

Tasks collapse the intrinsic dimensionality of activity in non-selective cortex

Erica L. Busch (erica.busch@yale.edu)

Department of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520 USA

Tristan S. Yates (tristan.yates@yale.edu)

Department of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520 USA

Nicholas B. Turk-Browne (nicholas.turk-browne@yale.edu)

Department of Psychology & Wu Tsai Institute Yale University, 100 College St., New Haven, CT 06520 USA

Abstract

Human brain activity collected from fMRI is measured in high voxel dimensions. However, the intrinsic dimensionality of neural population activity is often much lower than this measured dimensionality. How intrinsic dimensionality is modulated by stimuli and tasks remains unknown. Here, we explore this question during viewing of two naturalistic movies and during rest. We used T-PHATE, a nonlinear manifold learning approach designed for high-dimensional, timeseries data (Busch et al., 2023), to estimate the intrinsic dimensionality of fMRI activity across the brain using searchlights. Visually responsive brain regions showed high dimensionality during both movie viewing and rest. However, other brain regions had lower dimensionality during movie than rest. Thus, movie viewing appears to collapse the dimensionality of non-selective brain regions, whereas in a baseline state there is relatively unconstrained dimensionality across the brain.

Keywords: Manifold learning; Naturalistic stimuli; Dimensionality reduction

Introduction

The brain encodes rich, dynamic information about the world in distributed neural population codes that communicate via both local and global patterns of activity. Modern neuroimaging studies collect high-throughput samples of neural activity across experimental conditions to understand where and how different stimuli, tasks, and behaviors are processed in the human brain.

Direct neural recordings in nonhuman primates have shown that the measured dimensionality of neural activity generally exceeds the activity's intrinsic dimensionality — that is, the degrees of freedom required to describe the overall signal of the population — as there is substantial covariance among neural units tuned toward particular stimuli or tasks (Cunningham & Yu, 2014). The intrinsic dimensionality of neural population activity can provide insight into the complexity of the neural processing underlying a given task (Altan, Solla, Miller, & Perreault, 2021; Jazayeri & Ostojic, 2021). For instance, intrinsic dimensionality can be determined by incoming stimuli (Churchland et al., 2010; Mazzucato, Fontanini, & La Camera, 2016), ongoing behaviors (Stringer et al., 2019), and complex latent variables like experience and expectations (Jazayeri & Ostojic, 2021).

How is the intrinsic dimensionality of human brain activity altered by engaging in a task? In the past, estimating the intrinsic dimensionality of single-subject, task-based fMRI, such as during naturalistic movie viewing, has been hampered by methods unsuited for disentangling spatio-temporal signal and noise in the data. We use a nonlinear manifold learning approach called T-PHATE to estimate and compare intrinsic dimensionality across brain regions and between movie viewing and rest.

Materials and methods

Movie dataset We analyzed fMRI data from 12 subjects (7 female, age 18 to 32 years) who viewed two short, silent movies (“Aeronaut” and “Mickey”) while having their brains scanned. These same subjects also completed a short rest-fixation scan. More details about data collection and preprocessing are described in (Yates, Ellis, & Turk-Browne, 2023).

Reliability of neural responses To identify brain areas that reliably responded to the movies across subjects (Nastase, Gazzola, Hasson, & Keysers, 2019), we used a leave-one-subject-out intersubject correlation (ISC) analysis in volumetric searchlights (radius = 5) across the entire brain.

Manifold learning We recently introduced T-PHATE (*temporal potential of heat diffusion for affinity-based transition embedding*), a nonlinear manifold learning method designed to learn low-dimensional embeddings of high-dimensional, spatio-temporally noisy timeseries data (Busch et al., 2023). The T-PHATE algorithm learns a multi-view manifold: one view models the data's time-varying properties and the other view learns the data's geometry via the PHATE manifold learning method (Moon et al., 2019). These views are then combined into a single T-PHATE diffusion operator, which defines a manifold based on both dynamic and geometric data properties.

