

Spontaneous Behaviors Drive Multidimensional, Brain-wide Activity

Carsen Stringer^{*1,2}, Marius Pachitariu^{*1,3,4}, Nicholas Steinmetz^{5,6},
Charu Bai Reddy⁵, Matteo Carandini^{†5} and Kenneth D. Harris^{†3,4}

¹ HHMI Janelia Research Campus, Ashburn, Virginia, 20147, USA.

² Gatsby Computational Neuroscience Unit, UCL, London W1T 4JG, UK.

³ UCL Institute of Neurology, London WC1E 6DE, UK.

⁴ UCL Department of Neuroscience, Physiology, and Pharmacology, London WC1E 6DE, UK.

⁵ UCL Institute of Ophthalmology, London EC1V 9EL, UK.

⁶ Present address: Department of Biological Structure, University of Washington, Seattle, WA 98195, USA.

* equal first authors, † equal senior authors.

Neuronal populations in sensory cortex produce variable responses to sensory stimuli, and exhibit intricate spontaneous activity even without external sensory input. Cortical variability and spontaneous activity have been variously proposed to represent random noise, recall of prior experience, or encoding of ongoing behavioral and cognitive variables. Recording over 10,000 neurons in mouse visual cortex, we observed that spontaneous activity reliably encoded a high-dimensional latent state, which was partially related to the mouse’s ongoing behavior and was represented not just in visual cortex but across the forebrain. Sensory inputs did not interrupt this ongoing signal, but added onto it a representation of external stimuli in orthogonal dimensions. Thus, visual cortical population activity, despite its apparently noisy structure, reliably encodes an orthogonal fusion of sensory and multidimensional behavioral information.

In the absence of sensory inputs, the brain produces structured patterns of activity, which can be as large as or larger than sensory-driven activity (1). Ongoing activity exists even in primary sensory cortices, and has been hypothesized to reflect recapitulation or expectation of sensory experience. This hypothesis is supported by studies that found similarities between sensory-driven and spontaneous firing events (2–7). Alternatively, ongoing activity could be related to behavioral and cognitive states. The firing of sensory cortical and sensory thalamic neurons correlates with behavioral variables such as locomotion, pupil diameter, and whisking (8–23). Continued encoding of non-visual variables when visual stimuli are present could at least in part explain the trial-to-trial variability in cortical responses to repeated presentation of identical sensory stimuli (24).

The influence of trial-to-trial variability on stimu-

lus encoding depends on its population-level structure. Variability that is independent across cells – such as the Poisson-like variability produced in balanced recurrent networks (25) – presents little impediment to information coding, as reliable signals can still be extracted as weighted sums over a large enough population. In contrast, correlated variability has consequences that depend on the form of the correlations. If correlated variability mimics differences in responses to different stimuli, it can compromise stimulus encoding (26, 27). Conversely, correlated variability in dimensions orthogonal to those encoding stimuli has no adverse impact on coding (28), and might instead reflect encoding of signals other than visual inputs.

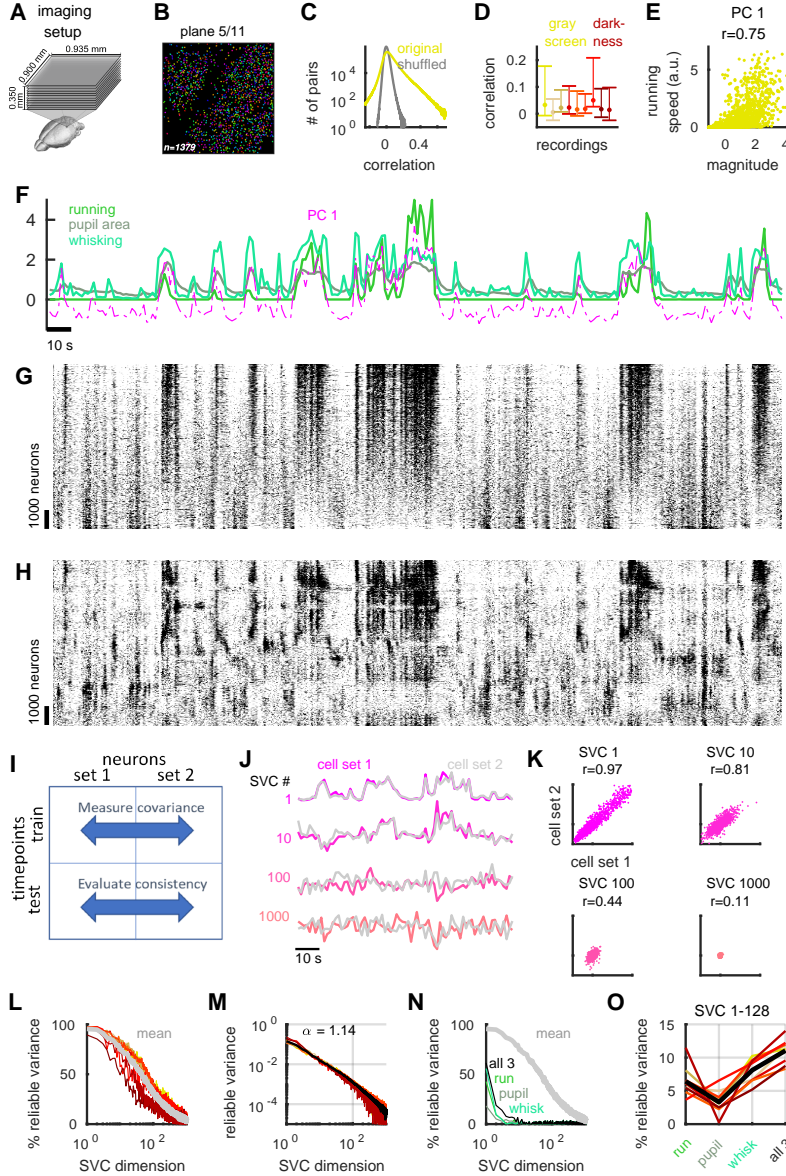


Figure 1: Structured ongoing population activity in visual cortex. (A,B) Two-photon calcium imaging of $\sim 10,000$ neurons in visual cortex using multi-plane resonance scanning of 11 planes spaced $35 \mu\text{m}$ apart. (C) Distribution of pairwise correlations in ongoing activity, computed in 1.2 second time bins (yellow). Gray: distribution of correlations after randomly time-shifting each cell's activity. (D) Distribution of pairwise correlations for each recording (showing 5th and 95th percentile). (E) First PC versus running speed in 1.2 s time bins. (F) Example timecourse of running speed (green), pupil area (gray), whisking (light green), first principal component of population activity (magenta dashed). (G) Neuronal activity, with neurons sorted vertically by 1st PC weighting, same time segment as F. (H) Same neurons as in G, sorted by a manifold embedding algorithm. (I) Shared Variance Component Analysis (SVCA) method for estimating reliable variance. (J) Example timecourses of SVCs from each cell set in the test epoch (1.2 s bins). (K) Same as J, plotted as scatter plot. Title is Pearson correlation between cell sets: an estimate of that dimension's reliable variance. (L) % of reliable variance for successive dimensions. (M) Reliable variance spectrum, power law decay of exponent 1.14. (N) % of each SVC's total variance that can be reliably predicted from arousal variables (colors as in E). (O) Percentage of total variance in first 128 dimensions explainable by arousal variables.

Spontaneous cortical activity reliably encodes a high-dimensional latent signal

To distinguish between these possibilities, we characterized the structure of neural activity and sensory variability in mouse visual cortex. We simultaneously recorded from $11,262 \pm 2,282$ (mean \pm s.d.) excitatory neurons, over nine sessions in six mice using 2-photon imaging of GCaMP6s in an 11-plane configuration (29) (Fig. 1A,B, Movie S1). These neurons' responses to classical grating stimuli revealed robust orientation tuning as expected in visual cortex (Fig. S1).

We began by analyzing spontaneous activity in mice free to run on an air-floating ball. Six of nine recordings were performed in darkness, but we did not observe differences between these recordings (shown in red on all plots) and recordings with gray screen (yellow on all plots). Mice spontaneously performed behaviors such as running, whisking, sniffing, and other facial movements, which we monitored with an infrared camera.

Ongoing population activity in visual cortex was highly structured (Fig. 1C-H). Correlations between neuron pairs were reliable (Fig. S2), and their spread was larger than would be expected by chance (Fig. 1C,D), suggesting structured activity (30). Fluctuations in the first principal component (PC) occurred over a timescale of many seconds (Fig. S3), and were coupled to running, whisking, and pupil diameter. These arousal-related variables correlated with each other (Fig. S4A,B), and together accounted for approximately 50% of the variance of the first neural PC (Fig. 1E, Fig. S4C). Correlation with the first PC was positive or negative in approximately similar numbers of neurons ($57\% \pm 6.7\%$ SE positive), indicating that two large sub-populations of neurons alternate their activity (Fig. 1F,G). The slowness of these fluctuations suggests a different underlying phenomenon to previously-studied up and down phases (5, 31–34), which alternate at a much faster timescale (~ 100 ms instead of multiple seconds) and correlate with most neurons positively. Indeed, up/down phases could not even have been detected in our recordings, which scanned the cortex every 400 ms.

Spontaneous activity had a high-dimensional structure, more complex than would be predicted by a single factor such as arousal. We sorted the raster diagram so that nearby neurons showed strong correlations (Figs. 1H, S5). Position on this continuum

bore little relation to actual distances in the imaged tissue (Fig. S6), suggesting that this activity was not organized topographically.

Despite its noisy appearance, spontaneous population activity reliably encoded a high-dimensional latent signal (Fig. 1I-K). We devised a method to identify dimensions of neural variance that are reliably determined by common underlying signals, termed Shared Variance Component Analysis (SVCA). We divided the recorded neurons into two spatially segregated sets, and divided the recorded timepoints into two equal halves (training and test; Fig. 1I). The training timepoints were used to find the dimensions in each cell set's activity that maximally covary with each other. These dimensions are termed Shared Variance Components (SVCs). Activity in the test timepoints was then projected onto each SVC (Fig. 1J), and the correlation between projections from the two cell sets (Fig. 1K) provided an estimate of the reliable variance in that SVC (see Methods and Appendix). The fraction of reliable variance in the first SVC was 97% (Fig. 1K,L), implying that only 3% of the variance along this dimension reflected independent noise. The reliable variance fraction of successive SVCs decreased slowly, with the 50th SVC at $\sim 50\%$ reliable variance, and the 512th at $\sim 9\%$ (Fig. 1L).

The magnitude of reliable spontaneous variance was distributed across dimensions according to a power law of exponent 1.14 (Fig. 1M). This value is larger than the power law exponents close to 1.0 seen for stimulus responses (35), but still indicates a high-dimensional signal. The first 128 SVCs together accounted for $86\% \pm 1\%$ SE of the complete population's reliable variance, and $67\% \pm 3\%$ SE of the total variance in these 128 dimensions was reliable. Arousal variables accounted for $11\% \pm 1\%$ SE of the total variance in these 128 components (16% of their reliable variance), and primarily correlated with the top SVCs (Fig. 1N,O). Thousands of neurons were required to reliably characterize activity in hundreds of dimensions, and the estimated reliability of higher SVCs increased with the number of neurons analyzed (Fig. S7A-E), suggesting that recordings of larger populations would identify yet more dimensions.

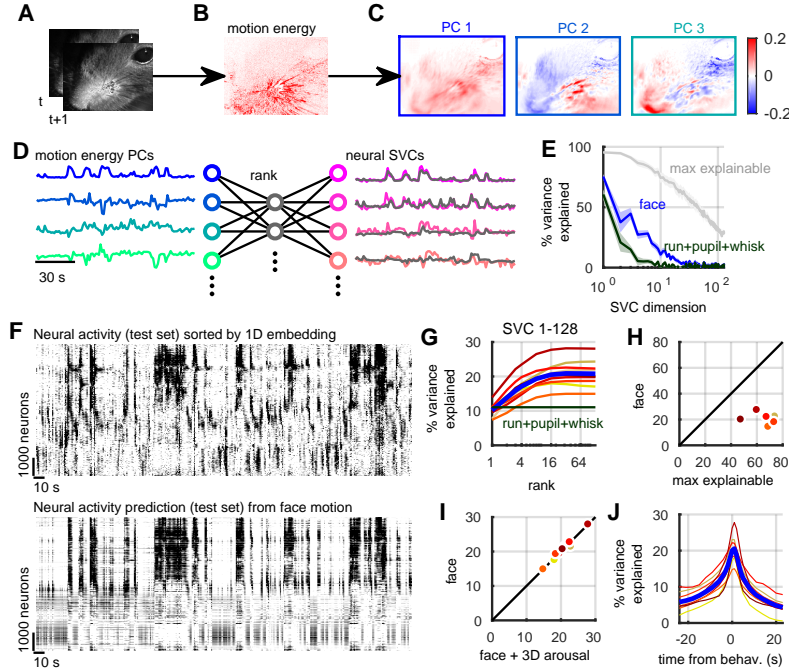


Figure 2: Multi-dimensional behavior predicts neural activity. (A) Frames from a video recording of a mouse's face. (B) Motion energy, computed as the absolute value of the difference of consecutive frames. (C) Spatial masks corresponding to the top three principal components (PCs) of the motion energy movie. (D) Schematic of reduced rank regression technique used to predict neural activity from motion energy PCs. (E) Cross-validated fraction of successive neural SVCs predictable from face motion (blue), together with fraction of variance predictable from running, pupil and whisking (green), and fraction of reliable variance (the maximum explainable; gray; cf. Fig. 1L). (F) Top: raster representation of ongoing neural activity in an example experiment, with neurons arranged vertically as in Fig. 1H so correlated cells are close together. Bottom: prediction of this activity from facial videography (predicted using separate training timepoints). (G) Percentage of the first 128 SVCs' total variance that can be predicted from facial information, as a function of number of facial dimensions used. (H) Prediction quality from multidimensional facial information, compared to the amount of reliable variance. (I) Adding explicit running, pupil and whisker information to facial features provides little improvement in neural prediction quality. (J) Prediction quality as a function of time lag used to predict neural activity from behavioral traces.

Ongoing neural activity encodes multidimensional behavioral information

Although arousal measures only accounted for a small fraction of the reliable variance of spontaneous population activity, it is possible that higher-dimensional measures of ongoing behavior could explain a larger fraction (Fig. 2A-C, Movie S2). We extracted a 1,000-dimensional summary of the motor actions visible on the mouse's face by applying principal component analysis to the spatial distribution of facial motion at each moment in time (36). The first PC captured motion anywhere on the mouse's face, and was strongly correlated with explicit arousal measures (Fig. S4B), while higher PCs distinguished different types of facial motion. We predicted neuronal population activity from this behavioral signal using reduced rank regression: for any N , we found the N dimensions of the video signal predicting the largest fraction of the reliable spontaneous variance (Fig. 2D).

