101. R. Bartolo, R. C. Saunders, A. R. Mitz, B. B. Averbeck, Information-limiting correlations in large neural populations. *J. Neurosci.* **40**, 1668–1678 (2020).
102. K. J. Lakshminarasimhan, A. Pouget, G. C. DeAngelis, D. E. Angelaki, X. Pitkow, Inferring decoding strategies for multiple correlated neural populations. *PLoS Comput. Biol.* **14**, e1006371 (2018).
103. D. A. Ruff, M. R. Cohen, Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nat. Neurosci.* **22**, 1669–1676 (2019).
104. P. Fries, Rhythms for cognition: Communication through coherence. *Neuron* **88**, 220–235 (2015).
105. T. Womelsdorf, P. Fries, The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* **17**, 154–160 (2007).
106. E. K. Miller, T. J. Buschman, Cortical circuits for the control of attention. *Curr. Opin. Neurobiol.* **23**, 216–222 (2013).