Intrinsic dimensionality We estimated intrinsic dimensionality (ID) over the T-PHATE diffusion operator. In the T-PHATE algorithm, the diffusion operator is powered by t , a parameter that represents the intrinsic dimensionality of the data. t provides a trade-off between encoding of local versus global information in the T-PHATE embedding space, and is optimized by minimizing the von Neumann entropy of the diffusion operator as a function of various values of t . This provides a measure of the number of significant eigenvalues of the diffusion operator. For more information about this calculation,

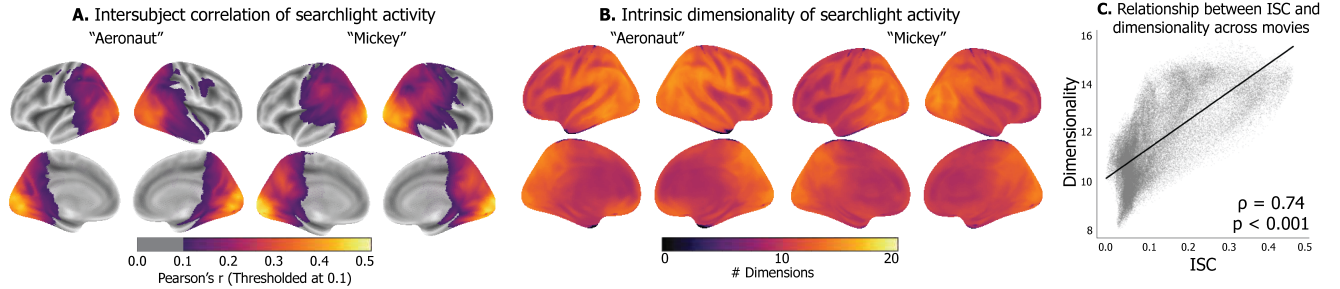


Figure 1: (A) Searchlight ISC was computed with leave-one-out cross-validation and presented as the average across subjects, thresholded arbitrarily at $r = 0.1$. (B) Intrinsic dimensionality for each searchlight pattern was computed with T-PHATE at the subject level, and presented as the group average across subjects. (C) Spatial maps of ISC and dimensionality calculated from independent data show a strong positive correlation. Each point represents a single searchlight.

see (Moon et al., 2019).

To probe how task conditions alter the ID of brain activity patterns across the brain, we applied T-PHATE in searchlights as done for the ISC analysis above. All searchlights had 1331 voxels, which places an upper limit on the dimensionality of a searchlight pattern. All analyses were performed independently for each subject and task.

For comparison, we repeated our analyses by estimating ID with two alternative methods: vanilla PHATE (which excludes the explicit temporal view) and PCA (retaining the number of components that capture 90% variance).

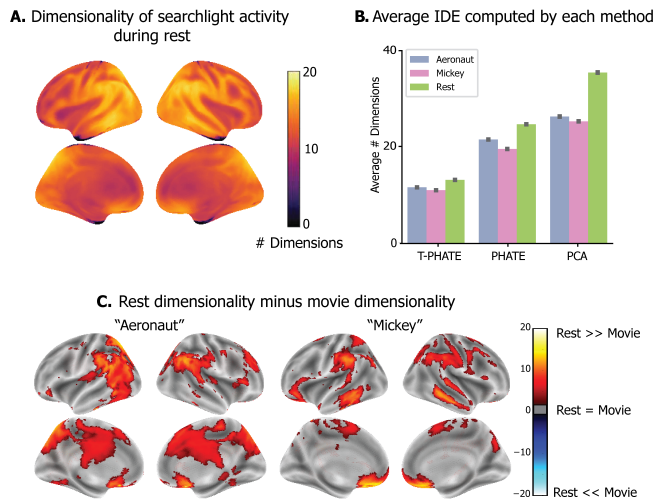


Figure 2: (A) T-PHATE ID of each searchlight during rest. (B) Average ID with T-PHATE, PHATE, and PCA across subjects and searchlights, for each task. (C) Difference in T-PHATE ID for rest minus movie tasks, thresholded at $q_{FDR} < 0.05$.