This multidimensional behavior measure predicted approximately twice as much variance as the three arousal variables (Fig. 2D-J, Movie S3). To visualize how multidimensional behavior predicts ongoing population activity, we compared a raster representation of raw activity (vertically sorted as in Fig. 1H) to the prediction based on multidimensional videography (Fig. 2F, see Fig. S5 for all recordings). To quantify the quality of prediction, and the dimensionality of the behavioral signal encoded in V1, we focused on the first 128 SVCs (accounting for 86% of the population's reliable variance). The best one-dimensional predictor extracted from the facial motion movie captured the same amount of variance as the best one-dimensional combination of whisking, running, and pupil (Fig. 2G). Prediction quality continued to increase with up to 16 dimensions of videographic information (and beyond, in some recordings), suggesting that visual cortex encodes at least 16 dimensions of motor information. These dimensions together accounted for $21\% \pm 1\%$ SE of the total population variance ($31\% \pm 3\%$ of the reliable variance; Fig. 2H), substantially more than the three-dimensional model of neural activity using running, pupil area and whisking ($11\% \pm 1\%$ SE of the total variance, $17\% \pm 1\%$ SE of the reliable variance). Moreover, adding these three explicit predictors to the video signal increased the explained variance by less than 1% (Fig. 2I), even though the running signal provided information not derived from video. A neuron's

predictability from behavior was not related to its cortical location (Fig. S8). The timescale with which neural activity could be predicted from facial behavior was ~ 1 s (Figs. 2J, S7H), reflecting the slow nature of these behavioral fluctuations.

Behaviorally-related activity is spread across the brain

Patterns of spontaneous V1 activity were a reflection of activity patterns spread across the brain (Fig. 3A-E). To show this, we used 8 Neuropixels probes (37) to simultaneously record from frontal, sensorimotor, visual and retrosplenial cortex, hippocampus, striatum, thalamus, and midbrain (Fig. 3A,B). The mice were awake and free to rotate a wheel with their front paws. From recordings in three mice, we extracted 2,296, 2,668 and 1,462 units stable across ~ 1 hour of ongoing activity, and binned neural activity into 1.2 s bins, as for the imaging data.

Neurons correlated most strongly with others in the same area, but also correlated with neurons in other areas, suggesting non-localized patterns of neural activity (Fig. 3C). All areas contained neurons positively and negatively correlated with the arousal-related top facial motion PC, although neurons in thalamus showed predominantly positive correlations (Fig. 3D, $p < 10^{-8}$ two-sided Wilcoxon sign-rank test). Sorting the neurons by correlation revealed a rich activity structure (Fig. 3E). All brain areas contained a sampling of neurons from the entire continuum (Fig. 3E, right), suggesting that a multidimensional structure of ongoing activity is distributed throughout the brain. This spontaneous activity spanned at least 128 dimensions, with 35% of the variance of individual neurons reliably predictable from population activity (Fig. S9).

Similar to visual cortical activity, the activity of brainwide populations was partially predictable from facial videography (Fig. 3F-H). Predictability of brain-wide activity again saturated around 16 behavioral dimensions, which predicted on average across areas 21.9% of the total variance (40% of the estimated maximum possible) (Fig. 3F). The amount of behavioral modulation differed between brain regions, with neurons in thalamus predicted best (35.7% of total variance, 59% of estimated maximum). The timescale of videographic prediction was again broad: neural activity was best predicted from instantaneous

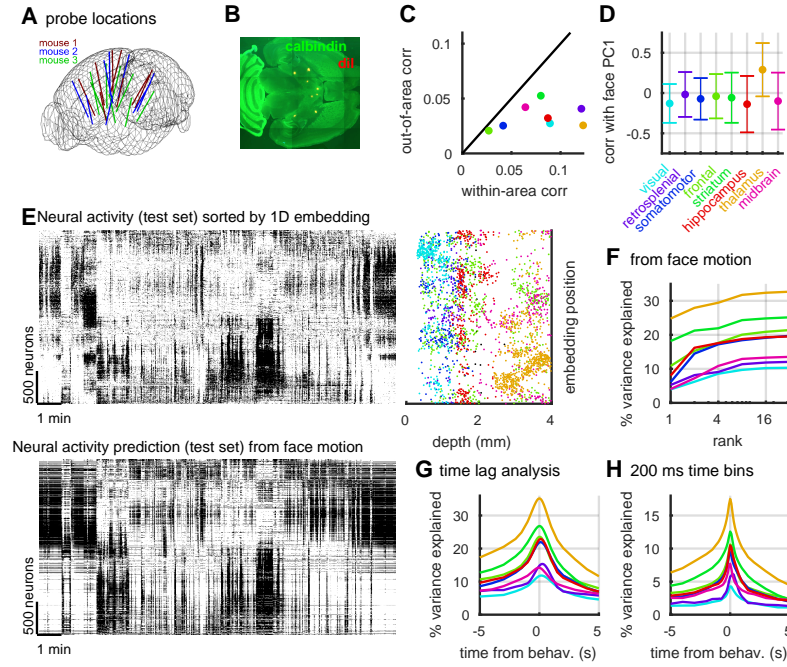


Figure 3: Behaviorally-related activity across the forebrain in simultaneous recordings with 8 Neuropixels probes. (A) Reconstructed probe locations of recordings in three mice. (B) Example histology slice showing orthogonal penetrations of 8 electrode tracks through a calbindin-counterstained horizontal section. (C) Comparison of mean correlation between cell pairs in a single area, with mean correlation between pairs with one cell in that area and the other elsewhere. Each dot represents the mean over all cell pairs from all recordings, color coded as in panel D. (D) Mean correlation of cells in each brain region with first principal component of facial motion. Error bars: standard deviation. (E) Top: Raster representation of ongoing population activity for an example experiment, sorted vertically so nearby neurons have correlated ongoing activity. Bottom: prediction of this activity from facial videography. Right: Anatomical location of neurons along this vertical continuum. Each point represents a cell, colored by brain area as in C,D, with x-axis showing the neuron's depth from brain surface. (F) Percentage of population activity explainable from orofacial behaviors as a function of dimensions of reduced rank regression. Each curve shows average prediction quality for neurons in a particular brain area. (G) Explained variance as a function of time lag between neural activity and behavioral traces. Each curve shows the average for a particular brain area. (H) Same as G in 200ms bins.

behavior, decaying slowly over time lags of multiple seconds (Figs. 3G-H, S10), with a full-width at half-max of 2.5 ± 0.4 s (mean \pm SE). Neural population activity showed coherent structure at timescales faster than this behavioral correlation (280 ± 43 ms, mean \pm SE) (Fig. S10). The fast-timescale structure modulated nearly all neurons in the same direction, leading to rapid fluctuations in the total population rate (“up and down phases”); by contrast, the structure seen at lower time scales was dominated by alternation in the activity of different neuronal populations, and steadier total activity (Figs. S10, S11, S12 (38)).

Stimulus-evoked and ongoing activity overlap along one dimension

We next asked how ongoing activity and behavioral information relates to sensory responses (Fig. 4A-B). We thus interspersed blocks of visual stimulation (flashed natural images, presented 1 per second on average) with extended periods of spontaneous activity (gray screen), while imaging visual cortical population activity (Fig. 4A). During stimulus presentation, the mice continued to exhibit the same behaviors as in darkness, resulting in a similar distribution of facial motion components (Fig. 4B).

There were not separate sets of neurons encoding stimuli and behavioral variables; instead, representations of sensory and behavioral information were mixed together in the same cell population. The fractions of each neuron’s variance explained by stimuli and by behavior were only slightly negatively correlated (Fig. S13; $r = -0.18$, $p < 0.01$ Spearman’s rank correlation), and neurons with similar stimulus responses did not have more similar behavioral correlates (Fig. S13; $r = -0.005$, $p > 0.05$).

The subspaces encoding sensory and behavior information overlapped in only one dimension (Fig. 4C-E). The space that encoded behavioral variables contained 11% of the total stimulus-related variance, 96% of which was contained in a single dimension (Fig. 4C) with largely positive weights onto all neurons (85% positive weights, Fig. 4D). Similarly, the space of ongoing activity, defined by the top 128 principal components of spontaneous firing, contained 23% of the total stimulus-related variance, 86% of which was contained in one dimension (85% positive weights). Thus, overlap in the spaces encoding sensory and behavioral variables arises primarily because

both can change the mean firing rate of the population: the precise patterns of increases and decreases about this change in mean were essentially orthogonal (Fig. 4E). Analysis of electrophysiological recordings confirmed that the relationship between stimulus-driven and spontaneous activity was dominated by a single shared dimension: the correlation between spontaneous and signal correlations was greatly reduced after projecting out this one-dimensional activity (Fig. S14).

Stimulus decoding analysis further confirmed that information about sensory stimuli was concentrated in the stimulus-only subspace. To show this, we trained a linear classifier to identify which stimulus had been presented, from activity in different 32-dimensional neural subspaces. Decoding from the stimulus space yielded a cross-validated error rate of 10.1 ± 4.0 %; activity in the spontaneous- or behavior-only spaces yielded errors of 53.1 ± 6.4 % and 56.8 ± 6.7 %, no better than randomly-chosen dimensions (Fig. 4F).

To visualize how the V1 population integrated sensory and behavior-related activity, we examined the projection of this activity onto three orthogonal subspaces: a multidimensional subspace encoding only sensory information (stimulus-only); a multidimensional subspace encoding only behavioral information (behavior-only); and the one-dimensional subspace encoding both (stimulus-behavior shared dimension) (Fig. 4G; Fig. S15). During gray-screen periods there was no activity in the stimulus-only subspace, but when the stimuli appeared this subspace became very active. By contrast, activity in the behavior-only subspace was present prior to stimulus presentation, and continued unchanged when the stimulus appeared. The one-dimensional shared subspace showed an intermediate pattern: activity in this subspace was weak prior to stimulus onset, and increased when stimuli were presented. Similar results were seen for the spontaneous-only and stimulus-spontaneous spaces (Fig. 4G, lower panels). Across all experiments, variance in the stimulus-only subspace was 119 ± 81 SE times larger during stimulus presentation than during spontaneous epochs (Fig. 4H), while activity in the shared subspace was 19 ± 12 SE times larger; activity in the face-only and spontaneous-only subspaces was only modestly increased by sensory stimulation (1.4 ± 0.13 SE and 1.7 ± 0.2 SE times larger, respectively).

Trial-to-trial variability in population responses

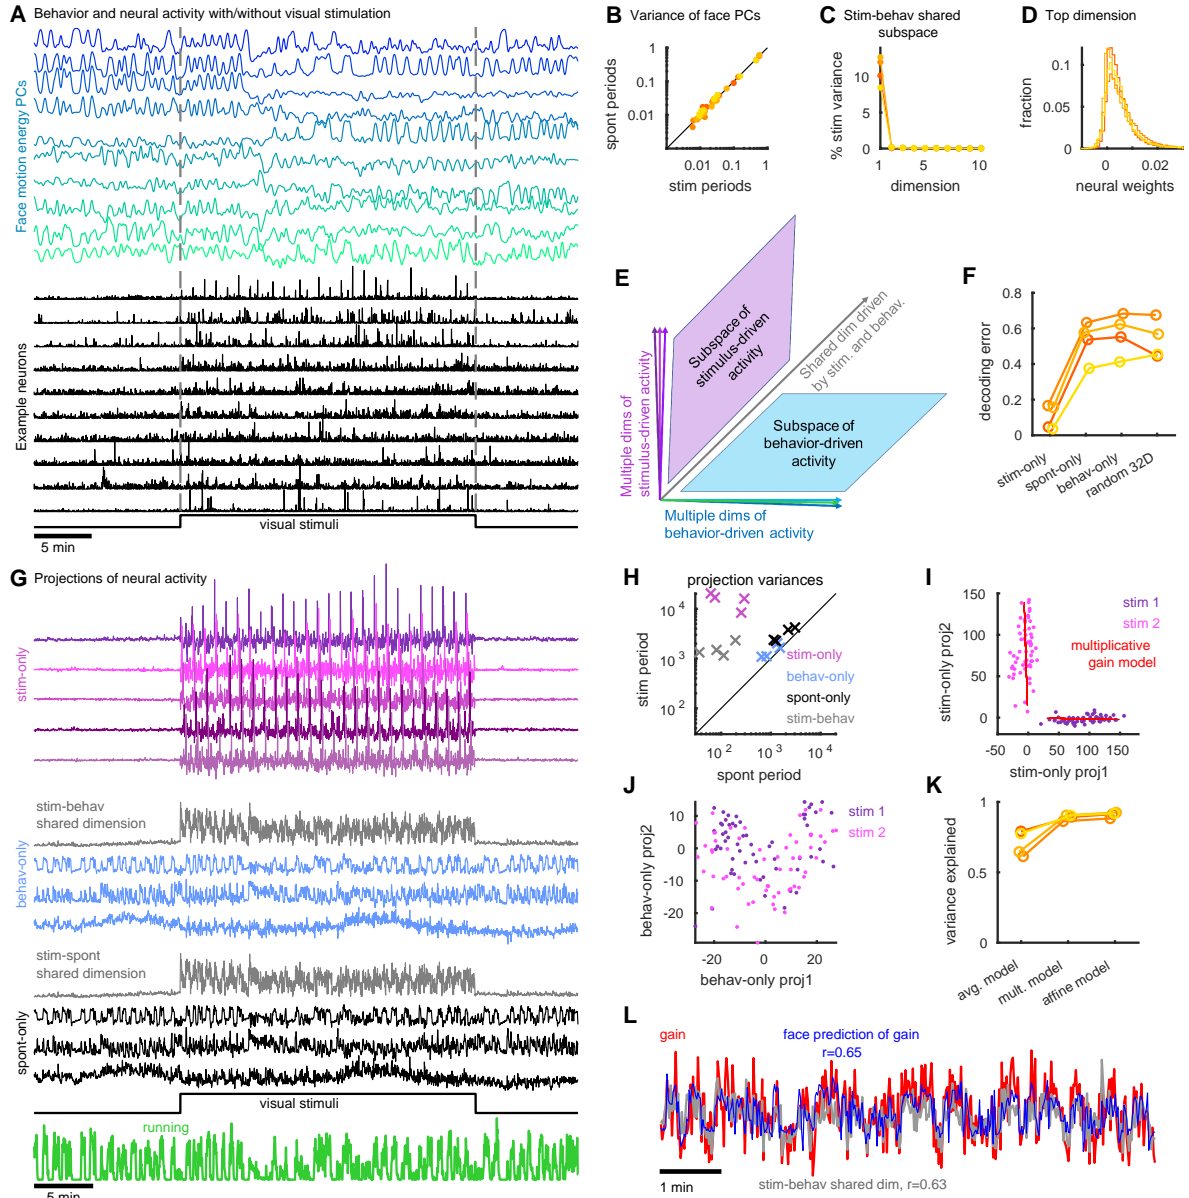


Figure 4: Neural subspaces encoding stimuli and spontaneous/behavioral variables overlap along one dimension. (A) Principal components of facial motion energy (top) and firing of ten example V1 neurons (bottom). (B) Comparison of face motion energy for each PC during stimulus presentation and spontaneous periods. Color represents recording identity. (C) The percentage of stimulus-related variance in each dimension of the shared subspace between stimulus- and behavior-driven activity. (D) Distribution of cells' weights on the single dimension of overlap between stimulus and behavior subspaces. (E) Schematic: stimulus- and behavior-driven subspaces are orthogonal, while a single dimension (gray; characterized in panels C,D) is shared. (F) Stimulus decoding analysis for 32 natural image stimuli from 32 dimensions of activity in the stimulus-only, behavior-only, and spontaneous-only subspaces, together with randomly-chosen 32-dimensional subspaces. Y-axis shows fraction of stimuli that were identified incorrectly. (G) Example of neural population activity projected onto these subspaces. (H) Amount of variance of each of the projections illustrated in G, during stimulus presentation and spontaneous periods. Each point represents summed variances of the dimensions in the subspace corresponding to the symbol color, for a single experiment. (I) Projection of neural responses to two example stimuli into two dimensions of the stimulus-only subspace. Each dot is a different stimulus response. Red is the fit of each stimulus response using the multiplicative gain model. (J) Same as I for the behavior-only subspace. (K) Fraction of variance in the stimulus-only subspace explained by: constant response on each trial of the same stimulus (avg. model); multiplicative gain that varies across trials (mult. model); and a model with both multiplicative and additive terms (affine model). (L) The multiplicative gain on each trial (red) and its prediction from the face motion PCs (blue).

