Off-manifold coding in visual cortex revealed by sleep

Eliezyer Fermino de Oliveira[1,*], Soyoun Kim[1,*], Tian Season Qiu[1], Adrien Peyrache[2], Renata Batista-Brito[1], and Lucas Sjulson[1,3,4]

[*]These authors contributed equally to this work.
[1]Dominick P. Purpura Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY, 10461, USA.
[2]Montreal Neurological Institute, McGill University, Montreal, QC, H3A 2B4, Canada.
[3]Department of Psychiatry and Behavioral Sciences, Albert Einstein College of Medicine, Bronx, NY, 10461, USA.
[4]Correspondence should be addressed to luke@sjulsonlab.org

Low-dimensional dynamics and movement-related activity are found throughout the brain. However, primary visual cortex contains high-dimensional sensory representations, raising the question of how low-dimensional dynamics and high-dimensional representations coexist. Here we approached this question by analyzing neuronal activity during slow-wave sleep, which provided reliable estimates of low-dimensional, internally-generated manifold structure. We found that movements and visual scenes were both encoded in the on-manifold subspace, which accounts for more variance than chance during sleep, but visual scenes were also encoded in the off-manifold subspace, which accounts for less variance than chance during sleep. This off-manifold subspace contains sparse activity in the neurons with the strongest low-dimensional modulation by movement, which paradoxically prevents movement-evoked activity from interfering with high-dimensional stimulus representations. These results reveal an unexpected link between low-dimensional dynamics and sparse coding, suggesting that these phenomena play a role in creating and accessing an off-manifold coding space for high-dimensional representations.

Intrinsic manifold | Low-dimensional dynamics | High-dimensional visual representations | Sparse coding

Introduction

Recent studies using large-scale neuronal recordings have reported that brainwide spontaneous neural activity contains low-dimensional representations of movements and internal states (1–3), even in sensory areas such as primary visual cortex (V1). Many studies have suggested that the brain operates in a low-dimensional regime (4–9), but high-dimensional coding of sensory stimuli has also been reported (10) and proposed to confer computational advantages (11, 12). How the brain reconciles low- and high-dimensional representations in the same neuronal population is an open question.

Here we investigated this question using large-scale recordings of neurons during slow-wave sleep (SWS) to probe the intrinsic population dynamics. Although previous work has examined replay during SWS (13–15), replay events account for only a small percentage of neuronal activity in SWS (16, 17). We aimed instead to use SWS as a window into internally-generated population activity that is uncontaminated by the influence of the ongoing behavior and sensory inputs present during wakefulness (18–20).

SWS activity provided reliable estimates of internally-generated low-dimensional manifold structure, enabling us to analyze awake activity in the on-, non-, and off-manifold subspaces defined by the amount of variance they account for during SWS. We found that wakefulness is associated with increased off-manifold activity in many brain regions, and in V1 this off-manifold activity encodes high-dimensional visual stimuli, but not movements. Further analysis and modeling of off-manifold activity revealed an unexpected link between sparse coding and dimensionality, suggesting that an unappreciated function of low-dimensional manifolds and sparse coding may be to create and access an off-manifold coding space that can store high-dimensional representations without interference from movement-evoked activity.

Results

Neural activity in slow-wave sleep reveals low-dimensional internally-generated population structure

Internally-generated population structure is mostly preserved in any brain state (7, 18, 19, 21), but activity evoked by uncontrolled movements and sensory inputs during wakefulness could contaminate estimates of this population structure. We reasoned that SWS is maybe closer to a true "ground state" of the brain and could provide a better estimate. We tested this by first investigating if the population structure observed in sleep is preserved in wakefulness. For this purpose, we performed recordings with long periods of SWS and awake states in V1. We used cross-validated Principal Component Analysis (cvPCA) to estimate the population structure in the data and found that SWS population geometry is mostly preserved in awake states (Fig. 1A), with high-variance SWS dimensions explaining the same variance in awake population activity in several brain areas (Fig S1 - 97.1 ± 0.7%; mean ± s.e.m., n = 122 recordings, multiple brain areas), even after subsampling to account for differences in firing rates (90.5 ± 3.4%; mean ± s.e.m., n = 122 recordings, multiple brain areas Fig. S2).

We next tested whether SWS activity would provide a more reliable estimate of the intrinsic population geometry than

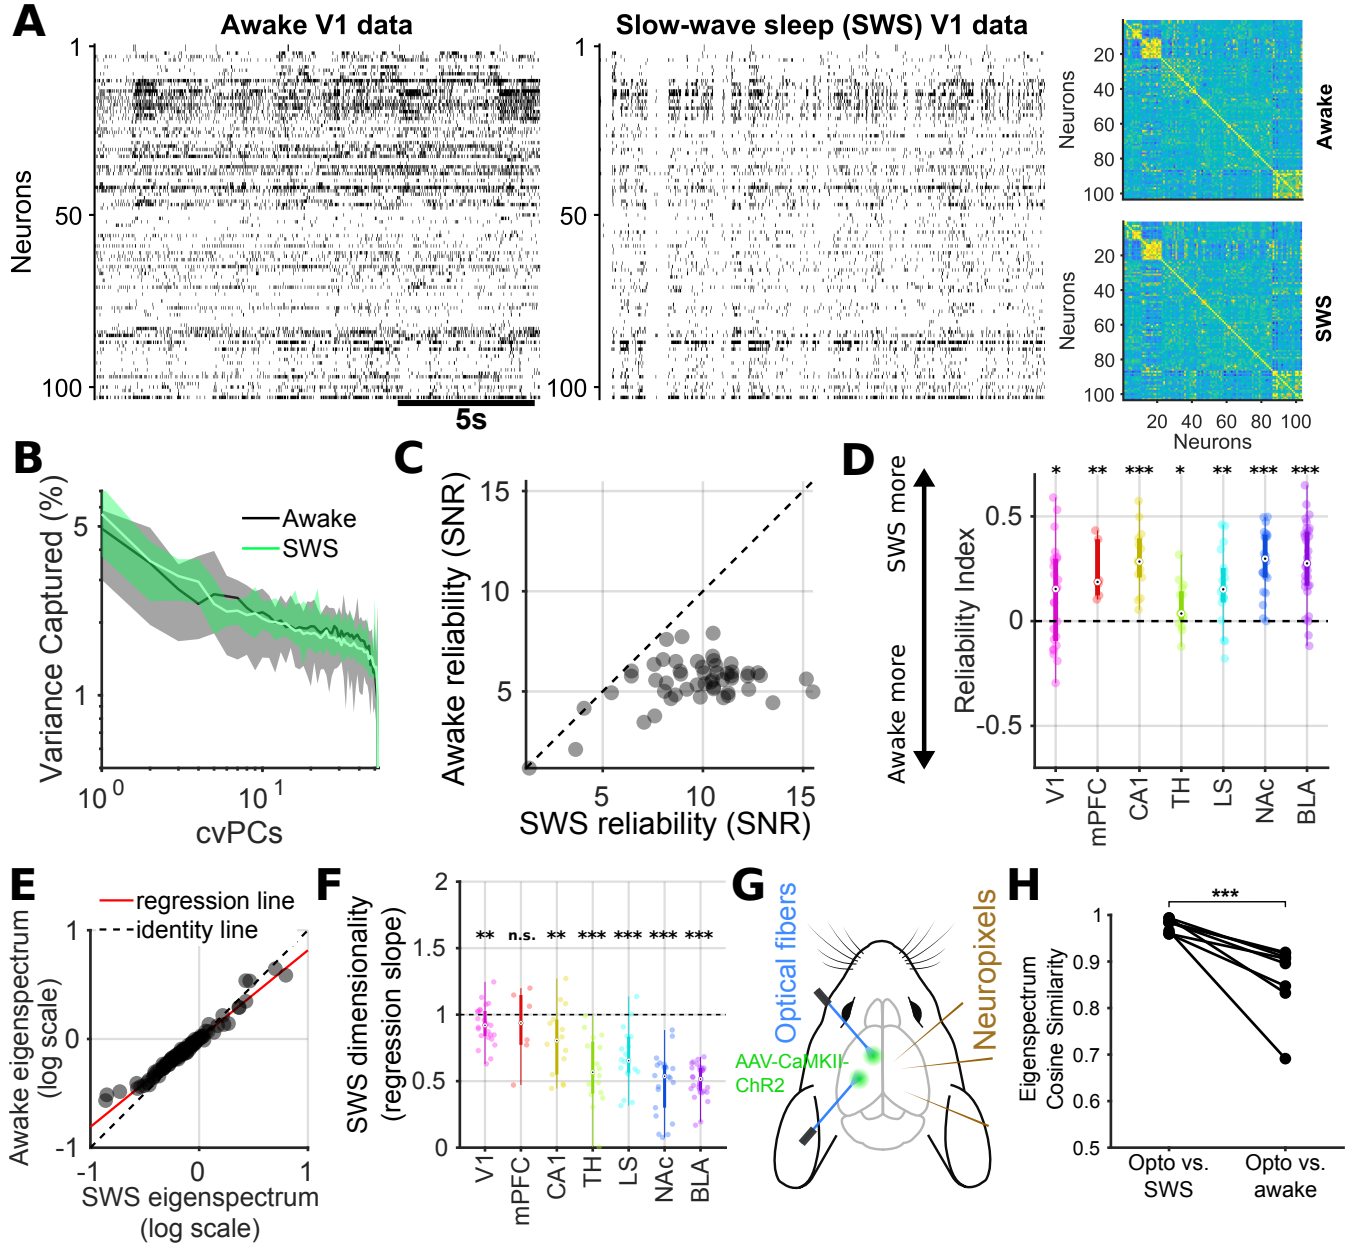

Fig. 1. Slow-wave sleep (SWS) enables reliable estimation of internally-generated low-dimensional population structure. (**A**) Typical segment of data illustrating that the correlation structure of activity in V1 is similar during wakefulness and slow-wave sleep (SWS). (**B**) Cross-validated principal components analysis (cvPCA) eigenspectra estimated from five adjacent segments of data show more segment-to-segment variability in wakefulness than in SWS (note wider confidence interval), suggesting SWS provides a more reliable estimate of population geometry. (**C**) Individual eigenvalues from panel B are estimated more reliably from SWS than from awake spontaneous activity. (**D**) cvPCA eigenspectra can be estimated more reliably from SWS in all datasets tested (V1: primary visual cortex, mPFC: medial prefrontal cortex, HPC: hippocampus, TH: thalamus, LS: lateral septum, NAc: nucleus accumbens, BLA: basolateral amygdala). (**E**) cvPCA eigenspectra decay faster in SWS than wakefulness (regression slope < 1), indicating that SWS is lower-dimensional. (**F**) Linear regression slopes of awake and SWS eigenspectra show that most brain areas tested have higher dimensionality during awake states. (**G**) We reasoned that if SWS population structure is generated internally, it should be evoked by diffuse, nonspecific optogenetic stimulation. We tested this by stimulating in S1 and M1 while recording in the contralateral hemisphere. (**H**) The eigenspectrum of optostimulation-evoked activity was more similar to SWS than wakefulness, indicating similar population structure (multi-brain area pooled, paired t-test, $p < 0.001$) [two-tailed t-test, *p<0.05, **p<0.01, ***p<0.001, n.s. not significant ($p > 0.05$)].

awake spontaneous activity. To accomplish this, we made independent estimates of the population activity eigenspectrum from short segments of the data (14.98 ± 0.78 min, mean ± s.e.m., n = 122 recordings, multiple brain areas), finding that SWS eigenspectrum estimates showed lower segment-to-segment variability in V1 (Fig. 1B-C). We used this variability to derive an index of reliability (see methods), finding that in both our V1 recordings and in publicly available datasets (16, 18, 22–25), SWS provides more reliable estimates of population geometry than awake spontaneous activity (Fig. 1D). We reasoned further that if SWS is closer to a ground state, then the activity should be lower-dimensional than in wakefulness. To test this, we compared the rates of eigenspectrum decay (10), finding that SWS activity was lower-dimensional in all brain regions tested, with the exception of mPFC (Fig. 1E-F). Interestingly, REM sleep was higher-dimensional than awake activity (Fig. S9), so we excluded it from our analyses.