Results

In both “Aeronaut” and “Mickey” movies, searchlight activity patterns showed the highest ISC (Figure 1A) and ID (Figure 2B) in visual regions, broadly defined. This suggests that, during movie viewing, ISC and ID are both selective to regions that are engaged by the task (in this case a silent movie). We

tested their relationship explicitly by correlating ISC and ID over searchlights across the entire brain. We performed this analysis between movies — e.g., $\text{corr}(\text{Aeronaut ID}, \text{Mickey ISC})$ — to ensure the independence of the measures and understand the generalizability of their relationship. There was a robust positive correlation between ISC and ID during movie viewing (Spearman's $\rho = 0.74$, $p < 0.001$) (Figure 1C).

Overall, ID was higher across the brain during rest vs. movie-viewing (Figure 2A). This pattern replicated with PHATE and PCA (Figure 2B). Manifold learning methods tend to inflate estimates of ID in the presence of noise or data non-linearities, if not properly modeled (Altan et al., 2021), which would explain why both PHATE and PCA show higher absolute ID than T-PHATE.

To quantify where in the brain dimensionality differs between tasks, we subtracted the ID maps for each of the movies from rest for every subject, resulting in 2 sets of 12 difference maps, respectively. The reliability of these differences was assessed with a one-sample t-test across subjects within movie, thresholded at $q_{FDR} < 0.05$ (Figure 2C). The overall higher dimensionality of rest was driven primarily by regions outside of visual cortex (which showed similarly high dimensionality in both rest and movie conditions). These mostly non-visual regions showing reduced dimensionality during movies included ventromedial prefrontal cortex, temporoparietal junction, and cingulate cortex.

In sum, the naturalistic task of movie viewing drives reliable, high-dimensional fMRI activity in visual regions, but may dampen the dimensionality of other regions that are unconstrained and high-dimensional at rest. We are extending this work to test whether this pattern of collapsed dimensionality in non-selective cortex generalizes to the auditory modality.

Acknowledgments

E.L.B. was supported by a NSF Graduate Research Fellowship (award no. 2139841). T.S.Y. was supported by a NSF Graduate Research Fellowship (award no. 1752134). N.B.T.B. was supported by the Canadian Institute for Advanced Research.

References

- Altan, E., Solla, S. A., Miller, L. E., & Perreault, E. J. (2021, November). Estimating the dimensionality of the manifold underlying multi-electrode neural recordings. *PLOS Computational Biology*, 17(11). doi: 10.1371/journal.pcbi.1008591
- Busch, E. L., Huang, J., Benz, A., Wallenstein, T., Lajoie, G., Wolf, G., . . . Turk-Browne, N. B. (2023). Multi-view manifold learning of human brain-state trajectories. *Nature Computational Science*, 1–14. doi: 10.1038/s43588-023-00419-0
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., . . . Shenoy, K. V. (2010, March). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369–378. doi: 10.1038/nn.2501
- Cunningham, J. P., & Yu, B. M. (2014, November). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11), 1500–1509. doi: 10.1038/nn.3776
- Jazayeri, M., & Ostojic, S. (2021, October). Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. *Current Opinion in Neurobiology*, 70, 113–120. doi: 10.1016/j.conb.2021.08.002
- Mazzucato, L., Fontanini, A., & La Camera, G. (2016). Stimuli Reduce the Dimensionality of Cortical Activity. *Frontiers in Systems Neuroscience*, 10.
- Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., . . . Krishnaswamy, S. (2019). Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, 37(12), 1482–1492. doi: 10.1038/s41587-019-0336-3
- Nastase, S. A., Gazzola, V., Hasson, U., & Keysers, C. (2019). Measuring shared responses across subjects using intersubject correlation. *Social Cognitive and Affective Neuroscience*, 14(6), 667–685.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019, April). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437). doi: 10.1126/science.aav7893
- Yates, T. S., Ellis, C. T., & Turk-Browne, N. B. (2023). *Functional networks in the infant brain during sleep and wake states*. bioRxiv. doi: 10.1101/2023.02.15.528718