to repeated stimulus presentations reflected a combination of multiplicative modulation in the stimulus-space, and additive modulation in orthogonal dimensions. To visualize how stimuli affected activity in these subspaces, we plotted population responses to multiple repeats of two example stimuli (Fig. 4I-J). When projected into the stimulus-only space, the resulting clouds were tightly defined with no overlap (Fig. 4I), but in the behavior-only space, responses to the two stimuli were directly superimposed (Fig. 4J). Variability within the stimulus subspace consisted of changes in the length of the projected activity vectors between trials, resulting in narrowly elongated clouds of points (Fig. 4I), consistent with previous reports of multiplicative variability in stimulus responses (39–42). A model in which stimulus responses are multiplied by a trial-dependent factor accurately captured the data, accounting for $89\% \pm 0.1\%$ SE of the variance in the stimulus subspace (Fig. 4K). Furthermore, the multiplicative gain on each trial could be predicted from facial motion energy ($r = 0.61 \pm 0.02$ SE, cross-validated), and closely matched activity in the shared subspace ($r = 0.73 \pm 0.06$ SE, cross-validated; Fig. 4L). Although ongoing activity in the behavior-only subspace and visual responses in the stimulus-only subspace added independently, we did not observe additive variability within the stimulus-only space itself: an “affine” model also including an additive term did not significantly increase explained variance over the multiplicative model ($p > 0.05$, Wilcoxon rank-sum test). Similar results were obtained when analyzing responses to grating stimuli rather than natural images (Fig. S16).

Discussion

Ongoing population activity in visual cortex reliably encoded a latent signal of at least 100 linear dimensions, and possibly many more. The largest dimension correlated with arousal and modulated about half of the neurons positively and half negatively. At least 16 further dimensions were related to behaviors visible by facial videography, which were also encoded across the forebrain. The dimensions encoding motor variables overlapped with those encoding visual stimuli along only one dimension, which coherently increased or decreased the activity of the entire population. Activity in all other behavior-related dimensions continued unperturbed regardless of sensory stimulation. Trial-to-trial variability of sensory responses

comprised additive ongoing activity in the behavior subspace, and multiplicative modulation in the stimulus subspace, resolving apparently conflicting findings concerning the additive or multiplicative nature of cortical variability (39–42).

Our data are consistent with previous reports describing low-dimensional correlates of locomotion and arousal in visual cortex (8, 10–16, 33), but suggest these results were glimpses of a much larger set of non-visual variables encoded by ongoing activity patterns. 16 dimensions of facial motor activity can predict 31% of the reliable spontaneous variance. The remaining dimensions and variance might in part reflect motor activity not visible on the face or only decodable by more advanced methods (43–48), or they might reflect internal cognitive variables such as motivational drives.

Many studies have reported similarities between spontaneous activity and sensory responses (2–7). We also observed a similarity, but found it arose nearly exclusively from one dimension of neural activity. This dimension summarized the mean activity of all cells in the population, and variations along it reflected both spontaneous alternation of up and down phases and differences in mean population response between stimuli. These results therefore demonstrate that the statistical similarity of firing patterns during stimulation and ongoing activity need not imply recapitulation of previous sensory experiences, merely that cortex exhibits mean rate fluctuations with or without sensory inputs. While our results do not exclude that genuine recapitulation could occur in other behavioral circumstances, they reinforce the need for careful statistical analysis before drawing this conclusion: even a single dimension of common rate fluctuation is sufficient for some previously-applied statistical methods to report similar population activity (49).

The brain-wide representation of behavioral variables suggests that information encoded nearly anywhere in the forebrain is combined with behavioral state variables into a mixed representation. We found that these multidimensional signals are present both during ongoing activity and during passive viewing of a stimulus. Recent evidence indicates that they may also be present during a decision-making task (50). What benefit could this ubiquitous mixing of sensory and motor information provide? The most appropriate behavior for an animal to perform at any moment depends on the combination of available sensory data,

ongoing motor actions, and purely internal variables such as motivational drives. Integration of sensory inputs with motor actions must therefore occur somewhere in the nervous system. Our data indicate that it happens as early as primary sensory cortex. This is consistent with neuroanatomy: primary sensory cortex receives innervation both from neuromodulatory systems carrying state information, and from higher-order cortices which can encode fine-grained behavioral variables (9). This and other examples of pervasive whole-brain connectivity (51–54) may coordinate the brain-wide encoding of behavioral variables we have reported here.

Acknowledgements

We thank Michael Krumin for assistance with the two-photon microscopes, and Andy Peters for comments on the manuscript. **Funding:** This research was funded by Wellcome Trust Investigator grants (095668, 095669, 108726, 205093, 204915), a European Research Council award (694401) and a grant from the Simons Foundation (SCGB 325512). CS was funded by a four-year Gatsby Foundation PhD studentship. MC holds the GlaxoSmithKline / Fight for Sight Chair in Visual Neuroscience. CS and MP are now funded by HHMI Janelia. NS

was supported by postdoctoral fellowships from the Human Frontier Sciences Program and the Marie Curie Action of the EU (656528). **Author contributions:** Conceptualization, C.S., M.P., N.S., M.C. and K.D.H.; Methodology, C.S. and M.P.; Software, C.S. and M.P.; Investigation, C.S., M.P., N.S. and C.B.R.; Writing, C.S., M.P., N.S., M.C. and K.D.H.; Resources, M.C. and K.D.H. Funding acquisition, M.C. and K.D.H. **Data and materials availability:** All of the processed deconvolved calcium traces are available on figshare (55) (https://figshare.com/articles/Recordings_of_ten_thousand_neurons_in_visual_cortex_during_spontaneous_behaviors/6163622), together with the behavioral traces (face motion components, whisking, running and pupil). All the spike-sorting electrophysiology and the corresponding behavioral traces are also on figshare (56) (https://janelia.figshare.com/articles/Eight-probe_Neuropixels_recordings_during_spontaneous_behaviors/7739750). The code for the analysis and the figures is available on github (<https://github.com/MouseLand/stringer-pachitariu-et-al-2018a>)

Declaration of interests

The authors declare no competing interests.

References and Notes

References

1. D. L. Ringach, *Current Opinion in Neurobiology* **19**, 439 (2009).
2. K. Hoffman, B. McNaughton, *Science* **297**, 2070 (2002).
3. T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, A. Arieli, *Nature* **425**, 954 (2003).
4. F. Han, N. Caporale, Y. Dan, *Neuron* **60**, 321 (2008).
5. A. Luczak, P. Barthó, K. D. Harris, *Neuron* **62**, 413 (2009).
6. P. Berkes, G. Orbán, M. Lengyel, J. Fiser, *Science* **331**, 83 (2011).
7. J. O'Neill, B. Pleydell-Bouverie, D. Dupret, J. Csicsvari, *Trends in neurosciences* **33**, 220 (2010).
8. C. M. Niell, M. P. Stryker, *Neuron* **65**, 472 (2010).
9. L. Petreanu, *et al.*, *Nature* **489**, 299 (2012).
10. P.-O. Polack, J. Friedman, P. Golshani, *Nature Neuroscience* **16**, 1331 (2013).
11. M. J. McGinley, S. V. David, D. A. McCormick, *Neuron* **87**, 179 (2015).
12. M. Vinck, R. Batista-Brito, U. Knoblich, J. A. Cardin, *Neuron* **86**, 740 (2015).
13. J. M. Pakan, *et al.*, *eLife* **5**, e14985 (2016).
14. M. Dipoppa, *et al.*, *Neuron* **98**, 602 (2018).
15. J. Reimer, *et al.*, *Nature Communications* **7**, 13289 (2016).
16. A. B. Saleem, A. Ayaz, K. J. Jeffery, K. D. Harris, M. Carandini, *Nature Neuroscience* **16**, 1864 (2013).
17. L. J. Gentet, *et al.*, *Nature Neuroscience* **15**, 607 (2012).
18. S. P. Peron, J. Freeman, V. Iyer, C. Guo, K. Svoboda, *Neuron* **86**, 783 (2015).
19. D. M. Schneider, A. Nelson, R. Mooney, *Nature* **513**, 189 (2014).
20. R. S. Williamson, K. E. Hancock, B. G. Shinn-Cunningham, D. B. Polley, *Current Biology* **25**, 1885 (2015).
21. S. Erisken, *et al.*, *Current Biology* **24**, 2899 (2014).
22. M. M. Roth, *et al.*, *Nature Neuroscience* **19**, 299 (2016).
23. L. M. Fernandez, *et al.*, *Cerebral Cortex* **27**, 5444 (2016).
24. A. Arieli, A. Sterkin, A. Grinvald, A. Aertsen, *Science* **273**, 1868 (1996).
25. C. van Vreeswijk, H. Sompolinsky, *Science* **274**, 1724 (1996).
26. D. A. Ruff, M. R. Cohen, *Nature Neuroscience* **17**, 1591 (2014).
27. R. Moreno-Bote, *et al.*, *Nature Neuroscience* **17**, 1410 (2014).

28. J. S. Montijn, G. T. Meijer, C. S. Lansink, C. M. Pennartz, *Cell reports* **16**, 2486 (2016).
29. M. Pachitariu, *et al.*, *bioRxiv* **061507** (2016).
30. M. R. Cohen, A. Kohn, *Nature Neuroscience* **14**, 811 (2011).
31. M. Okun, *et al.*, *Nature* **521**, 511 (2015).
32. T. A. Engel, *et al.*, *Science* **354**, 1140 (2016).
33. C. Stringer, *et al.*, *eLife* **5**, e19695 (2016).
34. C. M. Constantinople, R. M. Bruno, *Neuron* **69**, 1061 (2011).
35. C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, K. D. Harris, *bioRxiv* **374090** (2018).
36. K. Powell, A. Mathy, I. Duguid, M. Häusser, *eLife* **4** (2015).
37. J. J. Jun, *et al.*, *Nature* **551**, 232 (2017).
38. M. Okun, N. A. Steinmetz, A. Lak, M. Dervinis, K. D. Harris, *bioRxiv* **395251** (2018).
39. R. L. T. Goris, J. A. Movshon, E. P. Simoncelli, *Nature Neuroscience* **17**, 858 (2014).
40. I.-C. Lin, M. Okun, M. Carandini, K. D. Harris, *Neuron* **87**, 644 (2015).
41. M. C. Dadarlat, M. P. Stryker, *Journal of Neuroscience* **37**, 3764 (2017).
42. A. J. Christensen, J. W. Pillow, *bioRxiv* **214007** (2017).
43. A. E. X. Brown, E. I. Yemini, L. J. Grundy, T. Jucikas, W. R. Schafer, *Proceedings of the National Academy of Sciences* **110**, 791 (2013).
44. A. S. Machado, D. M. Darmohray, J. Fayad, H. G. Marques, M. R. Carey, *eLife* **4** (2015).
45. A. A. Robie, *et al.*, *Cell* **170**, 393 (2017).
46. A. B. Wiltchko, *et al.*, *Neuron* **88**, 1121 (2015).
47. A. Kurnikova, J. D. Moore, S.-M. Liao, M. Deschênes, D. Kleinfeld, *Current Biology* **27**, 688 (2017).
48. A. Mathis, *et al.*, *Nature Neuroscience* **21**, 1281 (2018).
49. M. Okun, *et al.*, *Journal of Neuroscience* **32**, 17108 (2012).
50. S. Musall, M. T. Kaufman, S. Gluf, A. K. Churchland, *bioRxiv* **308288** (2018).
51. R. Gămănuț, *et al.*, *Neuron* **97**, 698 (2018).
52. J. A. Harris, *et al.*, *bioRxiv* **292961** (2018).
53. M. N. Economo, *et al.*, *eLife* **5** (2016).
54. Y. Han, *et al.*, *Nature* **556**, 51 (2018).
55. C. Stringer, M. Pachitariu, C. Reddy, M. Carandini, K. D. Harris, Recordings of ten thousand neurons in visual cortex during spontaneous behaviors (2018). Doi:10.25378/janelia.6163622.v4.
56. N. Steinmetz, M. Pachitariu, C. Stringer, M. Carandini, K. D. Harris, Eight-probe neuropixels recordings during spontaneous behaviors (2019). Doi:10.25378/janelia.7739750.

57. N. A. Steinmetz, *et al.*, *eNeuro* **4**, ENEURO (2017).
58. J. Deng, *et al.*, *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (IEEE, 2009), pp. 248–255.
59. T.-W. Chen, *et al.*, *Nature* **499**, 295 (2013).
60. M. Pachitariu, C. Stringer, K. D. Harris, *Journal of Neuroscience* (2018).
61. J. Friedrich, P. Zhou, L. Paninski, *PLOS Computational Biology* **13**, e1005423 (2017).
62. K. A. Ludwig, *et al.*, *Journal of Neurophysiology* **101**, 1679 (2009).
63. M. Pachitariu, N. A. Steinmetz, S. N. Kadir, M. Carandini, K. D. Harris, *Advances in Neural Information Processing Systems* (2016), pp. 4448–4456.
64. A. J. Izenman, *Journal of Multivariate Analysis* **5**, 248 (1975).
65. K. D. Harris, J. Csicsvari, H. Hirase, G. Dragoi, G. Buzsáki, *Nature* **424**, 552 (2003).
66. J. W. Pillow, *et al.*, *Nature* **454**, 995 (2008).
67. M. London, A. Roth, L. Beeren, M. Häusser, P. E. Latham, *Nature* **466**, 123 (2010).