Finally, we asked whether the population geometry we observed during SWS was truly generated by internal dynamics as opposed to driven by highly structured inputs, such as descending motor activity encoding specific movements whose execution is blocked during SWS. To test this, we used head-fixed Neuropixels probes to record from multiple sites in one hemisphere during both SWS and wakefulness, then optogenetically stimulated nonspecific populations of neurons in contralateral M1 and S1 (Fig. 1G). Our rationale was that if SWS population geometry is internally-generated, then any diffuse, unstructured stimulation during wakefulness should evoke SWS-like activity patterns. In support of this, we found that the eigenspectrum was more similar to SWS than wakefulness during optostimulation. (Fig. 1H, S3).

Off-manifold dimensions are activated in wakefulness

For our subsequent analyses, we operationally defined three subspaces, which contain cvPCs that account for more ("on-manifold"), equal ("non-manifold"), or less ("off-manifold") variance than chance during SWS, as defined as the confidence interval of the eigenspectrum from shuffled SWS data (Fig. 2A). Although the high-variance cvPCs do not define a manifold per se, we use the term "on-manifold" to indicate a subspace that likely contains an internally-generated manifold (4, 5, 19, 20). The off-manifold subspace is also noteworthy: these dimensions are preserved in cross-validation and therefore are not merely noise. Instead, they represent patterns of neuronal activity that are reproducibly *less likely than chance* to occur during SWS.

To examine the similarities and differences between awake and SWS activity, we projected the awake data onto the SWS cvPCs (Fig. 2B) and determined the amount of variance captured by each cvPC. We found that the population structure was similar to SWS, but with increased activity in the off-manifold subspace (Fig. 2C). This increase was observed in the majority of brain areas analyzed (Fig. 2D), and it was not observed when the analysis was performed in the opposite direction (Fig. S4), suggesting it is not a trivial result of the fact that PCA minimizes variance in the off-manifold subspace during SWS. To guard further against this possibility, we developed a novel type of PCA, normalized contrastive PCA (ncPCA), which finds the dimensions exhibiting the largest increase in variance in the awake state. Unlike standard contrastive PCA (26), ncPCA is normalized to remove the bias toward high-variance dimensions. We found that the top ncPCs are preferentially off-manifold in all brain regions tested (Fig. 2E-F). Together, these results suggest that awake activity contains the same low-dimensional population structure seen in SWS, but with additional high-dimensional activity in the off-manifold subspace.

Movements and visual stimuli are both encoded on-manifold, but only stimuli are encoded off-manifold

Since natural scenes have high-dimensional representations in V1 (10), we next investigated how movements and natural visual scenes are encoded in the on/non/off-manifold subspaces (Fig. 3A). For this, we projected awake activity into each of the three subspaces and tested the ability of the projected data to predict movements and stimulus identity relative to a randomly-selected subspace of the same dimension. In our acute multisite Neuropixels recordings, we found that face motion, a multidimensional movement, is encoded in the on-manifold subspace (Fig. 3B). We also examined several publicly available datasets containing both SWS and freely-moving maze exploration, and we found that running speed is encoded in the on-manifold subspace as well (Fig. 3C). It has previously been suggested that distributed movement representations are created by efference copy and sensory reafference (2, 27). Although these inputs undoubtedly drive on-manifold activity, our results indicate that the manifold structure itself is created instead by the internal dynamics preserved during SWS.

Unexpectedly, we found that stimulus identity in V1 is encoded in both the on- and off-manifold subspaces, but not the non-manifold subspace (Fig. 3C, S6, S7). The on-manifold component is consistent with previous studies showing that stimulus-evoked activity in V1 is more similar to spontaneous activity than chance (7, 28, 29). However, our results indicate that natural scenes are also encoded off-manifold by patterns of activity which are less likely than chance to occur during SWS (Fig. 3C). Moreover, the decoding performance was better using on- and off-manifold activity combined than using either one alone (Fig. S8), indicating that these subspaces encode non-redundant information.

Off-manifold activity is composed of "chorister" neurons firing separately from the rest of the population

Given this surprising result, we next investigated which neuronal activity patterns constitute the off-manifold subspace. Since natural visual scenes are known to be encoded by sparse activity in V1 (30–32), we first quantified the population sparsity of each SWS cvPC by calculating the Gini coefficient (33). Our analysis revealed that off-manifold cvPCs have higher population sparsity than on- and non-manifold ones (Fig. 4A-B), meaning that fewer neurons participate in each dimension. It has been reported that cortical neurons exhibit a broad distribution of population coupling strengths ranging from "soloists," who fire decoupled from the rest of the population, to "choristers," whose activity is strongly coupled to

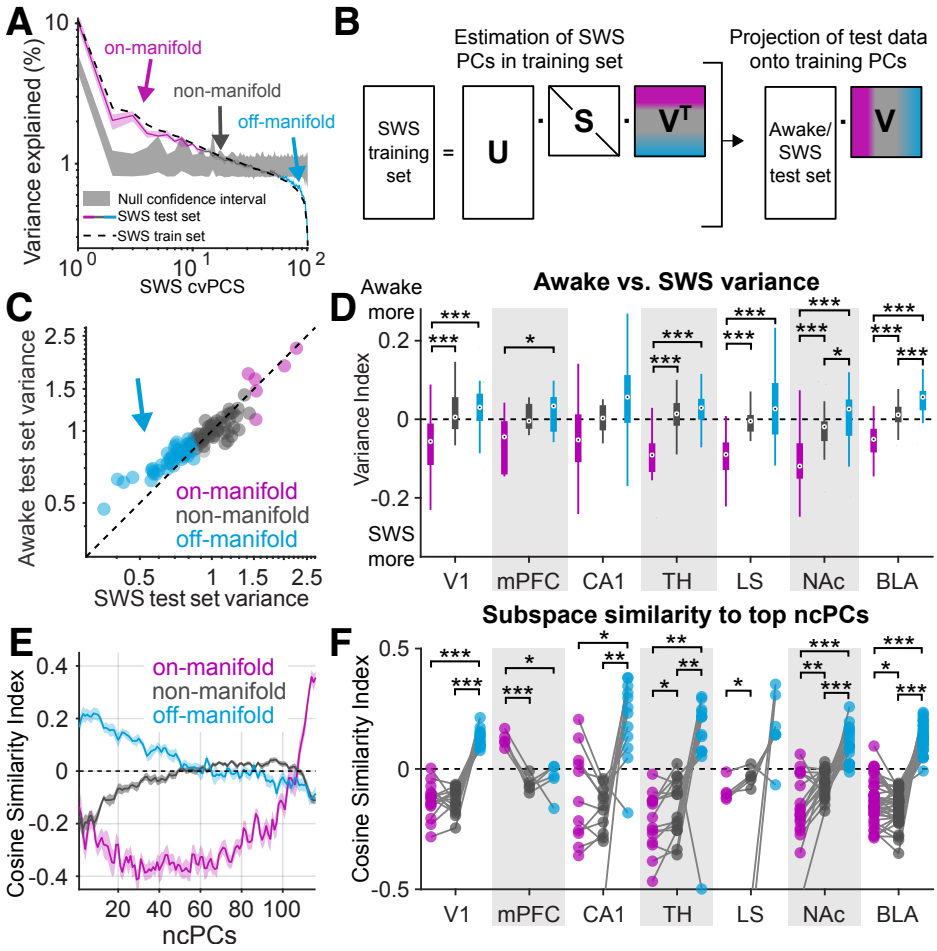

Fig. 2. Wakefulness exhibits increased off-manifold activity relative to SWS. (**A**) Using SWS cvPCA, we operationally defined three subspaces: on-manifold (magenta), non-manifold (gray), and off-manifold (cyan), which during SWS account for more, equal, or less variance than chance, respectively. Example shows a single session in mouse V1. (**B**) To measure the amount of variance in each subspace, we projected SWS and awake test sets onto cvPCs from an SWS training set. (**C**) Awake test set activity shows preferential activation of off-manifold dimensions relative to the SWS test set. (**D**) This was true in most brain regions tested, with the exception of dorsal CA1, which showed a trend that failed to reach significance. (**E**) Normalized contrastive PCA (ncPCA) finds the dimensions that differ most between awake and SWS data after normalization for total variance. ncPCA reveals that the dimensions that are most overrepresented in wakefulness share the greatest similarity with the off-manifold subspace. (**F**) The top 10% of ncPCs are more similar to off-manifold dimensions than on-/non-manifold for every brain area tested except for mPFC and LS [repeated measures, multiple comparisons *p<0.05, **p<0.01, ***p<0.001].

the population (8). We therefore asked whether off-manifold dimensions are sparse because they preferentially contain activity in soloist neurons. We found the opposite result: off-manifold dimensions were more likely to contain choristers, which are strongly coupled to the rest of the population (Fig. 4C) and modulated by movement (Fig. S5); soloist neurons were more likely to participate in non-manifold dimensions (Fig. 4C). The explanation for this seemingly contradictory result is that the off-manifold subspace contains population-sparse activity in neurons with low *average* population sparsity; in other words, it is when chorister neurons fire decoupled from the population (Fig. 4D). Because this is statistically unlikely to occur, it is observed at less-than-chance levels during SWS. In contrast, non-manifold activity occurs at chance levels because population-sparse activity is common in soloist neurons.