Supplementary materials

Materials and Methods

Supplementary Text

Figs. S1 to S17

References (57-67)

Materials and methods

All experimental procedures were conducted according to the UK Animals Scientific Procedures Act (1986). Experiments were performed at University College London under personal and project licenses released by the Home Office following appropriate ethics review.

Preparation for two-photon calcium imaging in visual cortex

The imaging methods were similar to those described elsewhere (14). Briefly, surgeries were performed in seven adult mice (P35-P125) in a stereotaxic frame and under isoflurane anesthesia (5% for induction, 0.5-1% during the surgery). We used mice bred to express GCaMP6s in excitatory neurons (1 EMX-CRE x Ai94 GCaMP6s mouse, 3 CamKII x tetO GCaMP6s mice, and 1 Rasgrf-CRE x Ai94 GCaMP6s mouse), or mice bred to express tdTomato in GAD+ inhibitory neurons (2 GAD-Cre x tdTomato mice), allowing inhibitory neurons to be identified and excluded from further analysis. We did not observe epileptiform activity in any of these mice (57).

Before surgery, Rimadyl was administered as a systemic analgesic and lidocaine was administered locally at the surgery site. During the surgery we implanted a head-plate for later head-fixation, made a craniotomy of 3-4 mm in diameter with a cranial window implant for optical access, and, in Gad-Cre x tdTomato transgenics, performed virus injections with a beveled micropipette using a Nanoject II injector (Drummond Scientific Company, Broomall, PA 1) attached to a stereotaxic micromanipulator. We used AAV2/1-hSyn-GCaMP6s, acquired from University of Pennsylvania Viral Vector Core. Injections of 50-200 nl virus (1-3 x10¹² GC/ml) were targeted to monocular V1, 2.1-3.3 mm laterally and 3.5-4.0mm posteriorly from Bregma. To obtain large fields of view for imaging, we typically performed 4-8 injections at nearby locations, at multiple depths (~500 μ m and ~200 μ m). Rimadyl was then used as a post-operative analgesic for three days, delivered to the mice via their drinking water.

Data acquisition

We optically recorded neural activity in head-fixed awake mice implanted with 3-4 mm cranial windows centered over visual cortex, obtaining ~10,000 neurons in all recordings. The recordings were performed using multi-plane acquisition controlled by a resonance scanner, with planes spaced 30-35 μ m apart in depth. In 9 recordings, 10 or 12 planes were acquired simultaneously at a scan rate of 3 or 2.5 Hz. For 3 further recordings we used single-plane configuration, with a scan rate of 30Hz. The mice were free to run on an air-floating ball and were surrounded by three computer monitors. Spontaneous activity was recorded in darkness (monitors off), or with a gray background or presented visual stimuli on these monitors arranged at 90° angles to the left, front and right of the animal, so that the animal's head was approximately in the geometric center of the setup.

For each mouse imaged, we typically spent the first imaging day finding a suitable recording location, where the following three conditions held:

- the GCaMP signal was strong, in the sense that clear transients could be observed in large numbers of cells
- a large enough field of view could be obtained for 10,000 neuron recordings,
- the receptive fields of the neuropil were clearly spatially localized on our three monitors.

In animals for which there was a choice over multiple valid recording locations, we chose either: 1) a horizontally and vertically central retinotopic location or 2) a lateral retinotopic location, at 90° from the center, but still centered vertically. We did not observe differences related to retinotopic location (central or lateral), and thus pooled data across recording locations. We also did not observe significant differences between recordings obtained from GCaMP transgenic animals and from virus injections, nor between recordings made in complete darkness or with a gray screen. Thus, we pooled data over all conditions.

Recording protocol and visual stimuli

Spontaneous multiple-plane recordings were performed with either constant gray screen background (3 recordings) or monitors switched off (6 recordings). The 9 multi-plane recordings were of lengths 162 min, 155 min, 120 min, 117 min, 105 min, 90 min, 70 min, 70 min, 70 min. We also performed 3 gray-screen recordings in single plane mode (30Hz frame rate), of lengths 93 min, 87 min, 49 min.

Natural images responses were recorded by presenting 90-114 repetitions of 32 images manually selected from the ImageNet database (58), from ethologically-relevant categories: “birds”, “cat”, “flowers”, “hamster”, “holes”, “insects”, “mice”, “mushrooms”, “nests”, “pellets”, “snakes”, “wildcat”. We chose images in which less than 50% of the image was a uniform background, and containing a balance of low and high spatial frequencies. Images were flashed on all three screens for 0.5 sec, with a randomized gray-screen inter-stimulus interval between 0.3 sec and 1.1 sec. The images were shown in a random order during each repeat block. In three sessions a 30 s period of gray screen period occurred in between each repeat block; each of these sessions lasted at least 110 min in total. In one session, we first recorded a gray screen period (10 min), then interleaved longer blocks of stimuli (25 min) and gray screen (15 min) three times, for a total of 130 minutes of recording.

In four recordings we presented full-field drifting gratings on all 3 monitors, at one of 32 directions evenly spaced at 11° , with a spatial frequency of 0.05 cycles per degree and temporal frequency of 2Hz. Each direction was presented 70-128 times with a duration of 0.5 s, and a gray-screen inter-stimulus interval randomly distributed between 0.3 and 1.1 s. Gratings were presented in random order in blocks of 32, and in between blocks there was an additional gray screen period of length 30 s (4 recordings) or 0.5 sec (1 recording). Each recording session was at least 90 min long.

In six electrophysiological recordings (six mice) we targeted visual cortex with neuropixels probes while presenting natural images from the ImageNet database. Stimuli were presented for 400ms, with a random gray-screen inter-stimulus interval of 300-700ms. We presented 700 different stimuli, two times each.

Calcium imaging processing

The pre-processing of all raw calcium movie data was done using a toolbox we developed called Suite2p, using the default settings (29). The software is available at www.github.com/Mouseland/suite2p.

Briefly, Suite2p first aligns all frames of a calcium movie using 2D rigid registration based on regularized phase correlation, subpixel interpolation, and kriging. For all recordings we validated the inferred X and Y offset traces, to monitor any potential outlier frames that may have been incorrectly aligned. In a few recordings, a very small percentage ($<0.01\%$) of frames that had registration artifacts were removed and the extracted traces were replaced with interpolated values at those frames. In all recordings, the registered movie appeared well-aligned by visual inspection. Next, Suite2p performs automated cell detection and neuropil correction. To detect cells, Suite2p computes a low-dimensional decomposition of the data, which is used to run a clustering algorithm that finds regions of interest (ROIs) based on the correlation of the pixels inside them. The extraction of ROIs stops when the pixel correlations of new potential ROIs drops below a threshold parameter, which is set as a fraction of the correlation in the high SNR ROIs; thus, it does not require the number of clusters to be set a priori. A further step in the Suite2p GUI classifies ROIs as somatic or not. This classifier learns from user input, reaching 95% performance on this data (29), thus allowing us to skip the manual step altogether for most recordings. We note that the 5% errors might be attributable to human labelling error, or to dendritic signals from backpropagating APs, reflecting the spiking of deeper cells. Thus, there is little risk of ROIs measuring signals other than neuronal action potentials.

We took great care to compensate cellular fluorescence traces for the surrounding neuropil signal (59). This contamination is typically removed by subtracting out from the ROI signal a scaled-down version of the neuropil signal around the ROI; the scaling factor was set to 0.7 for all neurons. Importantly, for computing the neuropil signal, we excluded all pixels that Suite2p attributed to an ROI, whether somatic or dendritic. After neuropil subtraction, we subtracted a running baseline of the calcium traces with a sliding window of 60 seconds to remove long timescale drift in baseline, then applied non-negative spike deconvolution using the OASIS algorithm with a

fixed timescale of calcium indicator decay of 2 seconds (60, 61). To further ensure out-of-focus fluorescence could not contribute to our results, we excluded neurons whose signal might span two planes by excluding neurons in sequential planes that had a greater than a 0.6 correlation (in 1.2 second bins) with each other, and whose centers were within 5 μm of each other in XY.

In addition, we ensured the cell sets used for reliable variance estimation (Fig. 1I) were spatially non-overlapping: we segregated the field of view into 16 strips in XY (encompassing all Z) of width 60 μm , and put cells from the odd strips in one group and the cells in the even strips in the other group. This ensured that no cells from different groups were at the same XY position but at a different depth. For peer prediction analyses (Fig. S9), we excluded all peer cells within 70 μm of the target (Euclidean distance in three-dimensional space).

Facial videography

Infrared LEDs (850nm) were pointed at the face of the mouse to enable infrared video acquisition in darkness. The videos were acquired at 30Hz using a camera with a zoom lens and an infrared filter (850nm, 50nm cutoff). The wavelength of 850nm was chosen to avoid the 970nm wavelength of the 2-photon laser, while remaining outside the visual detection range of the mice.

Running speed was not monitored videographically, but rather by optical mice placed orthogonally to the air floating ball on which the mouse stood.

Automated extraction of orofacial behaviors of mice

We developed a toolbox with a GUI for videographic processing of orofacial movements of mice. The software is termed FaceMap, and is available at www.github.com/MouseLand/FaceMap. The processing time taken by the software scales linearly with the number of frames, and runs 4x faster than real-time on 30 Hz videos.

Motion processing of regions of interest

To extract defined behavioral variables (e.g. pupil diameter, whisking), we used a graphical user interface which allows manual selection of face areas. The user can choose any region of the frame in which to compute the total absolute motion energy, the SVDs of the absolute motion energy or the SVDs of the raw frames.

The absolute motion energy at each time T is computed as the absolute value of the difference between consecutive frames, resulting in a matrix \vec{M} of size $N_{\text{pixels}} \times N_{\text{timepoints}}$. The whisker signal used in the current study (Fig. 1F) was defined to be the total motion energy summed over all pixels in a manually-defined region covering the whisker pad.

SVD computation for large matrices

To extract a high-dimensional representation of the facial signal, the toolbox applies singular value decomposition (SVD) to the raw movie, the motion energy movie, or both. The computation is identical in both cases.

The movie matrices are too large to decompose in their raw form. To compute their SVD, we first split the movie \vec{M} into temporal segments \vec{M}_i of length ~ 1 minute, and compute the SVD of each segment individually. Since the number of pixels is very large (> 1 million), we compute the SVD of each movie segment by computing the top 200 eigenvectors \vec{V}_i of its time by time covariance matrix. We then compute the spatial projections of the segment onto these components, $\vec{U}_i = \vec{M}_i \vec{V}_i$. Each matrix \vec{U}_i consists of the left singular vectors of \vec{M}_i , scaled by the singular values and is thus a 200-dimensional summary of the segment \vec{M}_i , related via an orthogonal projection. To estimate the SVD of the entire movie, we concatenate the \vec{U}_i for all segments of the movie, and re-compute the SVD: $[\vec{U}_1 \dots \vec{U}_n] = \vec{U} \vec{S} \vec{V}^\top$. The matrix \vec{U} represents the spatial components of the full movie, and we project the the movie onto the top 1000 components of it, to obtain their temporal profiles: $\vec{W}_{\text{motion}} = \vec{U}^\top \vec{M}$.

Pupil processing

To compute pupil area, the user first defines a region of interest using the FaceMap interface. The minimum value in this region is subtracted from all pixels for robustness across illumination changes. The darkest pixels in this region, identified by a user-selected threshold, correspond to the pupil. We estimate the pupil center as the center of mass of these dark pixels: $\vec{x} = \sum_{\vec{x}} \vec{x} R(\vec{x}) / \sum_{\vec{x}} R(\vec{x})$, where \vec{x} is the two-dimensional pixel location, $R(\vec{x})$ is that pixel's darkness level relative to the threshold, and the sum runs over all pixels \vec{x} darker than the threshold. We compute the covariance of a 2D Gaussian fit to the region of interest: $\sum_{\vec{x}} (\vec{x} - \vec{x})(\vec{x} - \vec{x})^T / N_{\vec{x}}$, where the sum runs over all pixels darker than the threshold and $N_{\vec{x}}$ is the number of such pixels. For robustness, this process is iterated 4 times after re-selecting only pixels that are 2 standard deviations away from the center, and recomputing the Gaussian covariance fit. The final result is an outline of the pupil defined by an ellipse 2 standard deviations from the center of mass.

Neuropixels recordings

Neuropixels electrode arrays (37) were used to record extracellularly from neurons in seven mice. In three of these mice, eight neuropixels probes were used to target multiple brain areas. The three mice were: 73 days old, male, wild-type (mouse 1); 113 days old, female, TetO-GCaMP6s;Camk2a-tTa (mouse 2); 99 days old, male, Ai32;Pvalb-Cre (mouse 3). In four of these mice, 1-4 neuropixels probes were used to target visual cortex. These mice were between 8 and 24 weeks old at the time of recording, and were of either gender. The genotypes of these mice were Slc17a7-Cre;Ai95, Snap25-GCaMP6s, Ai32;Pvalb-Cre, or Emx1-Cre;CaMKIIa-tTA;Ai94.

In all cases, a brief (<1 hour) surgery to implant a steel headplate and 3D-printed plastic recording chamber (~12mm diameter) was first performed. Following recovery, mice were acclimated to head-fixation in the recording setup. During head-fixation, mice were seated on a plastic apparatus with forepaws on a rotating rubber wheel. Three computer screens were positioned around the mouse at right angles. On the day of recording, mice were again briefly anesthetized with isoflurane while two to eight small craniotomies were made with a dental drill. After several hours of recovery, mice were head-fixed in the setup. Probes had a silver wire soldered onto the reference pad and shorted to ground; these reference wires were connected to a Ag/AgCl wire positioned on the skull. The craniotomies as well as the wire were covered with saline-based agar, which was covered with silicone oil to prevent drying. Each probe was mounted on a rod held by an electronically positionable micromanipulator (uMP-4, Sensapex Inc.) and was advanced through the agar and through the dura. Once electrodes punctured dura, they were advanced slowly (~10 μ m/sec) to their final depth (4 or 5 mm deep). Electrodes were allowed to settle for approximately 15 minutes before starting recording. Recordings were made in external reference mode with LFP gain=250 and AP gain=500, using SpikeGLX software. The mice were in a light-isolated enclosure and, during the spontaneous part of the recording, the computer screens were black. Data were preprocessed by re-referencing to the common median across all channels (62).

Spike sorting the Neuropixels data

We spike sorted the data using a modification of Kilosort (63), termed Kilosort2, that tracks drifting clusters (Fig. S17). Code is available at www.github.com/MouseLand/Kilosort2. Without the modifications, the original Kilosort and similar algorithms can split clusters according to drift of the electrode, which would confound our behavioral-related analyses. Kilosort2 tracks neurons across drift levels and for longer periods of time (~1 hour in our case). In addition, Kilosort2 performed automated splits and merges similar to what a human curator would do based on spike waveform similarity, the bimodality of the distribution of waveform features, and the spike auto- and cross-correlograms.

The final single units used were from several cortical areas (visual: 709, sensorimotor: 371, frontal: 842, retrosplenial: 122), hippocampal formation (753), striatum (310), thalamus (2743), midbrain (570).

Correlations

Pairwise correlations were computed after binning activity at 1.2-1.3 s (3 or 4 frames respectively for 12 and 10 plane recordings; 1.2 s bins for Neuropixels recordings). To compute shuffled correlations (Figure 1C), we circularly shifted each neuron’s activity in time by a random number of bins (at least ± 1000), and correlated all the shifted traces with all the original traces.

Arranging rasters by correlation

To visualize high-dimensional structure in raw data, raster plots were sorted vertically along a 1d continuum so that nearby neurons were most correlated. To do this, the binned activity of each neuron was first z-scored, and electrode data was high-pass filtered (100 s Gaussian kernel; this was not necessary for 2p data as traces had already been high-passed in preprocessing). Neurons were sorted using a generalization of scaled k-means clustering, where the clusters are ordered along a 1D axis to have similar means to their nearby clusters. Neurons were initially ordered based on their weights onto the first principal component of population activity, and divided into 30 equal-sized clusters along this ordering. On each iteration, we computed the mean activity of each cluster, smoothed it across clusters with a Gaussian window, then reassigned each neuron to the cluster whose smoothed activity it was most correlated with. This process was repeated for 75 iterations. The width of the Gaussian smoothing window was held at 3 clusters for the first 25 iterations, then annealed to 1 over the following 50 iterations. On the final pass, we upsampled the neurons’ correlations with each cluster by a factor of 100 via kriging interpolation with a smoothing constant of 1 cluster. This allowed us to determine sub-integer assignments of neurons to clusters, resulting in a continuous distribution of neurons along a 1D axis. The algorithm is available, implemented in Python and MATLAB at www.github.com/MouseLand/RasterMap. We ran the MATLAB version on the data here.

Although the electrode data was high-pass filtered to compute sorting, we display the original raw activity in Fig. 3E.

Shared Variance Component Analysis

The SVCA method gives an asymptotically unbiased lower-bound estimate for the amount of a neural population’s variance reliably encoding a latent signal. A mathematical proof of this is given in the appendix; here, we describe how the algorithm was implemented for the current study.

We first split the population into two spatially segregated populations. To do so, we divided the XY plane into 16 non-overlapping strips of width 60 μm , and assigned the neurons in the even strips to one group, and the neurons in the odd strips to the other group, regardless of the neuron’s depth. Thus, there did not exist neuron pairs in the two sets that had the same XY position but a different depth, avoiding a potential confound that a neuron could be predicted from its own out-of-focus fluorescence.

Neural population activity was binned at 1.2-1.3 s resolution (see above), and each neuron’s mean activity was subtracted from its firing trace. We divided the recording into training and test timepoints (alternating periods of 72 s each), thereby obtaining four neural activity matrices: \vec{F}_{train} , \vec{F}_{test} , \vec{G}_{train} , and \vec{G}_{test} of size $N_{\text{neurons}} \times N_{\text{timepoints}}$, where \vec{F} and \vec{G} represent activity of the two cell sets. We compute the covariance matrix between the two cell sets on the training timepoints as

$$\vec{C}_{\text{train}} = \vec{F}_{\text{train}} \vec{G}_{\text{train}}^{\top} / N_{\text{timepoints}}$$

We then compute the top 1024 left and right singular vectors of \vec{C} , yielding N_{neurons} -dimensional vectors \vec{u}_k and \vec{v}_k for $k = 1 \dots 1024$. These vectors are the shared variance components (SVCs) for each population. The amount of reliable variance in each SVC (Fig. 1M) is then estimated by the covariance of the SVC projections over the test samples:

$$\hat{S}_k = \vec{u}_k^{\top} \vec{F}_{\text{test}} \vec{G}_{\text{test}}^{\top} \vec{v}_k / N_{\text{timepoints}}$$

To obtain the fraction of reliable variance (Fig. 1L), we normalize this reliable variance by the arithmetic mean of the variances of the test set data for each cell set on the corresponding projections, $S_{k,\text{tot}} = (\vec{u}_k^\top \vec{F}_{\text{test}} \vec{F}_{\text{test}}^\top \vec{u}_k + \vec{v}_k^\top \vec{G}_{\text{test}} \vec{G}_{\text{test}}^\top \vec{v}_k) / 2$.

Predicting neural activity from behavioral variables

To estimate the fraction of neural variance that could be predicted from explicitly-computed arousal variables (Fig. 1N,O), we resampled their traces into the same 1.2-1.3 s bins as the neural data. The arousal variables (either single traces of running, whisking, pupil area or all three together) defined predictor matrices \vec{X}_{train} and \vec{X}_{test} for the training and test sets. We predicted the SVCs of neural activity $\vec{U}^\top \vec{F}_{\text{train}}$ and $\vec{V}^\top \vec{G}_{\text{train}}$ from the training-set behavior traces by unregularized multivariate linear regression, obtaining weight matrices \vec{A} and \vec{B} that minimized the squared errors $\|\vec{U}^\top \vec{F}_{\text{train}} - \vec{A} \vec{X}_{\text{train}}\|^2$ and $\|\vec{V}^\top \vec{G}_{\text{train}} - \vec{B} \vec{X}_{\text{train}}\|^2$. We then used these weight matrices to predict activity in the test set, and computed the covariance matrix of the residual error of each SVC:

$$S_{k,\text{res}} = \left(\vec{u}_k^\top \vec{F}_{\text{test}} - \vec{a}_k \vec{X}_{\text{test}} \right) \left(\vec{v}_k^\top \vec{G}_{\text{test}} - \vec{b}_k \vec{X}_{\text{test}} \right)^\top / N_{\text{samples}}$$

$S_{k,\text{res}}$ represents the amount of variance along SVC k that cannot be predicted by the behavioral traces, and $(\hat{S}_k - S_{k,\text{res}}) / \hat{S}_k$ represents the fraction of reliable variance that can be so predicted. To compute the fraction of total variance explainable by behavioral traces (Fig. 1N,O; Fig. 2E,G-J), we normalize instead by the total test-set variance: $(\hat{S}_k - S_{k,\text{res}}) / S_{k,\text{tot}}$.

To predict the fraction of neural variance that could be predicted from unsupervised videographic analysis (Fig. 2), we took a similar approach, but computed the weight matrices \vec{A} and \vec{B} by reduced-rank regression. Reduced-rank regression is a form of regularized linear regression, with the prediction weights matrix restricted to a specific rank (64), reducing the number of parameters and making it more robust to overfitting. Fig. 2E shows the fraction of total variance in successive dimensions that can be predicted by rank-16 prediction, while Fig. 2G shows how the predicted fraction of variance in the first 128 dimensions depends on the rank of the predictor.

Peer prediction analysis

The shared variance component analysis described above – like a related algorithm for estimating reliable stimulus coding (35) – provides unbiased estimates, but requires thousands of simultaneously-recorded neurons per brain area. Because this many neurons were not available in our Neuropixels recordings, we turned to another method to estimate the reliable variance in these data. This method is an adaptation of the previously-described “peer prediction” method (65, 66). Peer prediction analysis attempts to predict each neuron individually from the other simultaneously recorded cells (the neuron’s “peers”). In contrast, SVCA finds the dimensions of activity in a large population that can be most reliably predicted from a held-out set of neurons. Because a substantial fraction of a single neuron’s variance arises from independent noise, which is averaged out when projecting onto the SVCA dimensions, peer prediction gives systematically lower values of variance explained than SVCA.

To apply peer prediction to our data, we again binned neural activity with 1.2-1.3 s resolution, and divided these timepoints into a training set and a test set, consisting of alternating blocks of duration 72 s. Each neuron took a turn as target for prediction from the activity of simultaneously recorded “peer” cells, defined to be any cells on all other probes and cells on the same probe greater than 5 sites away (40 μm) for neuropixels recordings; for 2p recordings we used all neurons greater than 70 μm from the cell in 3D distance, in order to avoid potential optical contamination from the target neuron. We denote peer cell activity in the training and test sets by $N_{\text{cells}} \times N_{\text{timepoints}}$ matrices \vec{F}_{train} and \vec{F}_{test} , respectively, and target cell activity in the training and sets as $1 \times N_{\text{timepoints}}$ vectors \vec{g}_{train} and \vec{g}_{test} . We first computed the singular value decomposition of peer cell activity on the training set: $\vec{F}_{\text{train}} = \vec{U} \vec{S} \vec{V}^\top$. We then predicted the target neuron activity by ridge regression from n singular value components of peer cell

activity, where n took values $n = 1, 2, 4, 8, 16, \dots, 512, 1024$. The prediction weights were thus

$$\vec{w}_n = \left[(\vec{g}_{\text{train}} \vec{V}_n \vec{S}_n) \left((\vec{V}_n \vec{S}_n)^\top (\vec{V}_n \vec{S}_n) + \lambda I \right)^{-1} \right] \vec{U}_n^\top$$

where \vec{U}_n , \vec{V}_n , \vec{S}_n are matrices containing the top n singular vectors. Then the prediction of the single neuron activity on the test set was $\hat{g}_{\text{test}}^n = \vec{w}_n \vec{F}_{\text{test}}$, and the fraction of variance explained was $1 - \|\vec{g}_{\text{test}} - \hat{g}_{\text{test}}^n\|^2 / \|\vec{g}_{\text{test}}\|^2$. We chose λ to be 10 by hand.

Subspaces of stimulus and behavioral activity

The stimulus subspace (Fig. 4) was defined as the space spanned by the trial-averaged responses of each of the 32 stimuli presented. We computed the $N_{\text{neurons}} \times N_{\text{stimuli}}$ matrix of trial-averaged responses \vec{R} from one third of the stimulus responses, saving the other two thirds of the stimulus responses for variance estimation.

The behavior subspace was defined via the reduced-rank regression prediction method described above. We performed the regression on one-half of the spontaneous activity, leaving the other half of the spontaneous activity for variance estimation. This method produces a weight matrix of size $N_{\text{neurons}} \times N_{\text{facePCs}}$, that factorizes as a product of two matrices of sizes $N_{\text{neurons}} \times r$ and $r \times N_{\text{facePCs}}$, where r is the rank of prediction. We defined the behavior space as the space spanned by the first 32 columns of the former matrix, and we define a $N_{\text{neurons}} \times 32$ matrix \vec{E}_B whose columns contain an orthonormal basis for this space.

To determine dimensions inside the behavioral subspace that contain stimulus information, we found the sequence of orthogonal directions \vec{e}_i maximizing the sum of squared projections (the power) of the trial-averaged stimulus responses \vec{R} :

$$\vec{e}_i = \text{argmax}(\|\vec{e}_i^\top \vec{R}\|^2) \text{ such that } \|\vec{e}_i\|^2 = 1, \vec{e}_i \in \text{Span}(\vec{E}_B), \text{ and } \vec{e}_i \text{ orthogonal to } \vec{e}_1 \dots \vec{e}_{i-1}$$

The solution to this maximization problem is given by the left singular vectors of $\vec{E}_B^\top \vec{R}$. To determine the amount of stimulus variance in shared dimension i (Fig. 4C), we projected test data onto \vec{e}_i and quantified the stimulus-related variance in this projection using the unbiased method of (35) on the remaining two thirds of stimulus responses. We call \vec{e}_1 the “shared stimulus-behavior” dimension, because it contains significant stimulus variance, as opposed to $\vec{e}_2, \vec{e}_3, \dots$ which contain very little (Fig. 4C). A histogram of the weights of \vec{e}_1 on all neurons are plotted in Fig. 4D.

The behavior-only subspace was defined by projecting out the shared dimension \vec{e}_1 from all columns of \vec{E}_B . The stimulus-only subspace was defined by projecting out from all rows of \vec{R} the top right singular vector of the matrix $\vec{E}_B^\top \vec{R}$. Timecourses of neural activity projected into these subspaces is plotted in Fig. 4F. To quantify the amount of variance in each subspace (stimulus-only, behavior-only, stimulus-behavior) (Fig. 4G), we computed the total projected variance as the sum of squared projection lengths along each axis of an orthonormal basis.

An identical analysis was used to define the stimulus-spontaneous shared dimension, and spontaneous-only subspace, by replacing the subspace of the top 32 behavioral components with the subspace of the top 128 principal components of activity computed in one half of the spontaneous period.

To account for trial-to-trial variability in the stimulus subspace, we fit a multiplicative gain model. A gain parameter g_t was fit for each trial t , and the activity of dimension n in response to the stimulus σ_t shown on this trial was modelled as $\bar{f}_{n,\sigma_t}(1 + g_t \alpha_n)$. Here, $\bar{f}_{n,\sigma}$ represents the mean activity of dimension n to stimulus σ , and α_n represents the susceptibility of this dimension to gain fluctuations. Note that the mean of g_t across trials from the same stimulus is 0 by definition. We used an alternating optimization method to obtain the best fit \vec{g} given $\vec{\alpha}$, then $\vec{\alpha}$ given \vec{g} , repeating for 100 iterations. We also evaluated an affine model, allowing both the gain and offset of each neuron’s responses to change on a trial by trial basis: $\bar{f}_{n,\sigma_t}(1 + g_t \alpha_n) + a_t \beta_n$. Here, a_t is an additive offset on trial t , with each neuron scaling this offset by a factor β_n . The vector β can therefore describe directions of additive variability inside the stimulus subspace.

Stimulus decoding analysis

To predict stimulus identity from population responses, we fit a linear multi-class support vector machine (SVM) model, using MATLAB’s “fitcecoc” function. This function fits $K*(K-1)/2$ binary linear SVM models, where K is the number of different stimuli. Stimulus repeats were divided into equal thirds: one for estimating the stimulus subspace, one for training the classifier, and one for evaluating performance. For the 32 natural image stimuli, decoding error was defined as the fraction of responses not assigned to the correct stimulus. For drifting grating responses, we fit two decoders: one to direction (32 possibilities) and one to orientation (16 possibilities). We defined the error as the number of degrees the prediction was off, averaged across all the test set responses.

Population activity was always decoded from a 32-dimensional subspace. For the spont-only and behavior-only subspaces, we used the top 32 dimensions. Random 32-dimensional subspaces were found by generating a matrix of standard Gaussian variates of size $\text{neurons} \times 32$, and normalizing each column so the sum of its squares was 1. For each experiment we computed 20 different random subspaces, computed the decoding error and SNR in these subspaces, and averaged the results.

Supplementary Text

Proof that SVCA is an asymptotically unbiased estimate of reliable variance

Here we describe the SVCA method in more detail, and prove that the reliable variance it estimates has an expectation that is a lower bound for the true value, with bias vanishing in the limit of many neurons and stimuli.

The key to the SVCA method is to split spontaneous neuronal activity into a reliable component – which coherently and deterministically encodes an unobservable state variable – and an unreliable component, which is random and independent between neurons. The unreliable component is sometimes referred to as “Poisson noise”, as it would occur in a model where neurons fired a number of spikes drawn from a Poisson distribution whose rate deterministically encoded the state variable. Such behavior can also occur in networks where strong and balanced excitation and inhibition lead to chaotic variations in spike counts even in deterministic simulations (25, 33, 67).