Sparse activity and low-dimensional dynamics enable stimuli to access an off-manifold coding space

These findings imply that not all sparse activity in V1 is equal: sparse firing carries more information about stimulus identity when it occurs in chorister neurons, which are strongly population coupled and carry movement-related activity. However, movement and stimulus representations are orthogonal (1) in V1, presumably to prevent them from interfering with each other. If keeping movement and stimuli separate were the goal, it was not obvious to us why it would be advantageous to encode visual stimuli in the neurons with the strongest movement-related activity. To gain insight into this phenomenon, we built a model of Poisson-spiking neurons that receive a dense movement-evoked "noise" input with a distribution of input strengths ranging from low (soloists) to high (choristers) (Fig. 4E). A set of neurons was chosen to encode

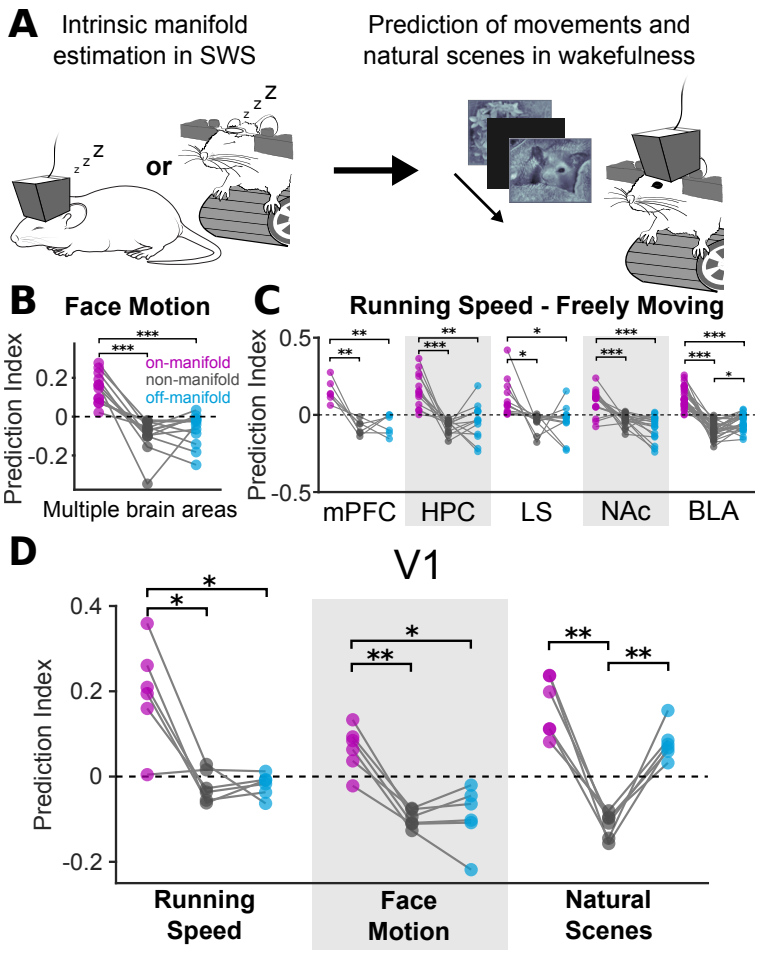

Fig. 3. Natural scenes, but not movements, are encoded in the off-manifold subspace. (**A**) We recorded both SWS and wakefulness in freely moving and head-fixed animals. We then used awake activity in the on/non/off-manifold subspaces, and decoded spontaneous movements and visual stimulus identity. (**B**) In Neuropixels recordings across multiple brain areas, face motion was encoded in the on-manifold subspace (multiple brain areas pooled). (**C**) In existing datasets from freely-moving animals, running speed was also encoded in the on-manifold subspace. (**D**) In V1, we found that running speed and face motion were also encoded on-manifold. However, natural scenes are encoded in both the on- and off-manifold subspaces, but not the non-manifold subspace. The prediction index quantifies prediction performance relative to a random subspace of the same dimensionality [repeated measures, multiple comparisons *p<0.05, **p<0.01, ***p<0.001].

visual stimuli ("signal"), with each stimulus activating a sparse subset of those neurons. We then systematically varied which neurons received the visual input, ranging from soloists to choristers, which creates coding that ranges from non-manifold to off-manifold (Fig. S12). Consistent with naive intuition, we found that when the signal-to-noise ratio is high, non-manifold stimulus coding by soloists yielded optimal stimulus decoding performance (Fig. 4F). However, when the signal-to-noise ratio decrease, off-manifold coding by choristers becomes optimal. An explanation for this counterintuitive result is that for soloist (i.e. non-manifold participating) neurons, spontaneous activity travels throughout the whole state space (Fig. 4G), preventing stimuli from being encoded separately from movements. However, for chorister (i.e. off-manifold participating) neurons, spontaneous activity is constrained to a low-dimensional manifold, opening an unused region of state space that can be used to encode the visual stimulus (Fig. 4H).

Discussion

This work aims to reconcile several disparate findings in the literature: brainwide representations of movement (1, 2), low-dimensional dynamics (4–6, 8, 9), and high-dimensional representations of sensory stimuli (10). The population structure of movement-evoked activity has generally been assumed to be generated by motor efference copy or sensory reafference (2, 27). However, we found that this structure is intact during sleep and evoked by diffuse optogenetic stimulation, suggesting that efference copy and sensory reafference are merely impinging upon internally-generated low-dimensional dynamics to generate movement-related activity. Our finding that visual scenes are encoded partly in the on-manifold subspace suggests that structured sensory inputs can also activate these internally-generated dynamics, consistent with previous studies showing that evoked and spontaneous activity in V1 are more similar than expected by chance (28, 29). However, we unexpectedly found that stimulus identity is also encoded

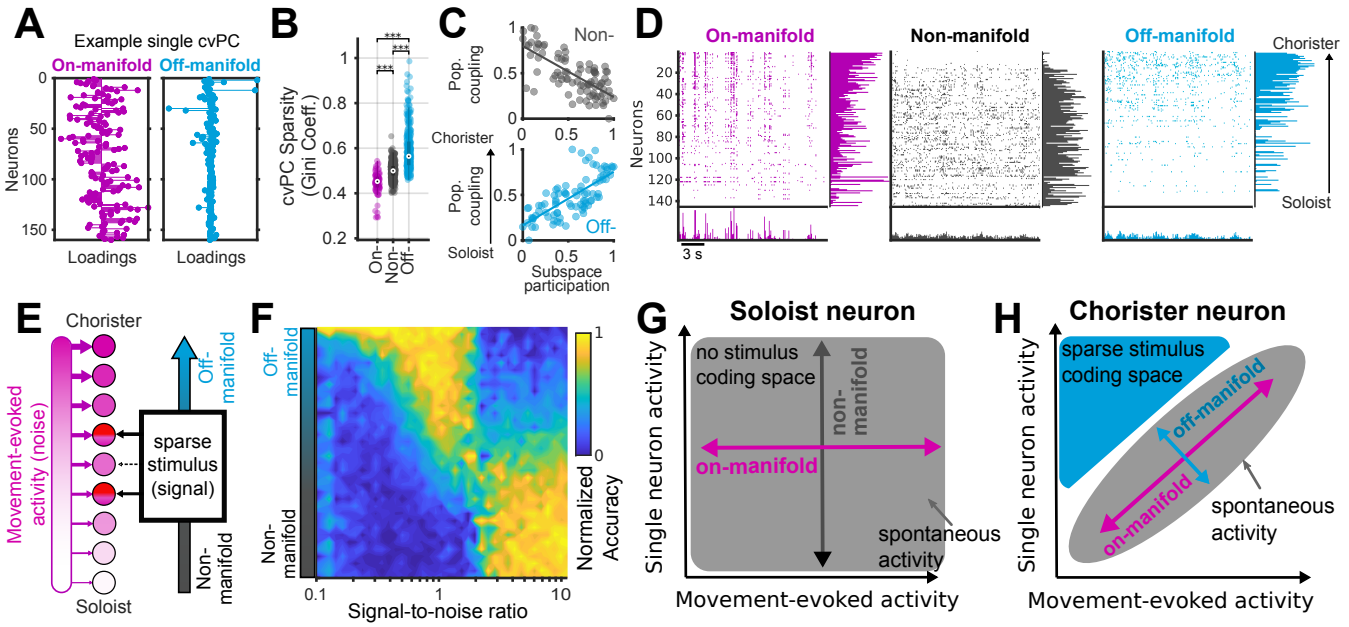

Fig. 4. The off-manifold subspace consists of population-sparse activity in "chorister" neurons, which have strong population coupling. (**A**) Fewer V1 neurons participate in a typical off-manifold cvPC (right) than a typical on-manifold cvPC (left). (**B**) Off-manifold cvPCs have higher population sparsity (i.e. fewer neurons participate) than on- or non-manifold cvPCs (Gini coefficient, one-way ANOVA, $p<0.001$). (**C**) Neurons with strong non-manifold participation have weaker average population coupling ("soloists"), and those with strong off-manifold participation have higher average population coupling ("choristers"). This indicates that although off-manifold cvPCs have higher population sparsity, the neurons that participate in them have lower average population sparsity (t-test $p<0.001$). (**D**) Off-manifold activity is when chorister neurons fire decoupled from the rest of the population. (left) Pseudo-raster plot of on-manifold activity in a sample session in V1. Rows are ordered by decreasing population coupling. Note that most activity is on the top of the plot, indicating that choristers have high on-manifold participation, and the bottom histogram has high peaks indicating low population sparsity. (center) Non-manifold activity in the same neurons. Most of the activity is in soloist neurons, and the flat bottom histogram indicates high population sparsity. (right) In the off-manifold subspace, chorister neurons have higher participation, but the bottom histogram is flat, indicating high population sparsity. (**E**) To understand possible advantages of off-manifold coding, we made a model of Poisson-spiking neurons that receive varying amounts of movement-evoked "noise." Neurons with the greatest noise were choristers (top), and those with the least noise were soloists (bottom). Visual stimuli were encoded by sparse inputs to a subset of neurons that were varied to change how much stimulus evoked activity was in the non- vs. off-manifold subspaces. (**F**) We found that for low signal-to-noise ratios, stimulus decoding performance was better when the stimulus was encoded off-manifold in choristers. This counterintuitively indicates that it is easier to separate signal and noise when they are encoded in the same neurons rather than different ones. (**G**) An intuitive explanation for this result is that in (weakly-coupled) soloist neurons, spontaneous activity enters the non-manifold space at chance levels, so there is no dedicated space to encode the stimulus. (**H**) For (strongly-coupled) chorister neurons, activity is coupled to movement, creating an off-manifold subspace (cyan) that spontaneous activity is unlikely to enter. This off-manifold coding space can be used to encode high-dimensional stimuli separate from movement-evoked activity.

in the off-manifold subspace, which contains activity patterns that occur at less-than-chance levels. This suggests that highly structured inputs can also evoke population activity patterns that are not generated spontaneously, providing a possible explanation for how high-dimensional representations can be compatible with low-dimensional dynamics. Whether these off-manifold patterns passively reflect the structure of their inputs or are stored in the local network and triggered by structured inputs will be an interesting question for future studies.

Our finding that off-manifold activity is composed of population-sparse activity in chorister neurons, which are the least likely to fire alone, reveals an unexpected link between dimensionality and sparse coding. Sparse representations have been reported in numerous brain regions of different species (30, 34–36), where they have been proposed as a mechanism to reduce energy consumption (37). Our results suggest that another important function of sparse coding may be to access the off-manifold subspace. Conversely, an underappreciated function of the low-dimensional dynamics found in many brain regions may be to open an off-manifold coding space for high-dimensional representations. Our finding that wakefulness is associated with increased off-manifold activity in many brain regions suggests this may be a general principle that holds outside of V1. However, the fact that we did not observe this in mPFC implies interesting functional differences between cortical areas.

This study has several important limitations, including the fact that we did not attempt to use nonlinear dimensionality reduction methods to determine the intrinsic dimension

or structure of the underlying manifold. Additionally, we did not attempt to determine which specific stimulus features are coded in the on- and off-manifold subspaces. Future studies addressing these questions could shed light on the possible roles of off-manifold activity in encoding higher-order stimulus features or unexpected deviations from predicted movement-evoked changes. Our findings highlight the benefits of using sleep as a tool to study awake brain function, even in primary sensory areas such as V1. Many previous studies have compared evoked activity to awake spontaneous activity (8, 9, 28), but our results suggest that SWS provides a more reliable estimate of internally-generated low-dimensional population structure than awake activity. We encourage others to incorporate sleep recordings into their experiments as a relatively low-effort means to gain insight into the interactions between internally-generated dynamics and representations of sensory and cognitive variables.