As we show in the main text, spontaneous activity encodes details of ongoing facial behavior – but the reliable component of spontaneous activity can also represent internal cognitive variables that cannot be directly measured. SVCA allows one to estimate the amount of neuronal population variance that coherently and reliably encodes an unobservable state variable, distinguishing it from the variance that is independently variable between neurons.

Formally, we consider the reliable component of spontaneous population activity to represent a state variable s , which is drawn from a set of possible states \mathcal{S} according to a probability distribution $\mathbb{P}(s)$. We model the activity of neuron n at timepoint t as the sum of a reliable response $\bar{f}(n, s_t)$ that depends deterministically on the state s_t at this timepoint, plus a noise term $\epsilon_{n,t}$ which is independent between neurons, and independent of the state s_t :

$$f_n(t) = \bar{f}(n, s_t) + \epsilon_{n,t}$$

To estimate the variance of reliable population activity, we divide the recorded neurons into two sets of size N and M . We denote the responses of the first set on trial t by $f_n(t)$ for $n = 1 \dots N$, and the second set by $g_m(t)$ with $m = 1 \dots M$, summarizing their activity in random vectors $\mathbf{f}(t) \in \mathbb{R}^N$ and $\mathbf{g}(t) \in \mathbb{R}^M$. Define the expected $N \times M$ cross-covariance matrix of these vectors as \mathbf{C} , with entries $C_{n,m} = \mathbb{E}_{s_t, \epsilon_t}[\bar{f}_n(t)g_m(t)]$ which is equal to $\mathbb{E}_s[\bar{f}(n, s)\bar{g}(m, s)]$ by independence of the noise terms. We write its singular value decomposition as $\mathbf{C} = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$, where the σ_i are non-increasing real numbers, and the \mathbf{u}_i and \mathbf{v}_i are orthonormal systems in \mathbb{R}^N and \mathbb{R}^M respectively.

To estimate the amount of reliable variance in an optimal k -dimensional subspace, we compute the first k left and right singular vectors $\hat{\mathbf{u}}_i$ and $\hat{\mathbf{v}}_i$ from the covariance matrix estimated from training timepoints, and estimate the k -dimensional reliable variance as

$$\hat{S}_{\leq k} = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^k \mathbf{f}(t) \cdot \hat{\mathbf{u}}_i \mathbf{g}(t) \cdot \hat{\mathbf{v}}_i$$

where the first sum runs over test set trials t . Because the training and test sets are independent, and because ϵ is independent of s ,

$$\mathbb{E}_{\text{test set}} [\hat{S}_{\leq k}] = \mathbb{E}_s \left[\sum_{i=1}^k \mathbf{f}(s) \cdot \hat{\mathbf{u}}_i \mathbf{g}(s) \cdot \hat{\mathbf{v}}_i \right] = \sum_{i=1}^k \hat{\mathbf{u}}_i^\top \mathbf{C} \hat{\mathbf{v}}_i$$

Now, because $\hat{\mathbf{u}}_i$ and $\hat{\mathbf{v}}_i$ are orthonormal systems, $\sum_{i=1}^k \hat{\mathbf{u}}_i^\top \mathbf{C} \hat{\mathbf{v}}_i \leq \sum_{i=1}^k \sigma_i$, with equality when $\hat{\mathbf{u}}_i$ and $\hat{\mathbf{v}}_i$ are the first k singular vectors of \mathbf{C} , which they approach when the number of training timepoints becomes large. Thus, the expectation of $\hat{S}_{\leq k}$ is a lower bound for $\sum_{i=1}^k \sigma_i$, that becomes accurate in the limit of a large number of training timepoints.

Finally, we show that as the number of recorded neurons increases, $\sum_{i=1}^k \sigma_i$ converges to the reliable variance of the entire population's top k dimensions. Recall that the singular values of \mathbf{C} are square roots of the eigenvalues of $\mathbf{C}\mathbf{C}^\top$. Concatenating the reliable rate vectors into $T \times N$ and $T \times M$ matrices \mathbf{F} and \mathbf{G} so $\mathbf{C} = \mathbf{F}^\top \mathbf{G}$, we see that σ_i^2 is the i^{th} eigenvalue of $\mathbf{F}^\top \mathbf{G} \mathbf{G}^\top \mathbf{F}$, which is also the i^{th} eigenvalue of $\mathbf{G} \mathbf{G}^\top \mathbf{F} \mathbf{F}^\top$ by cyclic permutability of eigenvalues. No matter how many neurons we record from, $\mathbf{F} \mathbf{F}^\top$ and $\mathbf{G} \mathbf{G}^\top$ are both $T \times T$ matrices, and we next consider their limit as the number of neurons becomes large.

To do so, we consider the neurons to be drawn from a hypothetically infinite population of neurons \mathcal{N} according to a probability distribution $\mathbb{P}(n)$. As the number N of neurons sampled from this distribution becomes large, the $(t_1, t_2)^{\text{th}}$ entry of the matrix $\frac{1}{N} \mathbf{F} \mathbf{F}^\top$ converges to

$$K_{t_1, t_2} = \mathbb{E}_n [\bar{f}(n, s_{t_1}) \bar{f}(n, s_{t_2})]$$

$\frac{1}{N} \mathbf{G} \mathbf{G}^\top$ converges to the same limit as the number of neurons M in the second set becomes large. Thus, the k first singular values of \mathbf{C} tend to the k first eigenvalues of \mathbf{K} , whose sum is the amount of reliable variance in the optimal k -dimensional subspace of population activity.




Figure S1: Robust orientation tuning in recorded V1 neurons. (A) Orientation tuning curve, showing population-averaged responses as a function of angle relative to each cell's preferred orientation (obtained from held-out responses). Stimuli were Gabor drifting gratings of spatial and temporal frequency of .05 cpd and 2 Hz. Each row shows responses a single recording session from a single mouse. (B) Orientation tuning curves of the 400 most tuned neurons in each experiment (as assessed by orientation selectivity index), arranged horizontally by preferred orientation (computed from held-out responses). (C-E) Orientation tuning curves of 400 neurons taken from the 75th, 50th, and 25th percentile of orientation selectivity, displayed similarly.




Figure S2: **Correlation matrices are structured and consistent across time.** (A,B) Pseudocolor representation of spontaneous correlation matrices for a subset of cells, computed independently from two halves of recording. (C) Scatter plot showing correlations of each cell pair in a single recording, for two independent halves. (D) Histogram showing Pearson correlation coefficient of pairwise correlations (as in C), for all recordings.

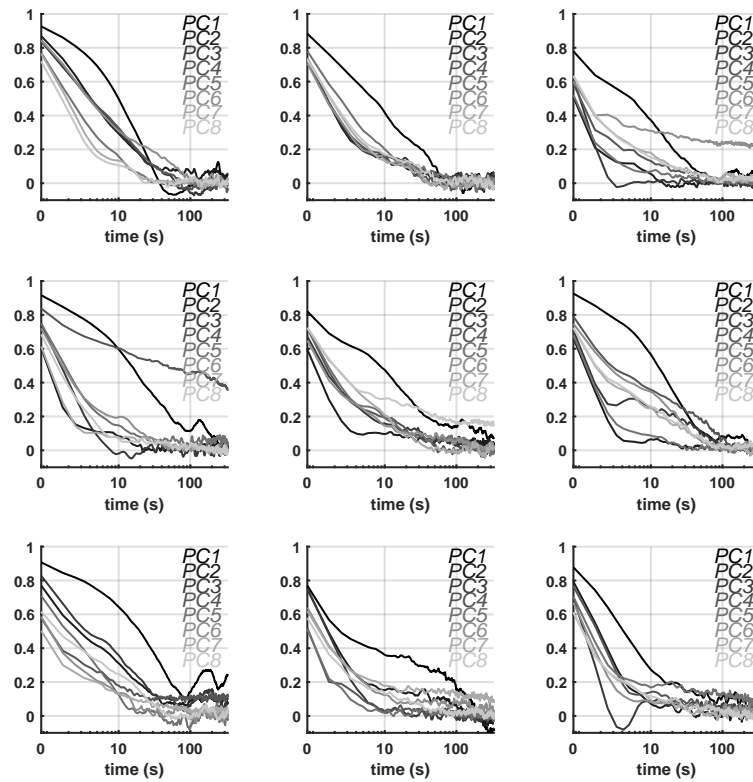


Figure S3: **Temporal autocorrelation of ongoing neural activity.** Each panel shows data for one recording; within each plot, each curve shows the temporal autocorrelation of a single principal component of ongoing population activity (1.2 second bins; log x-scale). Second plot on top row is for the example recording shown in Fig. 1.

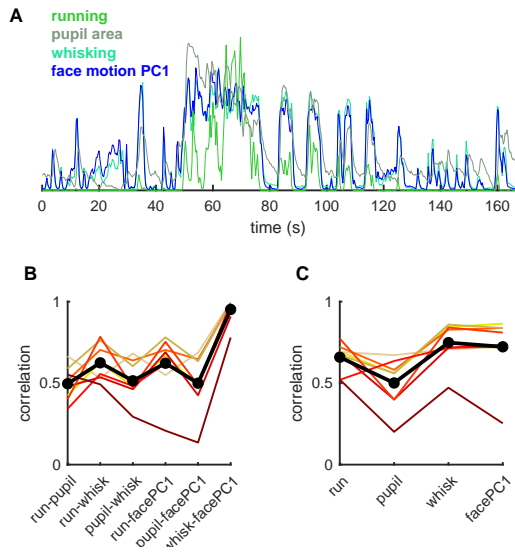


Figure S4: Correlations between arousal variables. (A) Time courses of running speed, pupil area, whisker motion energy, and first PC of motion energy, for a 160 s segment of an example experiment scaled to have the same maximum and minimum. (B) Pearson correlations for each pair of these variables. Each line denotes an individual experiment, black denotes average. (C) Correlation of the first neural PC with each behavioral variable.




Figure S5: Additional examples of predicting high-dimensional neural population activity from facial videography. For each panel, the top plot shows a raster diagrams sorted vertically to place correlated neurons together (cf. Fig. 1H). Bottom plots show predictions from facial videography (cf. Fig. 2F).

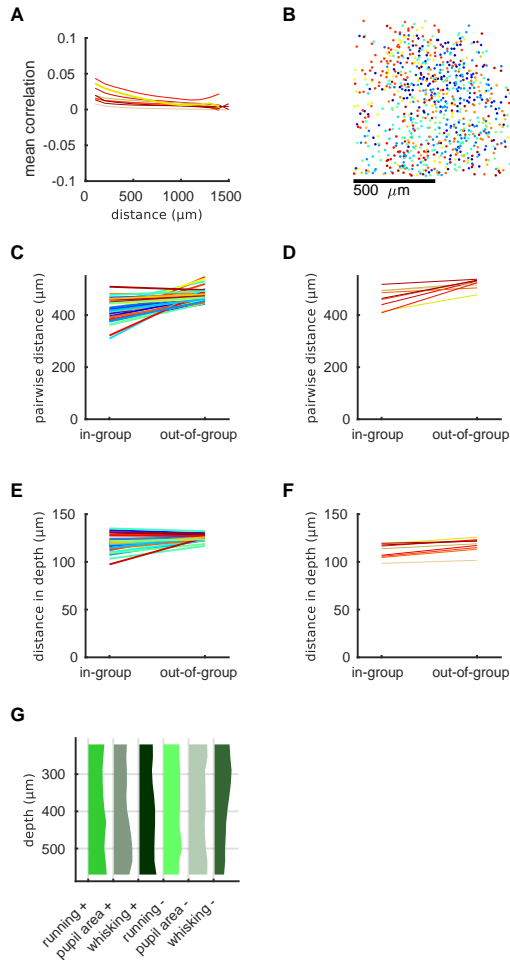


Figure S6: Spontaneous correlations bear little relationship to cortical distance. (A) Mean correlation over cell pairs, as a function of distance on the cortical surface. Each line represents a single experiment. (B) Neurons were assigned to 30 clusters according to a 1d manifold-embedding algorithm (Fig. 1H). Points show XY locations of cells in one example imaging plane, colored according to cluster. (C) Three-dimensional distances between neurons in the same or different clusters, for one experiment. Each line represents a cluster. (D) Average of (C) over all clusters. Each line represents an experiment. (E,F) Same as (C,D) but for one-dimensional pairwise distances in depth. (G) Distribution across cortical depth of six non-overlapping groups of neurons with different behavioral correlates. Neurons were grouped according to which arousal variable they best correlated with and the sign of this correlation.

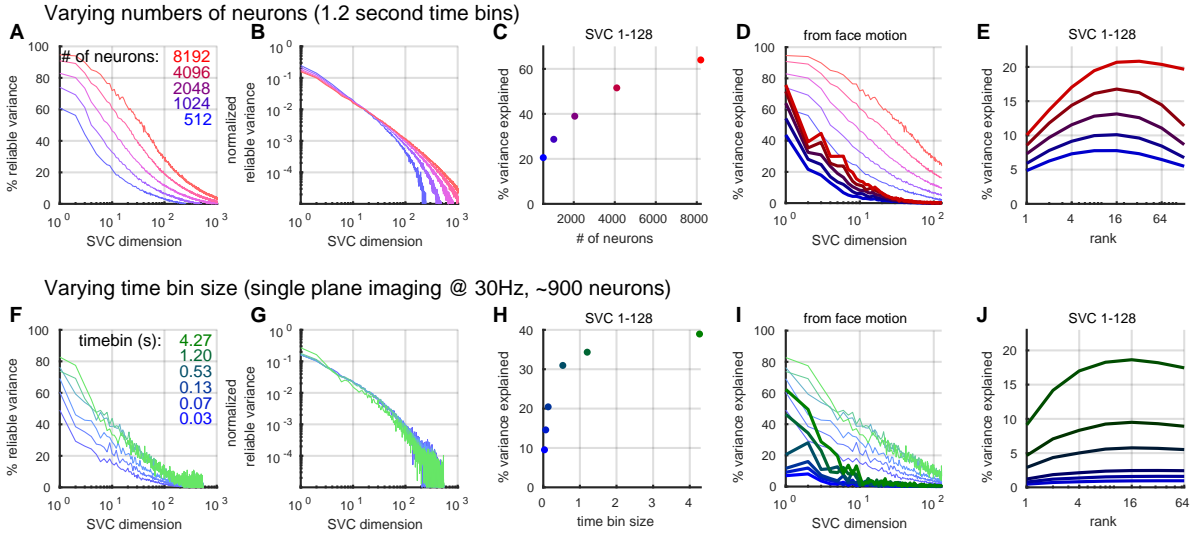


Figure S7: Varying the numbers of neurons and the time bin size. (A) SVCA analysis for randomly subsampled neural populations of different sizes. For each of the 9 recordings, we analyzed 10 random subsets of the recorded neurons, of sizes indicated by the color scheme. Each curve shows the percentage of neural variance explained as a function of SVC dimension (cf. Fig. 1L), averaged over all recordings. (B) Reliable variance spectrum (cf. Fig. 1M) as a function of subsampled population size, colors as before. (C) Percentage of reliable variance estimated in the first 128 SVCs, as a function of population size. (D) Prediction of SVCs by facial videography (cf. Fig. 2E), as a function of subsampled population size. Darker thick lines represent the prediction from face motion for each subset of neurons, and lighter thin lines show the fraction of reliable variance as in panel A. (E) Percentage of the first 128 SVCs' total variance that can be predicted from facial videography, as a function of number of facial dimensions used (cf. Fig. 2G) Colors indicate subsample size as in previous panels. (F-J) Same as (A-E), but now the number of neurons is fixed, and the size of the time bin varies. This analysis was performed on three single-plane recordings, which allowed a faster scan rate of 30 Hz, but yielded a smaller numbers of neurons: 817, 945, and 1685. Each curve shows an average over these three recordings.

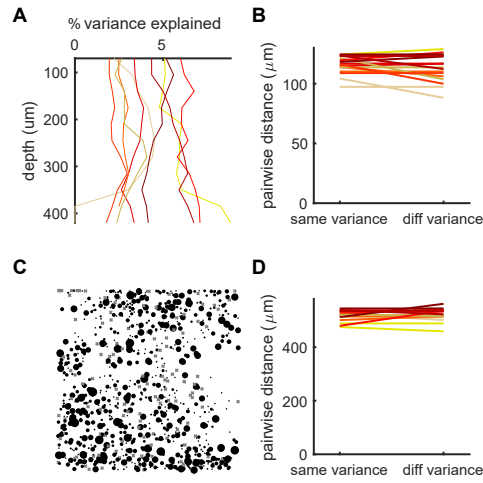


Figure S8: Predictability of neural activity from facial motion is spatially uniform in V1 recordings. (A) Percentage of single-neuron variance explained by facial motion, averaged over cells as a function of cortical depth (16-dimensional reduced rank regression, cross-validated, 1.2 second bins). Each line represents a different experiment. Note that single-neuron prediction explains a lower percentage of variance than SVC prediction (Fig. S9). (B) Neurons were split into two groups: low variance explained ($<3\%$) and high variance explained ($>10\%$). The average vertical distance between neurons in the same group (“same variance”) was similar to the distance between neurons with different variance levels (“diff variance”) ($115\ \mu\text{m}$ vs $117\ \mu\text{m}$, $p > 0.05$ Wilcoxon rank-sum test). The lack of difference indicates that predictability from facial motion does not depend systematically on cortical depth. (C) Fraction of variance explained as a function of XY position in V1, for an example recording plane. Each dot represents a cell, of size proportional to the explained variance; crosses indicate cells of negative explained variance on the test set. (D) Same as B, but for XYZ distance. The lack of difference ($516\ \mu\text{m}$ vs $525\ \mu\text{m}$, $p > 0.05$ Wilcoxon rank-sum test) indicates that variance explained does not depend systematically on XYZ position.

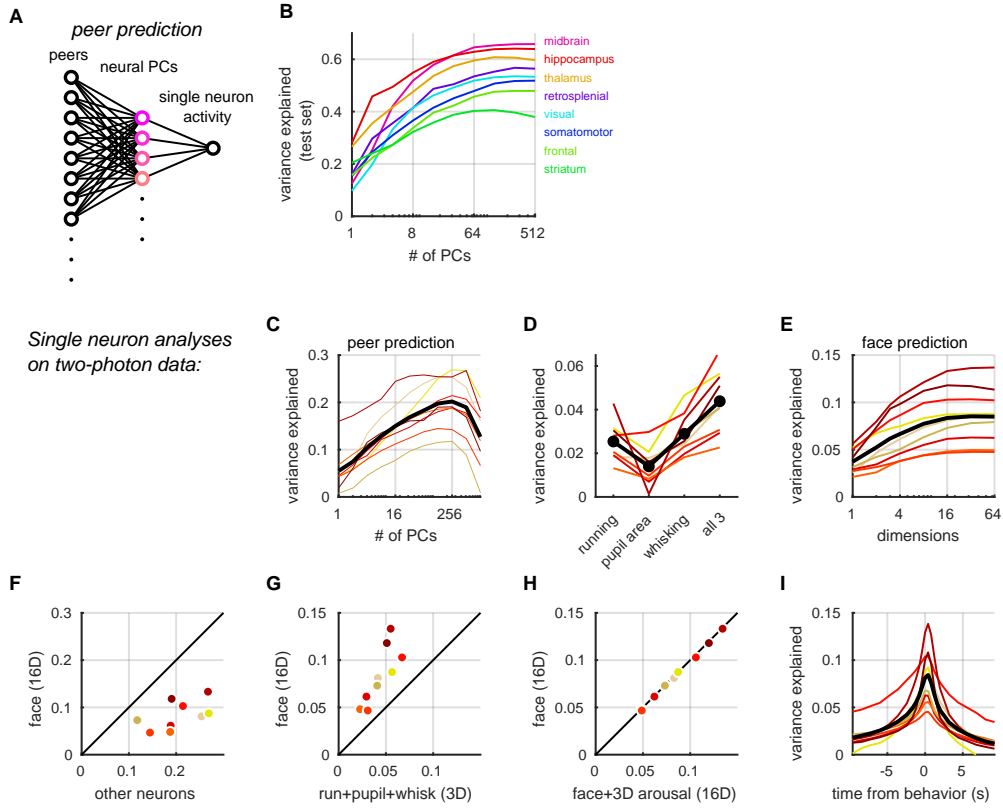


Figure S9: Explaining single neuron variance using peer prediction and behavioral variables. Shared variance component analysis – like a related algorithm for estimating reliable stimulus coding (35) – requires thousands of simultaneously-recorded neurons per brain area. Because this many neurons were not available in our Neuropixels recordings, we instead estimated the reliable variance predictable from facial behavior using an adaptation of the “peer prediction” method (65, 66). (A) In this method, each neuron takes a turn as target cell, and principal component analysis is applied to the activity of its simultaneously recorded peers by ridge regression, excluding spatially neighboring neurons. The activity of the target cell is predicted from these principal components, and prediction quality is assessed via cross-validation. (B) Average single neuron variance explained by peer prediction as a function of the number of principal components, for each brain area (neurons pooled across 3 mice). (C) Similar analysis for two-photon calcium imaging in V1. The peak predictability of $20.2\% \pm 1.7\%$ of variance explained is obtained when predicting from 256 peer PCs; because the independent noise in a single neuron’s activity cannot be predicted from other neurons, this is substantially lower than the 97% reliability of the first SVC, and the 67% reliability of the first 128 SVCs together. (D) Predicting single neuron activity from arousal variables; because the independent noise in a single neuron’s activity cannot be predicted from behavior, predictability is lower than when predicting SVCs (cf. Fig. 10). (E) Predicting single neuron activity from multidimensional behavior information using reduced rank regression (cf. Fig. 2G).

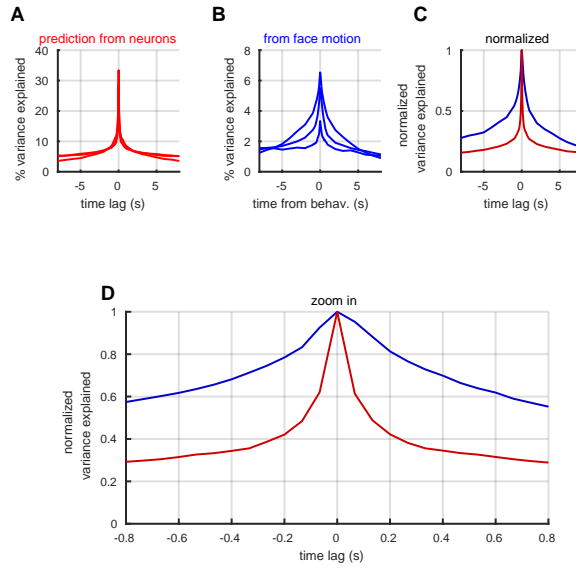


Figure S10: Timescales of neural activity. (A) Prediction of neural activity from other neurons in electrophysiological recordings using SVCA, with a time lag artificially inserted between the two cell sets to determine the timecourse of coherent population activity. Each curve represents a different recording in a different mouse. Neural activity was binned at 30 Hz; this smaller time bin size provides more temporal accuracy but lower percentages than with longer bins (cf. Fig. S7H). The sharp peak indicates coherent structure at fast timescales. (B) Prediction of same neural activity as in panel A using face motion PCs with varying time lags between the face motion and the neural activity. Neural activity and face motion PCs were binned at 30 Hz. (C) The averages of (A) and (B) normalized to a maximum of 1, plotted together to reveal differences in timescale. (D) Zoom in of (C), revealing that videographic prediction operates at slower timescales than neural prediction.

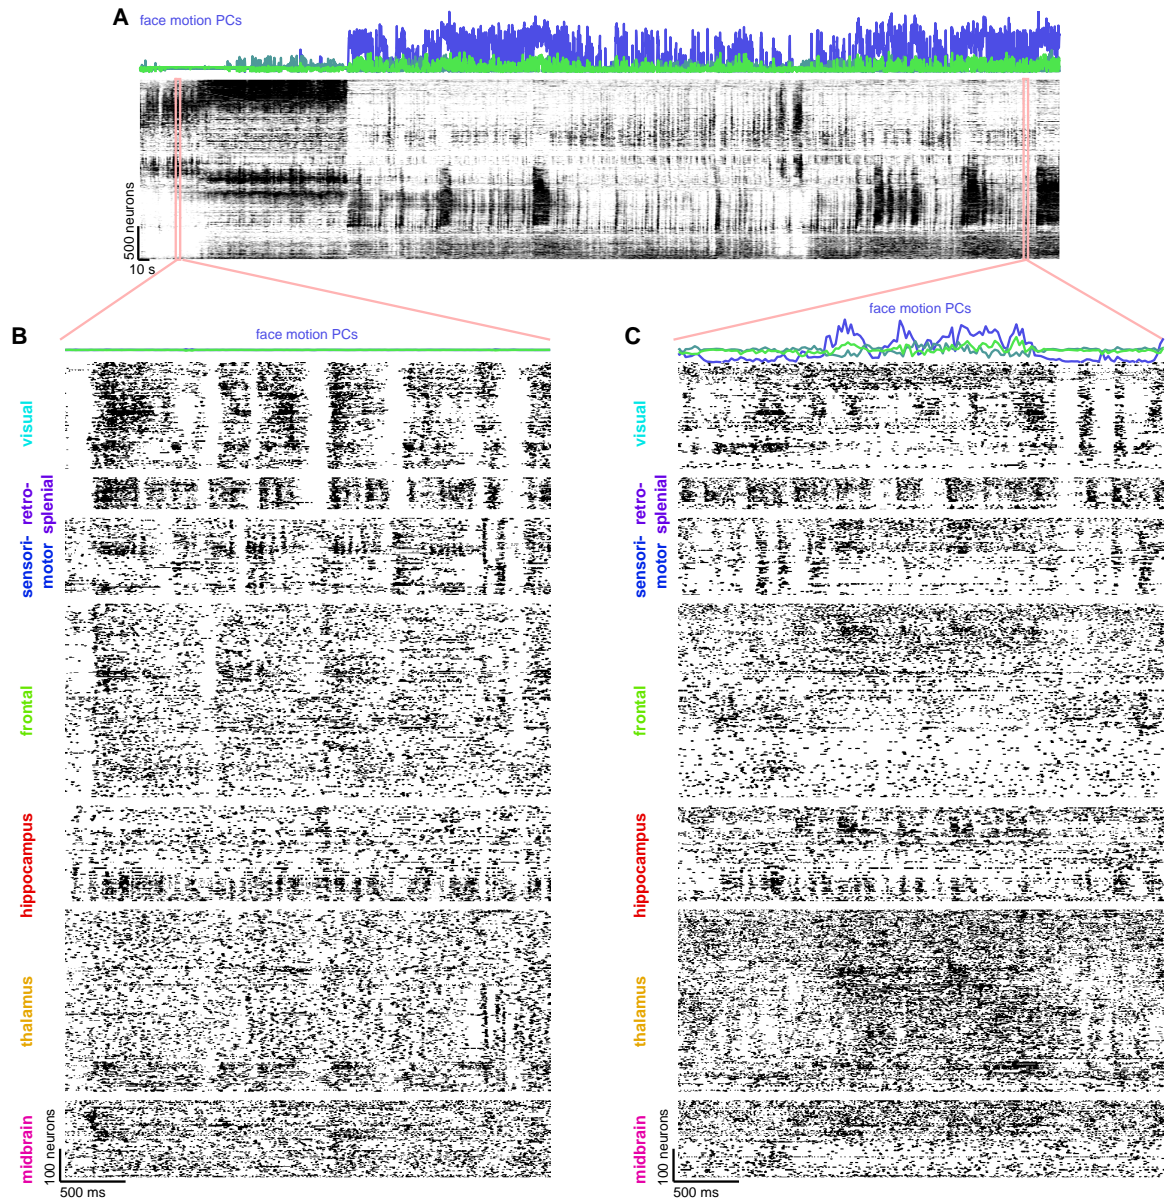



Figure S11: Behaviorally-related population patterns differ from up and down phases. (A) Raster representation of 14 minutes of electrophysiological recording across brain areas, sorted vertically so nearby neurons are correlated (cf. Fig. 3E). Top: timecourse of first three videographic behavior PCs. (B) Zoom in to a shorter time segment with no facial motion, with neurons arranged by brain region. Note that alternating periods of population activity and silence (up and down phases) are present, indicating they do not reflect behavioral events. (C) Zoom in to a shorter time segment with facial motion, showing slower changes in population activity but suppression of up and down phases.