References

1. Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):255–255, April 2019.
2. Simon Musall, Matthew T Kaufman, Ashley L Juavinett, Steven Gluf, and Anne K Churchland. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.*, 22(10):1677–1686, October 2019.
3. William E Allen, Michael Z Chen, Nandini Pichamoorthy, Rebecca H Tien, Marius Pachitariu, Liqun Luo, and Karl Deisseroth. Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*, 364(6437):253, April 2019.
4. J K Chapin and M A Nicolelis. Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. *J. Neurosci. Methods*, 94(1):121–140, December 1999.
5. Juan A Gallego, Matthew G Perich, Lee E Miller, and Sara A Solla. Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, June 2017.
6. Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, July 2012.
7. Tal Kenet, Dmitri Bibitchkov, Misha Tsodyks, Amiram Grinvald, and Amos Arieli. Spontaneously emerging cortical representations of visual attributes. *Nature*, 425(6961):954–956, October 2003.
8. Michael Okun, Nicholas Steinmetz, Lee Cossell, M Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, Sonja B Hofer, Thomas D Mrsic-Flogel, Matteo Carandini, and Kenneth D Harris. Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521(7553):511–515, May 2015.
9. Artur Luczak, Peter Barthó, and Kenneth D Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, May 2009.
10. Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D Harris. High-dimensional geometry of population responses in visual cortex. *Nature*, 571 (7765):361–365, July 2019.
11. Mattia Rigotti, Omri Barak, Melissa R Warden, Xiao-Jing Wang, Nathaniel D Daw, Earl K Miller, and Stefano Fusi. The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590, May 2013.
12. Stefano Fusi, Earl K Miller, and Mattia Rigotti. Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.*, 37:66–74, April 2016.
13. C Pavlides and J Winson. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.*, 9(8):2907–2918, August 1989.
14. M A Wilson and B L McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, July 1994.
15. David J Foster. Replay comes of age. *Annu. Rev. Neurosci.*, 40(1):581–602, July 2017.
16. Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I Wiener, and Francesco P Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.*, 12(7):919–926, July 2009.
17. Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.*, 10(1):100–107, January 2007.
18. Adrien Peyrache, Marie M Lacroix, Peter C Petersen, and György Buzsáki. Internally organized mechanisms of the head direction sense. *Nat. Neurosci.*, 18(4):569–575, April 2015.
19. Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.*, 22(9):1512–1520, September 2019.
20. Richard J Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells. *Nature*, pages 1–6, January 2022.
21. Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. December 2021.
22. Gabrielle Girardeau, Ingrid Inema, and György Buzsáki. Reactivations of emotional memory in the hippocampus–amygdala system during sleep. *Nat. Neurosci.*, 20(11):1634–1642, September 2017.
23. Lucas Sjulson, Adrien Peyrache, Andrea Cumpelik, Daniela Cassataro, and György Buzsáki. Cocaine place conditioning strengthens Location-Specific hippocampal coupling to the nucleus accumbens. *Neuron*, 98(5):926–934.e5, June 2018.
24. Yuta Senzai, Antonio Fernandez-Ruiz, and György Buzsáki. Layer-Specific physiological features and interlaminar interactions in the primary visual cortex of the mouse. *Neuron*, 101(3):500–513.e5, February 2019.
25. David Tingley and György Buzsáki. Transformation of a spatial map across the Hippocampal-Lateral septal circuit. *Neuron*, 98(6):1229–1242.e5, June 2018.
26. Abubakar Abid, Martin J Zhang, Vivek K Bagaria, and James Zou. Exploring patterns enriched in a dataset with contrastive principal component analysis. *Nat. Commun.*, 9(1):2134, May 2018.
27. Aman B Saleem, Aslı Ayaz, Kathryn J Jeffery, Kenneth D Harris, and Matteo Carandini. Integration of visual motion and locomotion in mouse visual cortex. *Nat. Neurosci.*, 16(12): 1864–1869, December 2013.
28. József Fiser, Chiayu Chiu, and Michael Weliky. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008):573–578, September 2004.
29. Pietro Berkes, Gergo Orbán, Máté Lengyel, and József Fiser. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013): 83–87, January 2011.
30. W E Vinje and J L Gallant. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276, February 2000.
31. Michael Weliky, József Fiser, Ruskin H Hunt, and David N Wagner. Coding of natural scenes in primary visual cortex. *Neuron*, 37(4):703–718, February 2003.
32. Emmanouil Froudarakis, Philipp Berens, Alexander S Ecker, R James Cotton, Fabian H Sinz, Dimitri Yatsenko, Peter Saggau, Matthias Bethge, and Andreas S Tolias. Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat. Neurosci.*, 17(6):851–857, June 2014.
33. Fabio Stefanini, Lyudmila Kushnir, Jessica C Jimenez, Joshua H Jennings, Nicholas I Woods, Garret D Stuber, Mazen A Kheirbek, René Hen, and Stefano Fusi. A distributed neural code in the dentate gyrus and in CA1. *Neuron*, 107(4):703–716.e4, August 2020.
34. M P Young and S Yamane. Sparse population coding of faces in the inferotemporal cortex. *Science*, 256(5061):1327–1331, May 1992.
35. Javier Perez-Orive, Ofer Mazor, Glenn C Turner, Stijn Cassenaer, Rachel I Wilson, and Gilles Laurent. Oscillations and sparsening of odor representations in the mushroom body. *Science*, 297(5580):359–365, July 2002.
36. Richard H R Hahnloser, Alexay A Kozhevnikov, and Michale S Fee. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902):65–70, September 2002.
37. W B Levy and R A Baxter. Energy efficient neural codes. *Neural Comput.*, 8(3):531–543, April 1996.
38. Hadas E Sloin, Lior Bikovski, Amir Levi, Ortal Amber-Vitos, Tomer Katz, Lidor Spivak, Shirly Someck, Roni Gattegno, Shir Sivroni, Lucas Sjulson, and Eran Stark. Hybrid offspring of C57BL/6J mice exhibit improved properties for neurobehavioral research. May 2022.
39. Eran Stark, Tibor Koos, and György Buzsáki. Diode probes for spatiotemporal optical control of multiple neurons in freely moving animals. *J. Neurophysiol.*, 108(1):349–363, July 2012.
40. Marius Pachitariu, Nicholas Steinmetz, Shabnam Kadir, Matteo Carandini, and Kenneth D Harris. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, page 061481, June 2016.
41. João D Semedo, Amin Zandvakili, Christian K Machens, Byron M Yu, and Adam Kohn. Cortical areas interact through a communication subspace. *Neuron*, 102(1):249–259.e4, April 2019.
42. Daniel Levenstein, György Buzsáki, and John Rinzel. NREM sleep in the rodent neocortex and hippocampus reflects excitable dynamics. *Nat. Commun.*, 10(1):2478, June 2019.

ACKNOWLEDGMENTS

We wish to thank Ruben Coen-Cagli, Adam Kohn, Daniel Levenstein, Michael Long, and Jacob Ratliff for helpful comments on the manuscript. We also wish to thank György Buzsáki, in whose lab the published datasets used in this study were collected. This work was supported by funds from the National Institute on Drug Abuse (DA051608 and DA036657 to L.S.), the Leon Levy Foundation (L.S.), the Brain and Behavior Research Foundation (L.S.), the Whitehall Foundation (L.S.), the McManus Charitable Trust (L.S.), the Feldstein Medical Foundation (L.S.), and Montefiore/Einstein startup funds (L.S.).

AUTHOR CONTRIBUTIONS

L.S. and A.P. conceived the project. E.F.O., S.K., R.B.B., and L.S. designed the experiments. E.F.O. collected all freely-moving electrophysiological data, and S.K. collected all head-fixed Neuropixels data. E.F.O., S.K., T.S.Q., and L.S. performed data analysis, and S.K. and E.F.O. performed computational modeling. L.S. supervised all aspects of the project. E.F.O., S.K., and L.S. wrote the manuscript with contributions from all authors.

COMPETING INTEREST STATEMENT

The authors declare no competing interests.

Methods

All experimental procedures were conducted in accordance with the Institutional Animal Care and Use Committee of Albert Einstein College of Medicine.

Animals and surgery

Experiments were conducted using C57BL/6J x FVB F1 hybrid mice (38). For the visual stimulation task, we chronically implanted a 64-site silicon probe (NeuroNexus) in the primary visual cortex (V1) (AP: 3.4 ML: 2.6 DV: 0.6 from brain surface) under isoflurane anesthesia, as described previously (39). Wires for reference and electrical ground were implanted above the cerebellum, and a copper mesh hat was built to shield the probes. Probes were mounted on microdrives that were used to move the probe farther into V1 for maximizing unit yield, though never beyond DV of 1.0 mm. Animals were housed individually after implantation and allowed to recover for at least one week before experiments. Recordings were performed using the Intan system at 30 kHz. Offline automatic spike sorting was performed using Kilosort 2.0 (1, 40), and all parameters used for sorting are presented in the Kilosort 2.0 wrapper repository in StandardConfig_KS2Wrapper.m. Manual adjustment of Kilosort outputs was performed using Phy. Isolated single units were assigned to putative principal neurons or fast-spiking interneurons based on through-to-peak time of waveforms, a criteria previously used in V1 (24).

We also used data that is publicly available on the Buzsáki laboratory website. We used the Matlab toolbox Buzcode to extract the data used for this manuscript . The methodology for collection of this data was described elsewhere (16, 18, 22–25). From all datasets, we selected only recording sessions where there were location/speed tracking and slow-wave sleep (SWS) states. A table with a brief summary of the dataset used is presented in table 1.

Brain region	Species	Animals	Sessions	Reference
V1	Mice	18	18	(24)
mPFC	Rat	4	82	(16)
TH	Mice	7	51	(18)
LS	Rat	5	45	(25)
HPC/NAc	Mice	6	34	(23)
BLA	Rat	4	39	(22)

Table 1. Summary of datasets used for analysis

Viral vectors

We used an adeno-associated viral vector (AAV) with a CaMKII promoter to express channelrhodopsin-2 in pyramidal cells in primary sensory and motor cortex. The recombinant AAV vector was pseudotyped with AAV5 capsid protein and packaged by Addgene.

Visual stimuli

Natural scene stimuli were the same as used in (1). In brief, 32 images were selected from ethological relevant categories, such as animals and plants. The images chosen were less than 50% uniform background, and with a balance of low and high spatial frequencies. We performed 60 repetitions with 0.5 s duration and an inter-stimulus interval varying from 0.7 to 1 seconds. Visual stimuli were presented on a screen facing the eye contralateral to the side being recorded.

Neuropixels data acquistion

Neuropixels electrodes were used to record extracellularly from neurons in multiple brain areas including V1, hippocampus, thalamus, striatum, and prefrontal cortex in head-fixed mice. On the day of recording, two small craniotomies were made with a dental drill. After recovery, animals were head-fixed on a custom made treadmill. Two neuropixels probes are inserted in each animal's right hemisphere (anterior-posterior (AP): -3.2 mm, medial-lateral (ML): 2.7 mm, depth(D): 3.6 mm, angle: horizontal 60 degrees, medial 15 degrees; AP: -0.2 mm, ML: 3.2 mm, D: 5.6 mm, angle: horizontal 60 degrees, medial 55 degrees; AP: 0.9 mm, ML: 1.1 mm, D: 3.8 mm, angle: horizontal 60 degrees, medial 55 degrees). Each probe was mounted on a rod held by a micromanipulator (uMP-4, Sensapex Inc.) and advanced slowly ($\sim 1\mu$m/s). Electrodes were allowed to settle for 30 min before starting recording. Recordings were made in external reference mode with LFP gain 250 and AP gain 500 using Open-Ephys software. Wires for reference and electrical ground were connected to the cerebellum.

Optogenetic stimulation

In each animal's left hemisphere, the optical fibers (400 µm core diameter, 0.5 NA) were inserted in S1 (anterior-posterior (AP): -1.5 mm, medial-lateral (ML) 1.5 mm, Depth(D): 0.6 mm) and M1 (AP: 1.5 mm, ML 1.7 mm, D: 1.2 mm). Optogenetic stimulation was at 450 nm (Osram PL450B laser diode) with power ranging from 0-6 mW. We use three types of optogenetic stimuli: 25 Hz white noise lasting 1 s, 10 Hz white noise lasting 1 s, and a 10 Hz sinusoid lasting 0.4 s.

Data analysis and exclusion

We used custom MATLAB custom scripts for analysis and plotting. In our analysis we used the following toolboxes: Buzcode, communication subspace (41) for Reduced Rank Regression, and GPML toolbox for Gaussian Process Regression. We used only sessions containing a minimum of 30 neurons and 20 minutes of sleep, unless otherwise stated. When analyzing running speed and face motion prediction, we excluded sessions where the full rank prediction was smaller than a threshold of $R^2 = 0.1$. In the classification of natural visual scenes we excluded low-quality sessions in which the total prediction accuracy was below 30%.

Cross-validated PCA

This method is similar to as described before (10). First, we separate the data into training and test sets (N_{tr} and N_{ts} respectively). The singular vectors V_n are calculated for the training set, and test set data is projected onto those singular vectors to estimate cross-validated scores $U_n = N_{ts}V_n$, with variance calculated from each cvPC score. When estimating the amount of variance, we performed cvPCA in five contiguous folds, i.e. folds that have the same duration but consist of continuous blocks of time rather than randomly selected time

points. We set each cell to contribute equally to the cvPCA by z-scoring them individually. Each state (awake and SWS) is z-scored separately. Since shortening the data in time can change the variance of each neuron, we z-scored the training set separately before PCA, but the test set was only z-scored once with the entire dataset.

To build a null distribution for the eigenspectrum, we shuffled the cell identities of the test set 10,000 times and calculated the variance of the shuffled test set's projection onto the original singular vectors. We defined the confidence interval as the lower $(100 \times p)\%$ and upper $(100 \times (1-p))\%$ values of this null distribution, where $p = \frac{1}{N}$ and N is the number of cells in the data. The confidence interval was used to define where on-/non-/off-manifold subspaces begin and end in the cross-validated eigenspectrum.

Normalized contrastive PCA

To determine the dimensions that are maximally different between SWS and awake states, we developed a new PCA method to find the contrastive dimensions between the two. In order to control for high variance dimensions dominating the contrastive PCs, we normalize the method by their variance. This method finds the eigenvector x that maximizes equation 1.

$$\frac{x^T(B^TB - A^TA)x}{x^T(B^TB + A^TA)x} \quad \textbf{(1)}$$

Where A and B are the neural activity in SWS and awake states.

Reliability index

The reliability index was calculated based on the coefficient of variation from the cross-validated eigenspectrum. For this purpose, we took segments of awake and SWS data of equal duration and separated the data from each state into five contiguous folds. We performed PCA on one fold and projected each other fold in the testing sets onto the training set PCs. Each fold got a turn as the training set, which allowed us to build a confidence interval out of the PCs. Using the standard deviation (σ) and mean (μ) we calculated the Reliability index by the ratio $\frac{\mu}{\sigma}$ which is equivalent to the inverse of the coefficient of variation $\frac{\sigma}{\mu}$.

Gaussian process regression and running speed prediction

We used gaussian process regression (GPR) to predict running speed from the neural activity. Running speed was log-transformed and z-scored before model training and prediction. The parameters used for GPR were mean function zero, covariance function rational quadratic, and exact inference of posterior probability with gaussian likelihood. All the hyperparameters were optimized by the log marginal likelihood.

Face motion extraction and prediction

We record the face of the mouse using a camera with a infrared filter. Infrared LEDs were directed to the face of the mouse. Videos were collected at 30 Hz and aligned with electrophysiology based on digital pulses that triggered frame collection.

To extract the face motion variables we used https://github.com/MouseLand/facemap. For face motion variables used here, we excluded ROIs around the eyes to avoid contamination by the pupil.

Reduced Rank Regression was used for prediction of face motion components by z-scored neural activity. The initial β used for the regression were estimated by Ridge Regression, as described previously (41).

Optostimulation prediction and Partial Least Squares regression

To determine the dimensionality of optostimulation-encoded population activity we used Partial Least Squares (PLS) regression. This method identifies dimensions that maximize the variance explained in the neural data by optostimulation. Similarly to PCA, PLS regression identifies orthogonal dimensions with decaying variance explained in a successive manner. We performed prediction of optostimulation by using the z-scored neural data projected into the identified PLS dimensions and using GPR to predict the z-scored optostimulation. For GPR we used the same parameters as described in the gaussian process regression section. This method was done separately for S1 and M1 optostimulation.

To define the dimensionality of optostimulation in the neural data, we picked the number of PLS dimensions that plateaued the cross-validated R^2. For that we picked the lowest number of dimensions that had prediction performance one s.e.m. away from peak prediction.

Natural visual scenes classification

The prediction of the stimulus identity was done by fitting a linear multiclass support vector machine (SVM); the SVM model fits a one-vs-one multiple binary classification to a total of K(K-1)/2 models. The average firing rate of neurons during each stimulation was z-scored and used as a variable for prediction. This same data was projected in different dimensions, when stated so, for some analyses performed in this article.

R^2 and accuracy

To estimate performance of regression analysis, we use the equation 2 for R^2.

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2} \quad \textbf{(2)}$$

Where y is the original variable, $\hat{y}$ is the model predicted and $\bar{y}$ is the mean of the original variable. In this R^2 equation, the model is compared to a constant model (the mean of the original variable). If the prediction is worse than the constant model, the R^2 can be negative, which is comparable to $R^2 = 0$.

The performance of classification for natural scenes was calculated using the accuracy equation 3.

$$\text{Accuracy} = \frac{(\text{True Positive} + \text{True Negative})}{(\text{All Samples})} \quad \textbf{(3)}$$

Prediction index and random projections

There was a different number of dimensions in each subspace (on-/non-/off-manifold), and that can influence the prediction quality of the variables we tested. To control for that, we generated random orthogonal vectors with matching numbers of dimensions to the on-/non-/off-manifold subspaces. We projected the data in both the random subspace and original subspaces to make predictions of running speed, face motion, and natural scenes. Using the R^2 of these two subspace predictions we calculated the prediction as indicated in equation 4.

$$\text{Prediction Index} = \frac{R_s^2 - R_{rp}^2}{R_s^2 - R_{rp}^2} \tag{4}$$

Where R_s^2 is the R^2 of the subspace and R_{rp}^2 the R^2 random projections of matching dimensionality. A positive prediction index means the prediction was better by the original dimensions, and a negative index means that random projections had better prediction. Because the R^2 cannot be negative for this index calculation, we zeroed every R^2 that was negative (see R^2 and accuracy section).

Single-neuron population coupling

Population coupling was calculated as described before (8); in brief, we summed the neural activity of N-1 neurons binned and zero-centered, then took the dot product of the left-out neuron's activity with that of the summed population activity. The resulting number is normalized by the L2-norm of the left-out neuron's activity. We normalized the population coupling to the range [0,1]. Unlike Okun et al., we did not smooth the data prior to the population coupling calculation.

Population sparseness

We used the Gini coefficient of the absolute loadings to estimate the population sparseness. This was calculated in each cvPC separately. To calculate the Gini coefficient, we first binned the absolute value from cvPCs loadings into 30 equally spaced bins, then estimated scores S_n by multiplying the amount of neurons in the bin to the loadings value the bin represents. With $f(y_i)$ representing the fraction of the population with those loadings, we estimated the Gini coefficient by equation 5.

$$\text{Gini coefficient} = 1 - \frac{\sum_{i=1}^{n} (f(y_i)(S_{i-1} + Si))}{S_n} \tag{5}$$

The coefficient varies from 0 to 1, where 0 is perfect equality (i.e. equal participation of all the neurons in a cvPC) and 1 is maximal inequality (only one neuron participates in the cvPC).

Overall firing rate fluctuation

To control for population firing rate fluctuation in our results, we removed it from the data when applicable. For that we estimate the overall firing rate fluctuation as the change in activity in a $[\mathbf{1}]$ vector of the same number of rows as the number of neurons. This vector is normalized by its $L2$-norm and used to project the z-scored neural data N (organized with time bins in rows and neurons in columns) and estimate the overall firing rate fluctuation in time FR_t. We then removed this from the data by subtracting this projection, as in $N - \text{FR}_t \cdot [\mathbf{1}]^{\mathbf{T}}$.

Up and down state detection

To control for up and down transitions in SWS affecting our estimates of population activity, we detected up/down states and restricted the PCA to up states alone. For detecting these states, we used the same methodology as described previously (42). In brief, during periods of SWS, the LFP delta (0.5-8 Hz) and gamma bands (30-600 Hz) are used for thresholding and detection of up (high gamma power) and down states (delta peak and low gamma power). Down states are then further restricted to periods where spiking activity is below threshold.

High firing rate spikes subsampled

To control for high firing rate neurons, we subsampled their spikes to match the number of spikes in the low firing rate neurons. We sorted the neurons according to their firing rate and subsampled the top half of neurons to match the bottom half. We randomized which top neuron would match which bottom neuron so that the first top neuron did not always match the first bottom neuron.

Single neuron subspace participation

To determine the participation of a neuron to a specific subspace, we calculated the average absolute loading of that neuron in all the cvPCs that span that specific subspace.

Simulation

We built a simple neural network to investigate how the overlap of both movement-evoked activity and visual stimuli influences the decoding of the visual stimulus identity in relation to the signal-to-noise ratio. The network has 100 principal neurons, and each neuron's firing patterns are Poisson-distributed with uncorrelated white noise added at each time step. Neurons were not connected to each other but received inputs encoding both movement "noise" and visual stimulus "signal", each with their own weight matrix. There are two distinct time epochs: a sleep epoch receiving only the movement-evoked white noise input, and a visual stimulus epoch receiving both the noise input and the visual stimulus signal input. For both epochs we used 100 ms time steps, and the total duration of each epoch is 50 s. During the sleep epoch, the white noise temporal profile was fed into the 60 principal neurons with a weight vector of normal distribution loadings in the range [0,1]. During the visual stimulus epoch, principal neurons also received the same white noise input as during the sleep epoch. In addition, ten different sparse input weight vectors representing ten visual stimuli were fed into the principal neurons with loadings in the range [-1,1]. 30 out of 100 neurons received visual stimuli, each with its own weight vector. 5 out of 30 neurons received all 10 visual stimulus inputs, and different combinations of 5 out of 25 neurons were stimulated with each visual stimulus input. Thus 10 different combinations of neurons are tuned for individual stimuli. In the simulations, two parameters were explored: the signal-to-noise ratio and the

signal/noise population overlap. The signal-to-noise ratio is varied by adjusting the amplitude of the white noise stimulus. The signal/noise overlap is defined by selecting 30 out of 100 neurons to receive the visual stimulus, and those neurons vary from having the highest to lowest movement-evoked activity. To test the network's performance encoding visual stimuli, prediction accuracy was computed using the accuracy formula above.