A spontaneous correlations



B autocorrelation of PC's

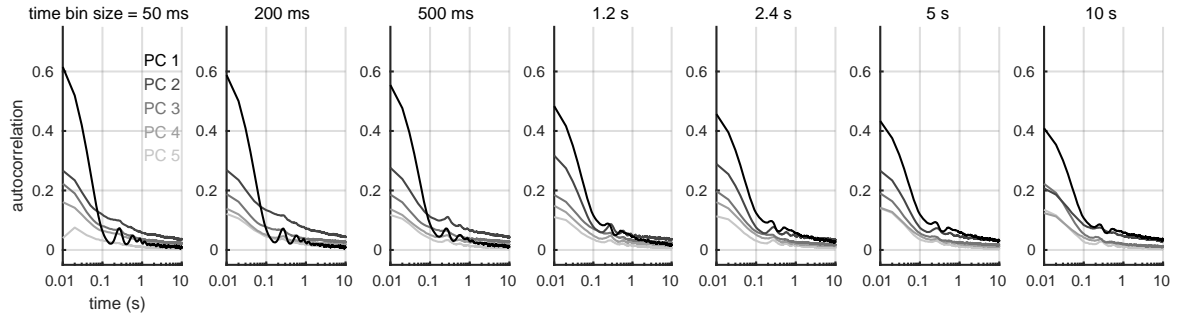


Figure S12: Structure of correlations varies across timescales in visual cortex. (A) Pairwise correlation matrix of spontaneous electrophysiologically-recorded V1 population, computed for different time bin sizes (above each matrix). Neurons were sorted by their weights onto the first principal component of activity binned at each resolution. With 50 ms bins almost all neurons are positively correlated, reflecting up and down phase modulation. With 10 s bins neurons are positively and negatively correlated, reflecting the slower-timecourse behaviorally-related fluctuations also seen in 2-photon recordings. Neurons with spontaneous rates of less than 0.1 Hz were excluded. (B) Temporal autocorrelation of population activity binned at 10 ms resolution, and projected onto principal component dimensions estimated from a range of bin sizes. Short timescale activity was most marked in the 1st PC computed at 50 ms bins (for which correlations were all positive, cf. part A).

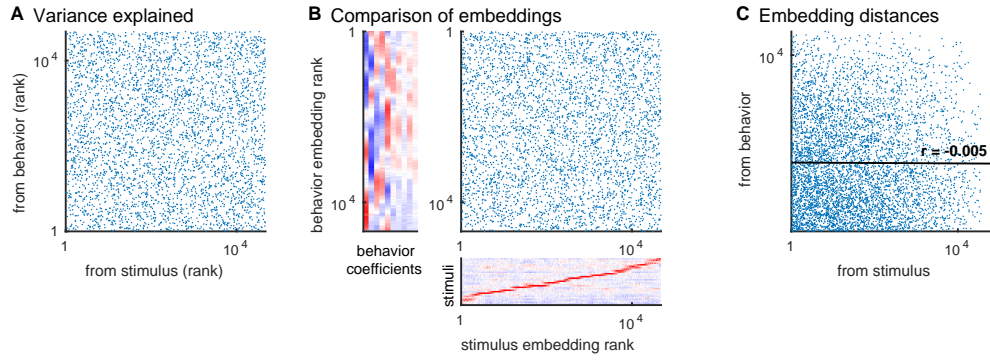


Figure S13: **Mixed representation of stimuli and behavior in visual cortex.** (A) Comparison of variance explained by stimuli and by behavior videography for single neurons. X- and y-coordinates represent each neuron's rank order from least to most predictable by stimulus and behavior, respectively. The extremely weak correlation ($r = -0.18$, $p < 0.01$ Spearman's rank correlation) indicates that the strength with which a neuron is modulated by stimuli is essentially unrelated to the strength with which it is modulated by behavior. (B) Comparison of tuning to behavior and to stimuli. Each point represents a neuron, and the x- and y-coordinates represent its rank in two separate 1D embeddings computed from stimulus responses, and from behavioral coefficients of the face prediction model (16D), respectively. (C) Pairwise comparison of visual-tuning similarity (x- axis) and behavioral-tuning similarity (y- axis). Each point represents a pair of cells; tuning similarities are defined as the pairwise distance in the 1D embeddings on the x- and y-axes of panel B. The lack of correlation ($r = -0.005$, $p > 0.05$) indicates that neurons which are tuned for similar stimuli are no more likely to be tuned for similar behavioral features.

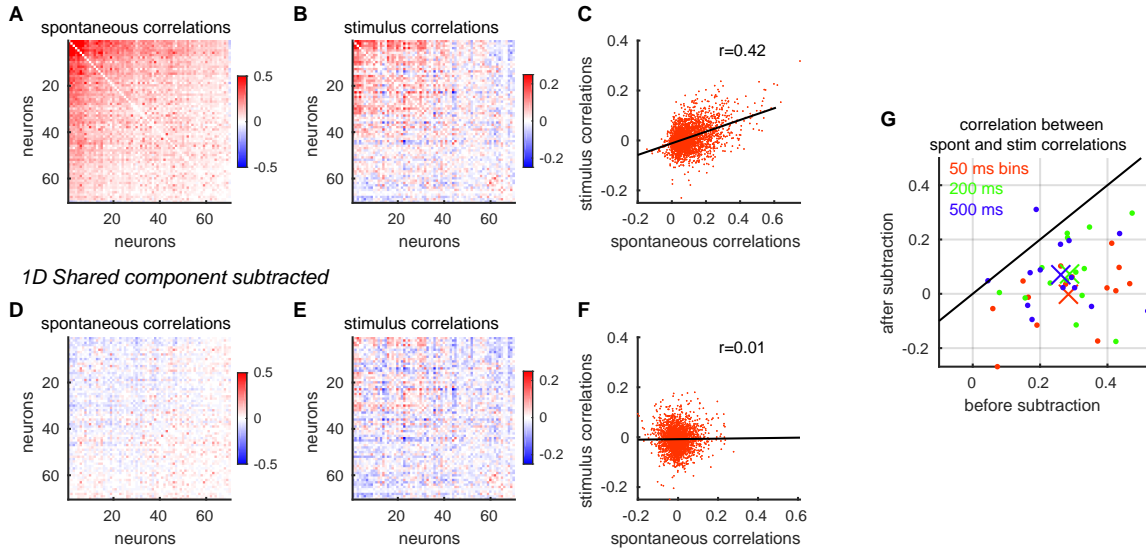


Figure S14: One shared dimension causes similarity of spontaneous and stimulus-driven correlations in electrophysiological recordings. (A) Pairwise correlations between neurons during periods of spontaneous, ongoing activity in 50ms bins (cf. Fig. S12A). Neurons of spontaneous firing rate less than 0.1 Hz or less than 2.5% signal-related variance were excluded. (B) Pairwise signal correlations for these neurons, computed from responses to presentation of 700 natural image stimuli. Stimulus responses were estimated in time bins 20-70 ms following the onset of each stimulus, and an unbiased estimate of signal correlation was obtained using the covariance of the neurons' responses to these stimuli across the two repeats (35). (C) Similarity of spontaneous correlations and signal correlations. Each point represents a pair of neurons from one of fourteen probes in six population recordings in visual cortex. Inset: Pearson correlation coefficient. (D-F) Same analyses as A-C, after projecting out the leading dimension of spontaneous activity estimated by principal component analysis with 50ms bins. The similarity of spontaneous and signal correlations has gone, indicating that this similarity was due to one dimension only. (G) Similarity of spontaneous and signal correlations, before and after projecting out the single shared dimension, for each experiment individually. Each dot represents a different electrophysiological recording. Colors represent 3 different possible bin sizes used to compute correlations of neural activity (in each case, stimulus response windows started 20 ms after stimulus onset). The average Pearson correlation values before subtraction were 0.28, 0.29, and 0.26 for 50, 200 and 500 ms bins respectively, and after subtraction were -0.002, 0.073, and 0.070. This reduction in correlation was significant for each time bin ($n=14$ recording probes, $p < 10^{-4}$, $p < 10^{-4}$, and $p < 10^{-2}$ Wilcoxon one-sided sign-rank test for 50, 200 and 500 ms bins respectively).




Figure S15: **Short-timescale view of projections onto activity subspaces, corresponding to a zoom into Fig. 4G.**

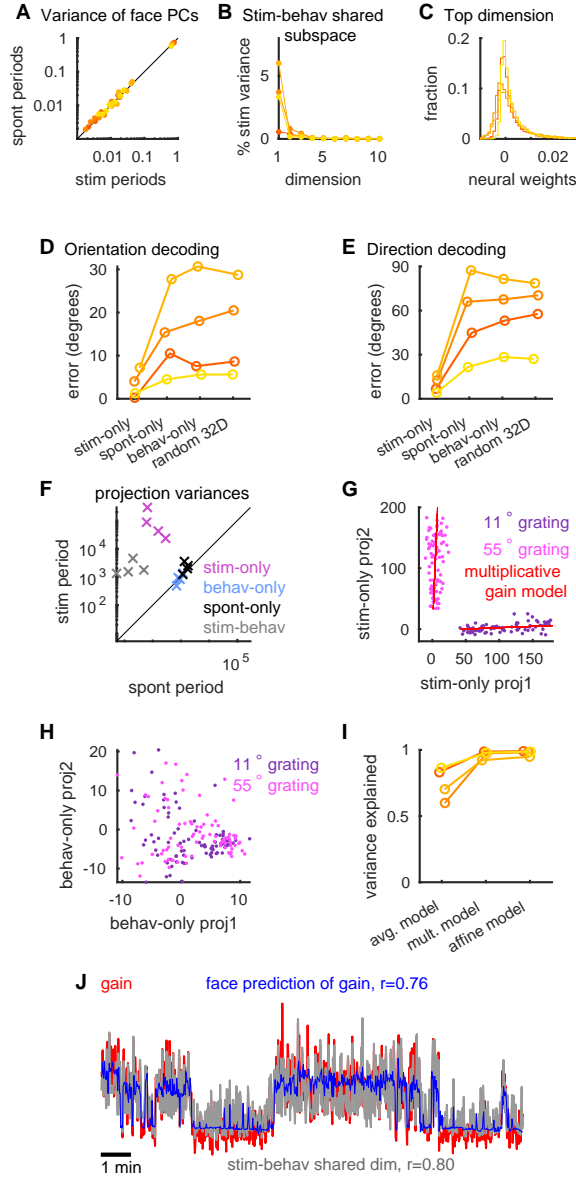
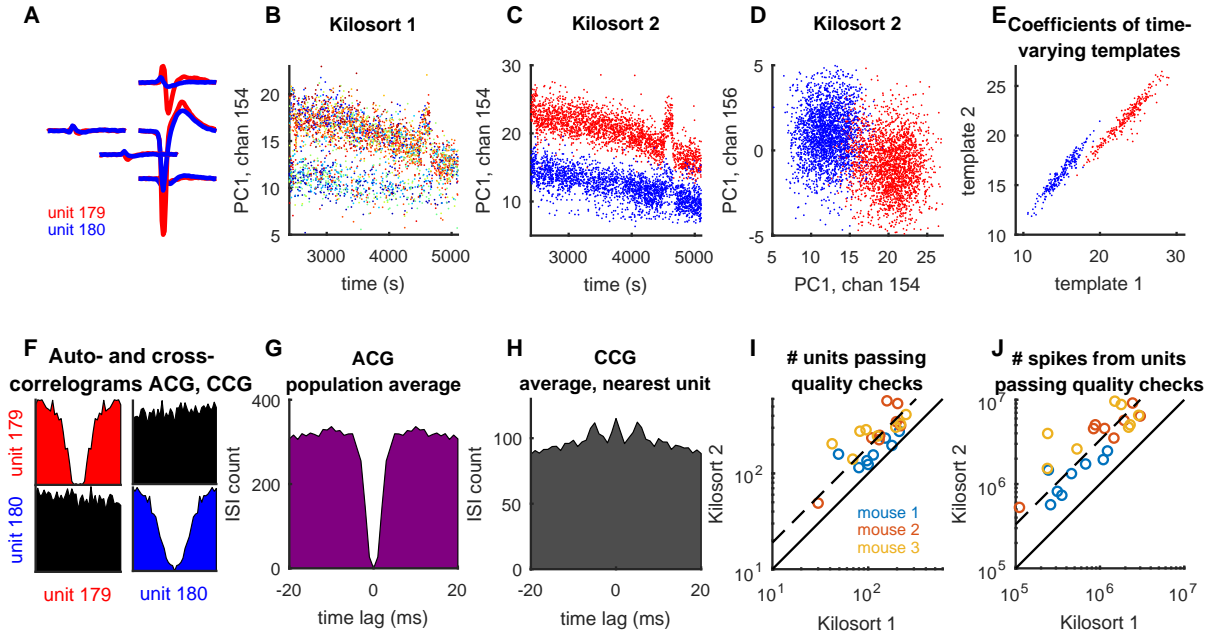


Figure S16: Responses to drifting gratings overlap in only one-dimension with spontaneous/behavioral variables (cf. Fig. 4). (A) Comparison of face motion energy for each PC during presentation of drifting gratings (16 orientations, 32 directions, 11 degree separation), and spontaneous periods. Color represents recording identity. (B) The percentage of stimulus-related variance in each dimension of the shared subspace between stimulus- and behavior-driven activity. (C) Distribution of cells' weights on the single dimension of overlap between stimulus and behavior subspaces. (D) Orientation decoding analysis, from 32 dimensions of activity in the stimulus-only, behavior-only, and spontaneous-only subspaces, together with randomly-chosen 32-dimensional subspaces. Y-axis shows orientation decoding error (degrees). (E) Same as D but for direction decoding. (F) Amount of variance of each of the projections. Each point represents summed variances of the dimensions in the subspace corresponding to the symbol color, for a single experiment. (G) Projection of neural responses to two example stimuli into two dimensions of the stimulus-only subspace. (H) Same as G for the behavior-only subspace. (I) Fraction of variance in the stimulus-only subspace explained by: constant response on each trial of the same stimulus (avg. model); multiplicative gain that varies across trials (mult. model); and a model with both multiplicative and additive terms (affine model). (J) The multiplicative gain on each trial (red) and its prediction from the face motion PCs (blue).



Fully automated spike-sorting with Kilosort2. (A) Waveforms of two example neurons. (B) Spike amplitude on largest channel (measured by waveform principal component 1), shown as a function of time for 20 subclusters found by Kilosort1 that were determined by a human curator to contain spikes from these neurons. (C) Same as B, showing the output of Kilosort2, which identified two single clusters found to account for these two neurons. (D) Spike amplitudes on two channels of largest amplitude shows incomplete separation of the two clusters. (E) Projection of spike waveforms onto the time-varying templates tracked by Kilosort2 shows complete separation of the two clusters. (F) Auto- and cross- correlograms (ACG, CCG) of the two neurons show very few refractory period violations (which would result in an incomplete dip in the ACG) and very little cross-contamination (which would result in a dip in the CCG). (G) Population average of ACGs for all units on one probe, again indicating very little refractory period violation. (H) Population average of CCGs for one probe, including for each good cluster its CCG to the other cluster of most similar waveform. (I) Number of units passing quality checks for Kilosort2 vs Kilosort1 (refractory violations < 20% than expected from Poisson process, and no CCGs with apparent refractory periods at < 20%). Dotted line shows the average ratio of 1.90. (J) Total number of spikes for units passing quality checks. Dotted line shows the average ratio of 3.33.

Figure S17:

Movie S1. Spontaneous neural activity of 10,000+ neurons in visual cortex of awake mice. Two-photon calcium imaging of 11 planes spaced $35\ \mu\text{m}$ apart. Movie speed is 10x real-time.

Movie S2. Multi-dimensional spontaneous behaviors. Movie speed is 5x real-time.

Movie S3. Spontaneous behaviors are correlated with spontaneous neural activity. Video of mouse face recorded simultaneously with neural activity. Movie speed is 10x real-time.