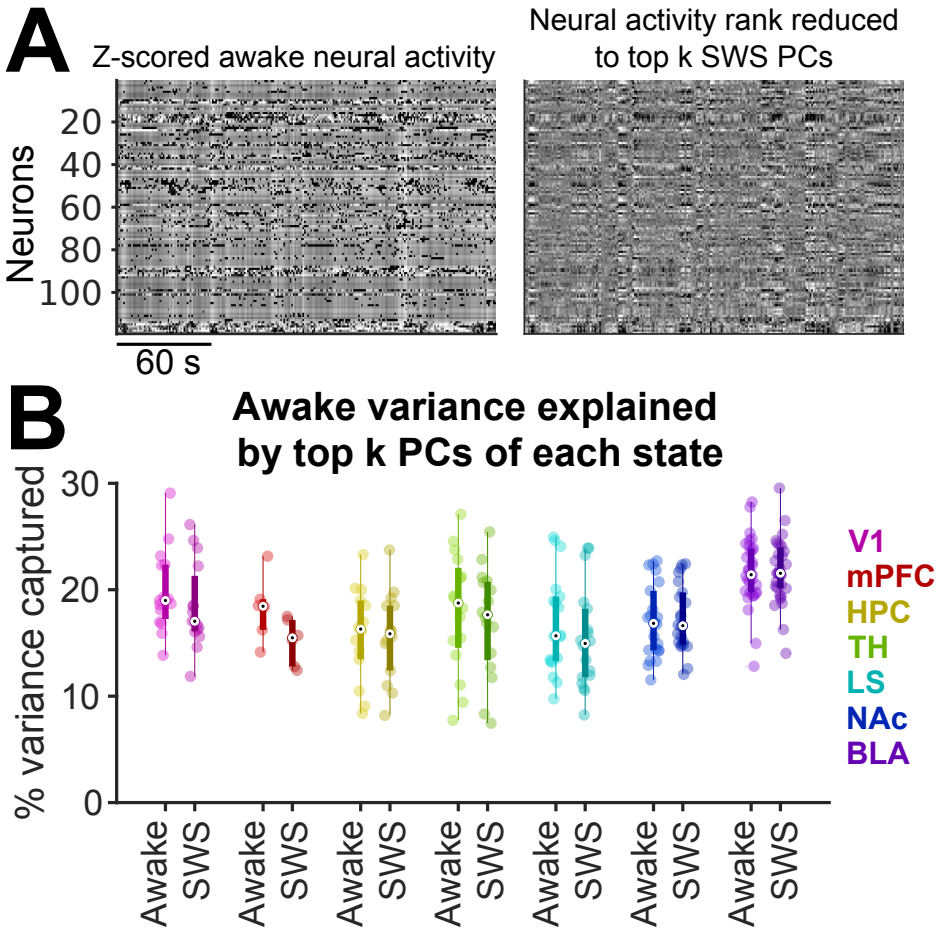

Supplementary Fig. 1. SWS and awake cvPCs explain a similar amount of variance in awake population activity. (**A**) SWS cvPCs capture much of the population structure of awake activity. (**B**) The top k cvPCs of awake and SWS activity (defined as the cvPCs with variance above chance) explain similar amounts of variance in the awake test set data. This suggests that low-dimensional population structure is conserved across SWS and wakefulness.

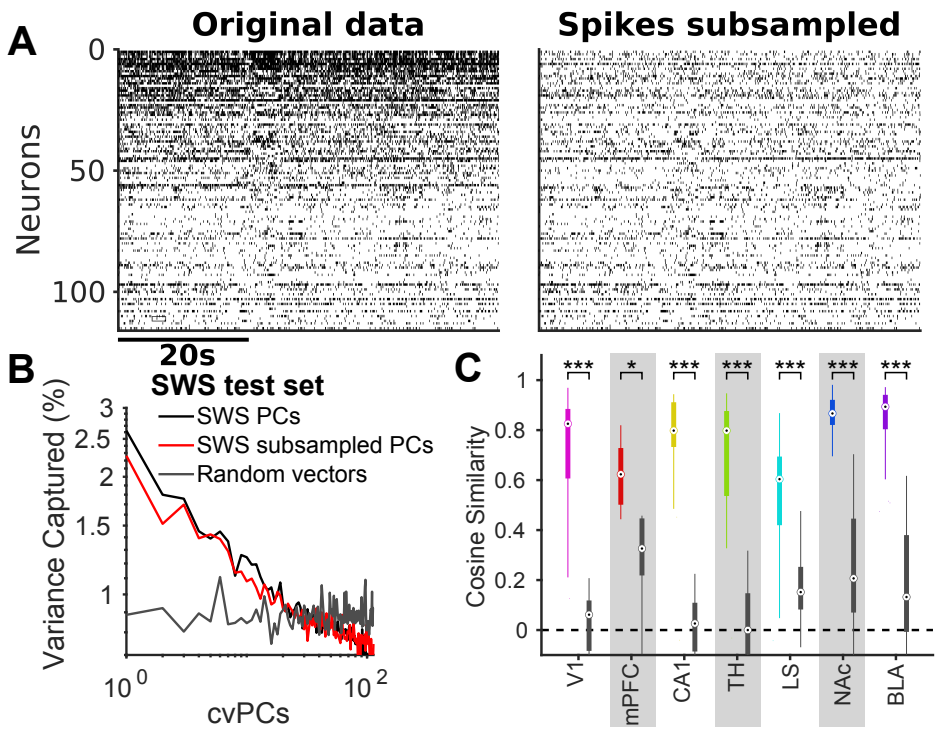

Supplementary Fig. 2. The population activity structure observed during SWS is not explained by high firing rate neurons dominating on-manifold dimensions. (**A**) To control for higher firing rate neurons being overrepresented in on-manifold dimensions, we subsampled the spikes of the top half of high-firing-rate neurons to match the firing rates of the bottom half. (**B**) Example eigenspectrum of original SWS test set projected onto PCs from SWS training set, subsampled SWS training set, and random vectors. (**C**) The SWS PC eigenspectrum is more similar to the subsampled SWS PC eigenspectrum (colored boxplots) than random vectors (grey boxplots) for all brain areas tested. [*p<0.05, ***p<0.001]

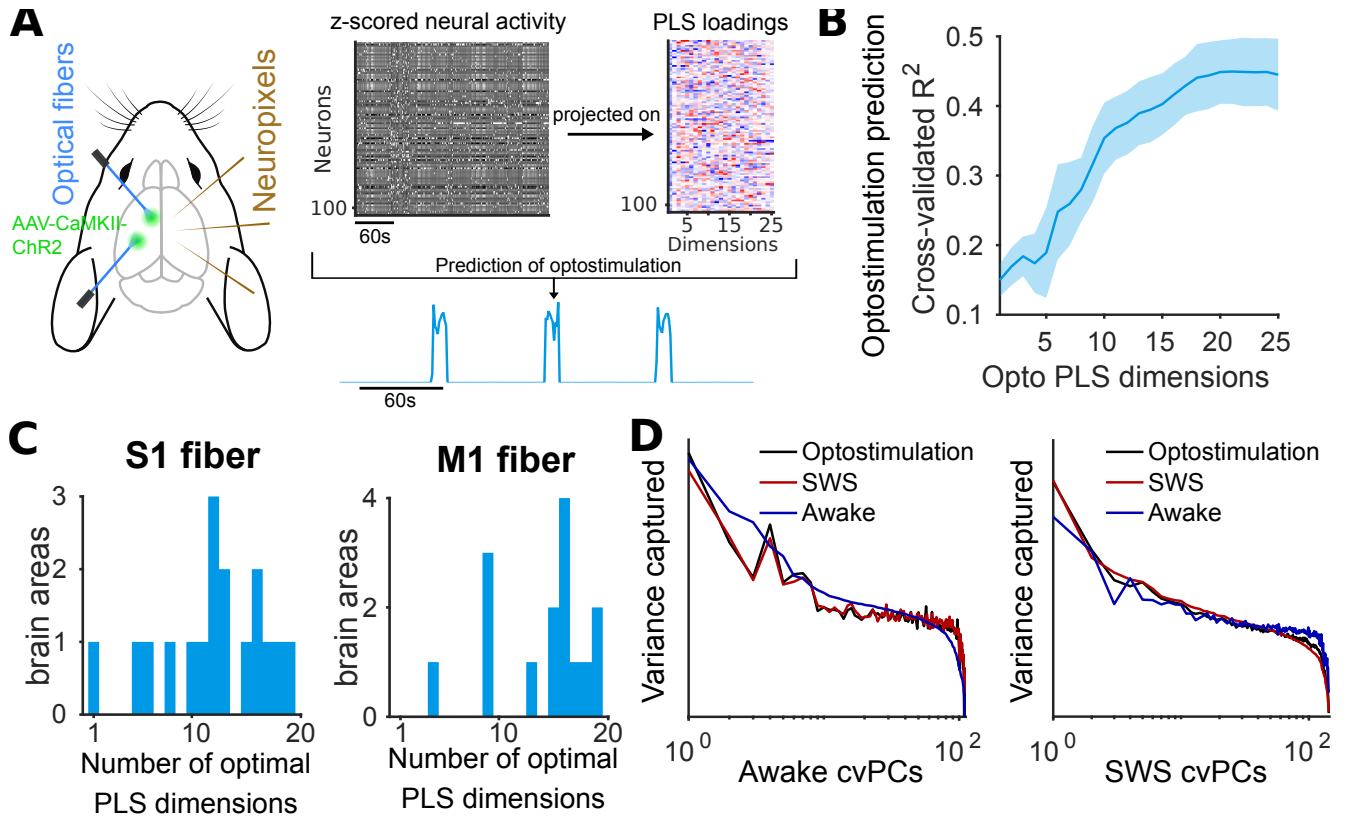

Supplementary Fig. 3. Optogenetic stimulation evokes multidimensional activity patterns similar to those observed during SWS. (**A**) Using partial least squares regression (PLS), we found neuronal activity dimensions that best predicted the waveform of diffuse optogenetic stimulation applied to the contralateral hemisphere. (**B**) Many PLS dimensions are required to predict the one-dimensional optogenetic stimulus. (**C**) S1 and M1 optostimulation generate multidimensional representations in multiple brain areas (S1: 12 ± 1.2, n = 17 brain areas; M1: 14.1 ± 1.2, n = 15 brain areas; mean ± s.e.m.). (**D**) Optostimulation during awake states evokes population activity that is more similar to SWS than awake states. The left eigenspectrum example shows the projection of test set data onto awake cvPCs, and the right eigenspectrum shows the projection onto SWS cvPCs. In both cases, optostimulation is closer to SWS than awake (group statistics shown in Fig. 1H).

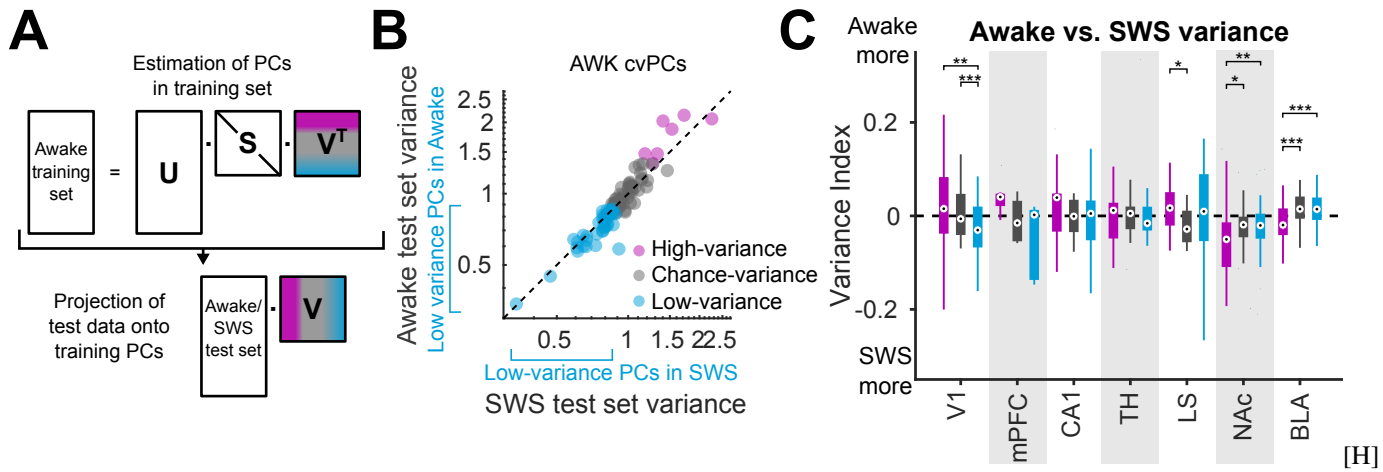

Supplementary Fig. 4. SWS activity does not show increased variance in dimensions that exhibit the lowest variance during wakefulness. (**A**) PCs were estimated from the awake training set, then awake and SWS test sets were projected onto these PCs. We built a null distribution using the awake training set to identify cvPCs that explain more, equal, or less variance than chance. (**B**) The SWS test set projected onto awake cvPCs does not show an increased variance in the low-variance dimensions, as is observed when an awake test set is projected onto SWS cvPCs (Fig. 2C-D). This suggests that the increased off-manifold activity observed during wakefulness is not a trivial consequence of the awake and SWS eigenspectra being different. (**C**) With the sole exception of the BLA, we did not observe any increased variance in the lowest-variance PC subspace. [*p<0.05, **p<0.01, ***p<0.001].

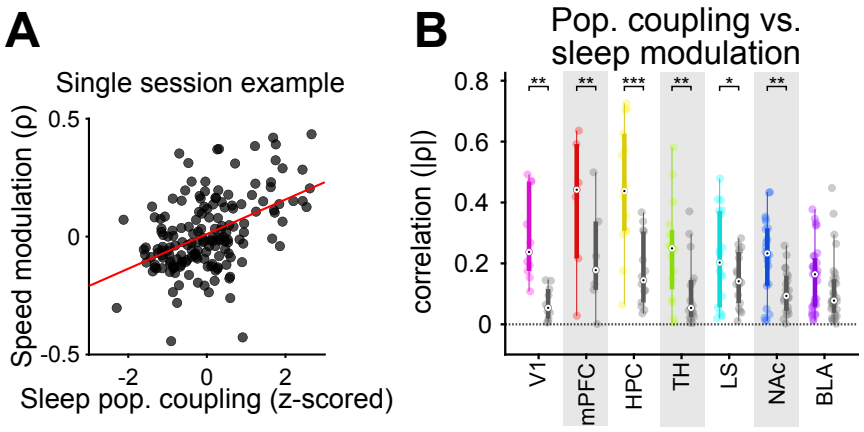

Supplementary Fig. 5. Population coupling during sleep is correlated with running speed speed modulation during wakefulness. (**A**) An example from a single recording session showing that sleep population coupling and awake running speed modulation are correlated. Each dot is a neuron. (**B**) Summary of several brain areas. The absolute correlation coefficient for sleep population coupling and speed modulation for each brain area (colored boxplot) is shown alongside the result from a null distribution where we shift running speed in time (gray boxplot). [*p<0.05, **p<0.01, ***p<0.001]

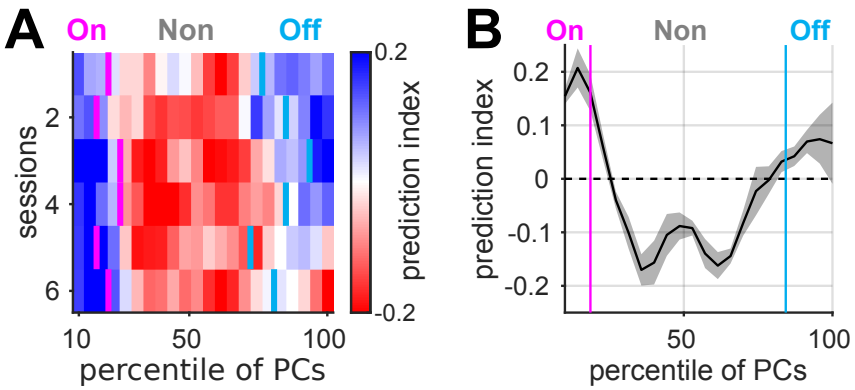

Supplementary Fig. 6. SWS on- and off-manifold boundaries roughly match changes in the encoding of natural scenes in V1. (**A**) On- and off-manifold PCs predict the identity of natural scene stimuli better than non-manifold PCs. (**B**) The averaged prediction index across all sessions shows that changes in decoding performance roughly match the on- and off-manifold boundaries, represented by magenta and cyan lines, respectively.

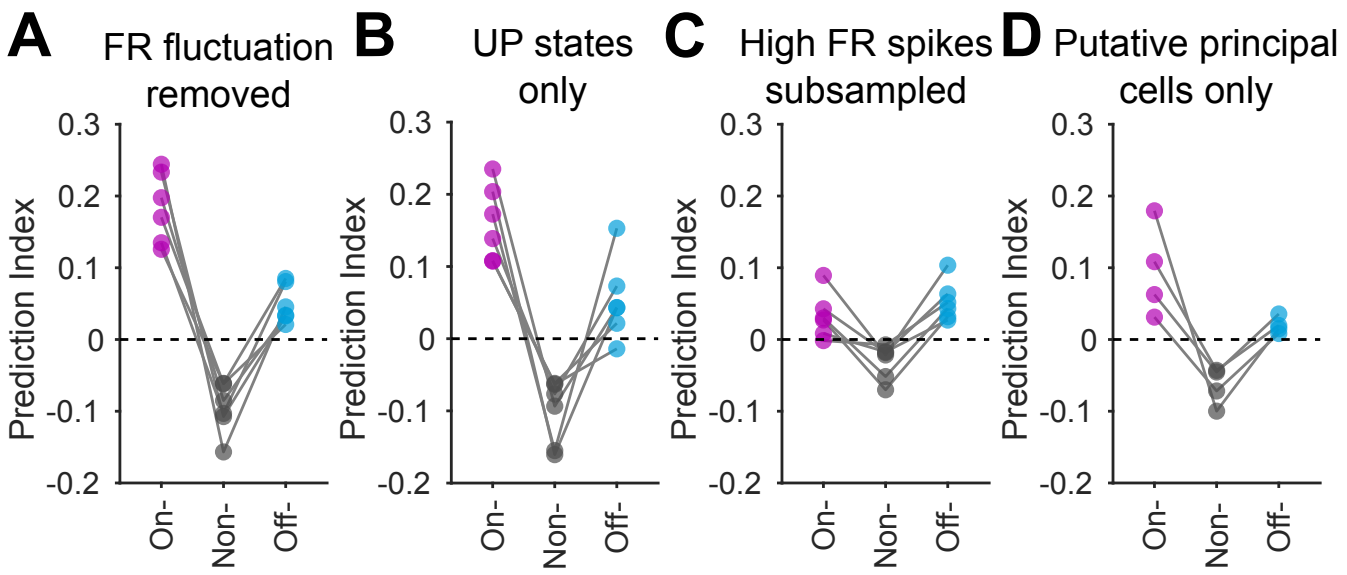

Supplementary Fig. 7. Off-manifold encoding of natural visual scenes is not explained by nonspecific effects. (**A**) Removal of firing rate fluctuation from each neuron in SWS. (**B**) Estimating SWS cvPCs using UP states periods only. (**C**) Subsampling the spikes from top half firing rate neurons to match the amount of spikes of bottom half. (**D**) Using solely putative principal cells.

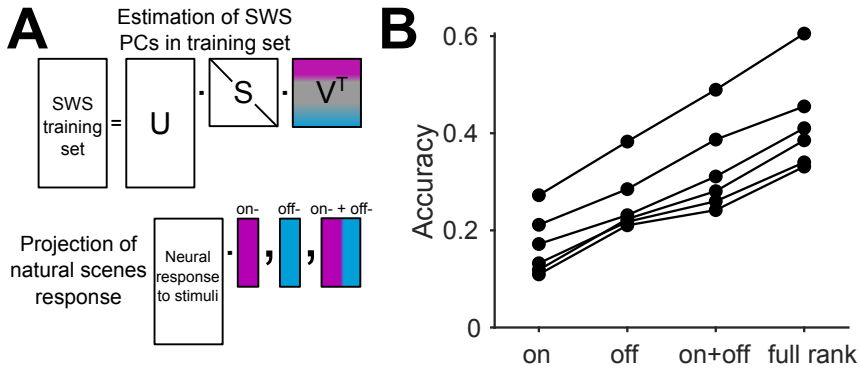

Supplementary Fig. 8. The on- and off-manifold subspaces encode complementary information about natural scene stimuli. (**A**) To test whether the on- and off-manifold subspaces encode redundant information, we tested the performance of a classifier in predicting stimulus identity from neural activity in the on-manifold subspace, the off-manifold subspace, or both. (**B**) Prediction performance is higher for both subspaces combined. Full rank prediction displayed for comparison.

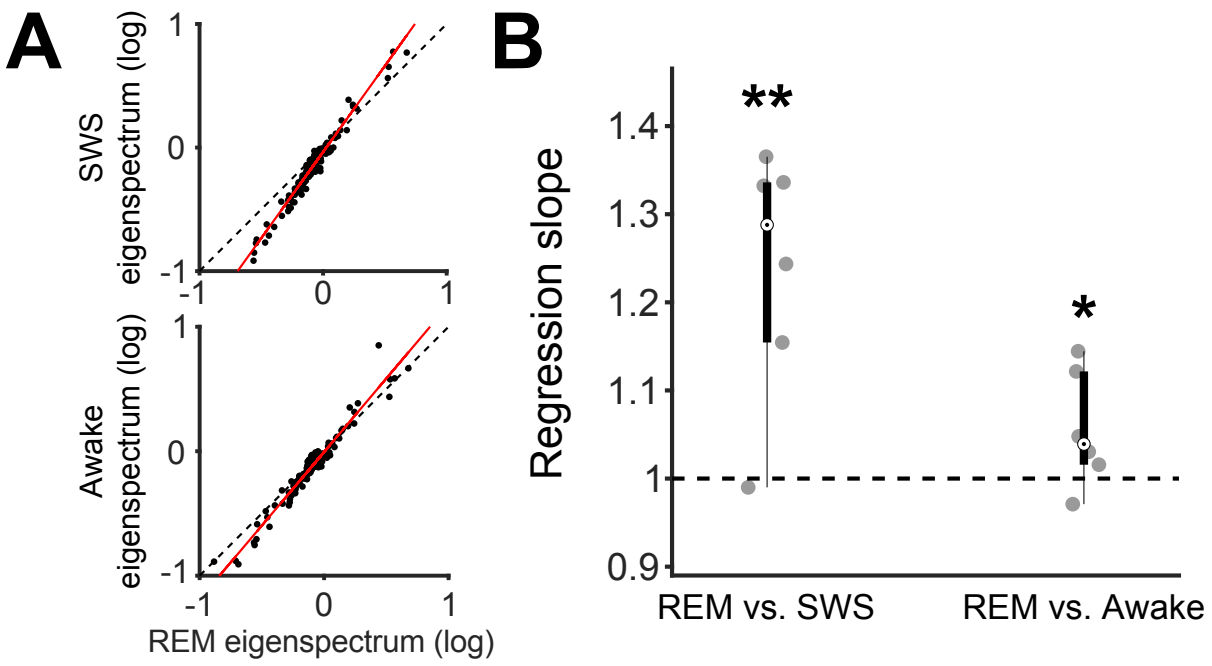

Supplementary Fig. 9. REM sleep is higher dimensional than SWS and wakefulness. (**A**) The eigenspectrum decays more slowly in REM, indicating higher-dimensional population activity than in SWS or wakefulness. (**B**) Summary of several recording sessions from V1 (SWS: p = 0.005, 1.24 ± 0.06, n=6 recordings; Awake: p=0.048, 1.06±0.03, n=6 recordings; mean ± s.e.m.).

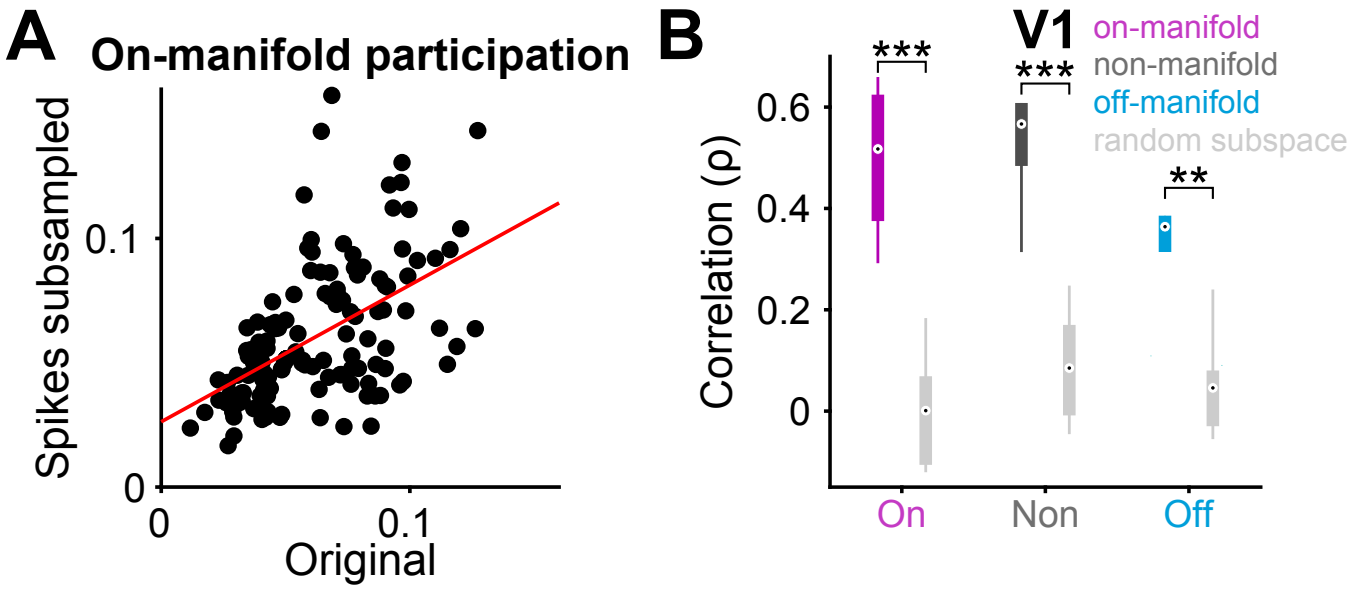

Supplementary Fig. 10. Subsampling spikes in high firing rate neurons does not drastically change subspace participation. (**A**) In an example session, the on-manifold participation of each neuron is similar before and after subsampling spike trains to control for firing rate. (**B**) Correlation coefficients for subspace participation in original and subsampled spike trains for all V1 recording sessions: on-manifold (magenta), non-manifold (dark gray), and off-manifold (cyan). Light gray boxplots are for randomly chosen vectors. [**p<0.01, ***p<0.001]

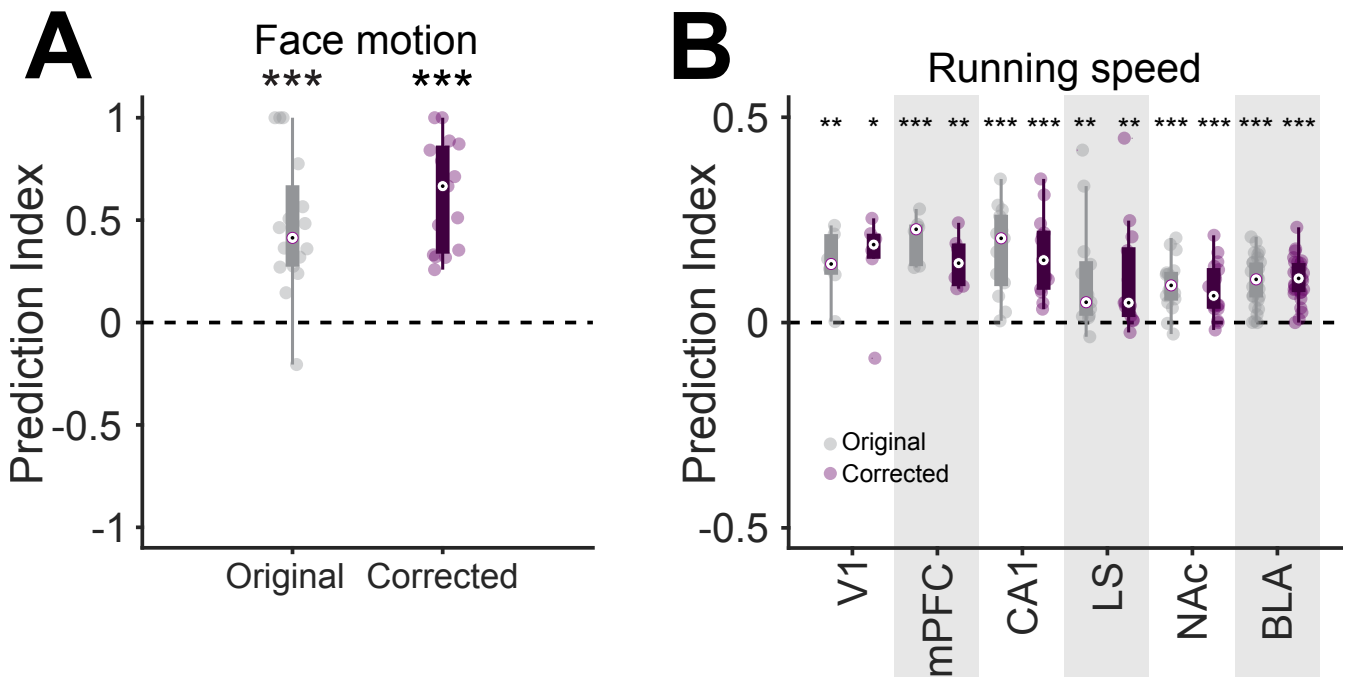

Supplementary Fig. 11. Controlling for the firing rate fluctuation does not change on-manifold prediction of movements. (**A**) Removing overall firing rate fluctuation of each neuron does not change on-manifold predictability of face motion. (**B**) Same as **A** but for running speed. [*p<0.05, **p<0.01, ***p<0.001]

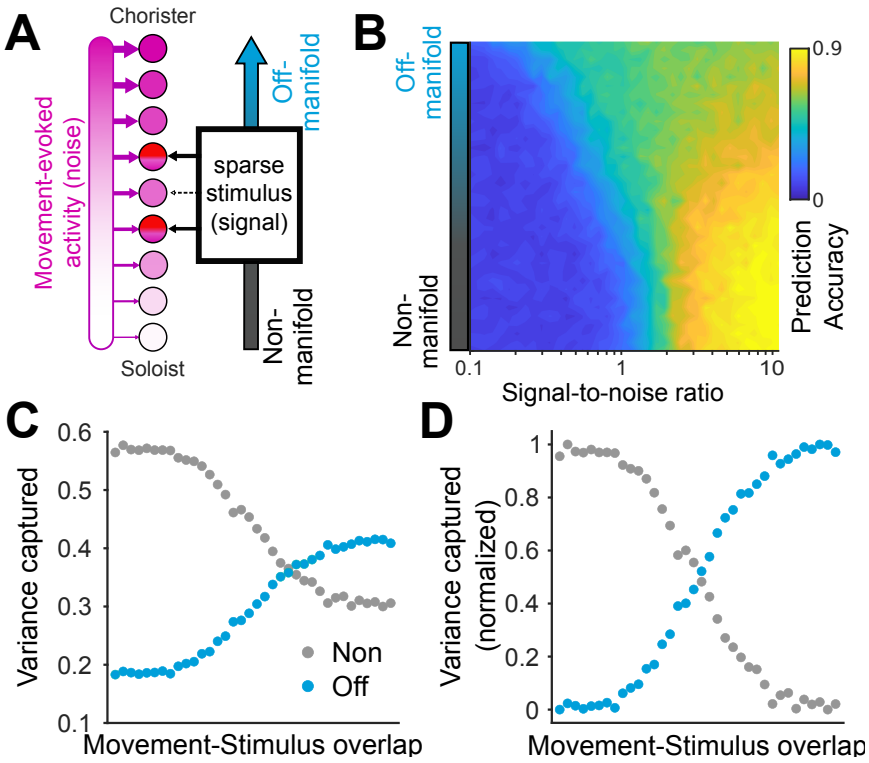

Supplementary Fig. 12. Stimulus-evoked activity in chorister neurons creates off-manifold encoding. (**A**) In our model, the stimulus provides inputs to a sparse subset of neurons that varies from low to high population coupling. (**B**) At low signal-to-noise ratios, off-manifold coding becomes advantageous. This is the same plot shown in Fig. 4F, but without normalization. (**C**) The proportion of stimulus-related variance captured by the non- and off-manifold subspaces (grey and cyan, respectively). As the sparse stimulus goes from overlapping with soloists to choristers, the coding of the stimulus switches from non-manifold to off-manifold. (**D**) Same as **C** but normalized to range.