Intrinsic dimensionality of brain activity manifolds across tasks and development

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

To investigate cognitive processing in the brain, researchers often focus on the localization or amplitude of fMRI activity. Intrinsic dimensionality – the degrees of freedom needed to characterize complex signals – may offer additional insights into where and how computations occur across the brain. Here, we investigated how tasks and development alter the intrinsic dimensionality of brain activity. We used a nonlinear manifold learning method to estimate the intrinsic dimensionality of fMRI activity in whole-brain searchlights from infant (3–24 months), child (3–12 years), and adult datasets during resting and naturalistic fMRI tasks (i.e., movies or narratives). These analyses revealed two principles of intrinsic dimensionality in the human brain: task processing reduces the intrinsic dimensionality of task-irrelevant brain regions relative to rest and development increases the intrinsic dimensionality of task-relevant brain regions. These findings hold implications for understanding information processing in the brain and studying the development of adult-like brain functioning.

Introduction

The human brain computes over and encodes information in distributed patterns of activity. Modern experiments capture these processes in high-throughput, whole-brain activity patterns while heterogeneous participants engage in different experimental conditions, resulting in high-dimensional and complex datasets. Yet, a growing body of research from neural recordings in animals suggests that the neural signals that emerge from brain recordings and code for behavior are lower dimensional – that is, their measured dimensionality (e.g., the number of fMRI voxels) generally vastly exceeds the activity's *intrinsic dimensionality*. Lower intrinsic dimensionality (ID) stems from the substantial covariance among neural units tuned toward particular stimuli or tasks (Cunningham & Yu, 2014). Relative changes in ID can provide insight into the complexity of the neural processing underlying a task (Altan, Solla, Miller, & Perreault, 2021; Jazayeri & Ostojic, 2021). Further, it can inform how tuned the neural population is toward a stimulus (Mazzucato, Fontanini, & La Camera, 2016), and how relevant it is for a given behavior (Stringer et al., 2019; Jazayeri & Ostojic, 2021).

Here we apply the concept of ID to human fMRI datasets to answer two questions: (1) How is the ID of activity across the human brain altered by tasks? (2) How do these task-related dynamics in ID develop from infancy to adulthood?

Materials and methods

Datasets

Narratives We used a subset of the *Narratives* dataset including adults ($N = 45$, 18–53 years) who listened to four audio narratives (stories from "The Moth" podcast: "Black", "Bronx", "Forgot", and "PiemanPNI"; each 7–13.3 minutes) while in the fMRI scanner (Nastase et al., 2021).

Rest/Movie The *Rest/Movie* dataset includes awake infants ($N = 26$, 3.9–24 months) and adults ($N = 12$; 18–32 years) watching a silent cartoon movie ("Aeronaut"; 3 minutes) while in the fMRI scanner. It also includes data from sleeping infants ($N = 20$, 3.9–24 months; 2–5 minutes) and resting adults (same sample as "Aeronaut"; 5 minutes) (Yates, Ellis, & Turk-Browne, 2023).

Partly Cloudy The *Partly Cloudy* dataset includes children ($N = 122$, 3–12 years) and adults ($N = 33$, 18–39 years) watching a silent cartoon movie (Disney Pixar's "Partly Cloudy"; 5.6 minutes) while in the fMRI scanner (Richardson, Lisandrelli, Riobueno-Naylor, & Saxe, 2018).

Analyses

Reliability of neural responses To identify brain areas that responded reliably to the stimuli 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 brain.

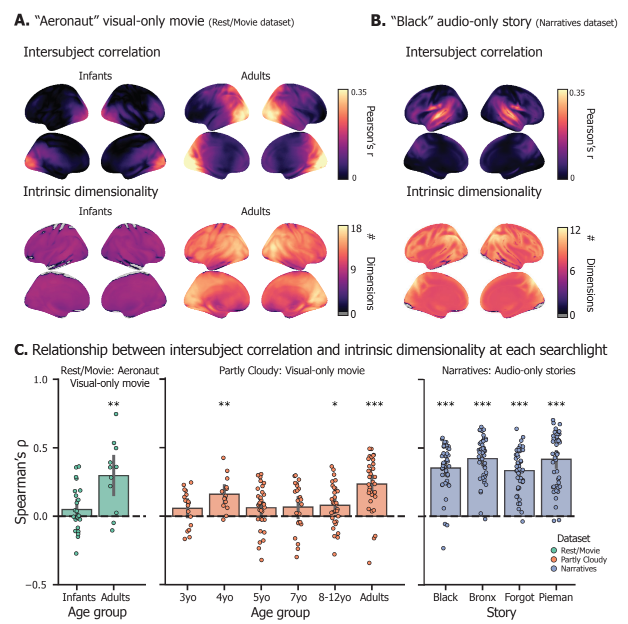

Figure 1: Relationship between task processing and ID. (A) Average intersubject correlation and intrinsic dimensionality in infants and adults watching a silent movie (Rest/Movie dataset). (B) Same as (A), but for one of the audio-only stories in the Narratives dataset. (C) Within-participant analysis of the relationship between intersubject correlation and intrinsic dimensionality at each brain searchlight. Dots are participants, bars are average across participants within age group (left, middle) or stimulus (right). $*p < 0.05$, $**p < 0.01$, $***p < 0.001$

Intrinsic dimensionality We estimated intrinsic dimensionality (ID) using T-PHATE (temporal potential of heat diffusion for affinity-based transition embedding), a nonlinear manifold learning method designed to learn latent embeddings of high-dimensional fMRI timeseries data (Busch et al., 2023). As with the ISC analysis, we estimated ID at every searchlight in the brain, within each participant's data. This resulted in one value for each location in the brain. The T-PHATE diffusion operator captures nonlinear pairwise affinities between the fMRI activity of a given set of voxels at each time point. ID is estimated as the number of significant eigenvalues of this diffusion operator. Dimensionality estimated with T-PHATE resemble estimates by linear methods (e.g., principal components analysis), but the T-PHATE estimates are generally lower; this indicates that the data manifold is highly nonlinear (which is common among neural data (Altan et al., 2021; Jazayeri & Ostojic, 2021)) and captured more accurately in nonlinear dimensions with T-PHATE. For more information about this method, we refer readers to (Moon et al., 2019; Busch et al., 2023).

Results

We first evaluated the reliability of brain responses to these naturalistic stimuli using intersubject correlation (ISC) at each searchlight, which confirmed that both infants and adults show reliable activity task-relevant regions (mainly visual cortex; Fig.1A). Then, we computed the ID of brain responses and found that infants had low and static dimensionality across the brain relative to high and variable dimensionality in adults (Fig. 2B). Areas with higher ISC had relatively higher dimensionality in the adult brain. We replicated this apparent association in adults during naturalistic audio-only stimuli (Fig.1B).

By correlating each participant's ISC and ID estimates, we confirmed infants show no relationship between task engagement and ID ($t(2) = 1.5,\ p = 0.15$), whereas adults showed a significant relationship ($t(11) = 3.8,\ p = 0.003$). We investigated this further in another movie dataset comprising a wider age range (Partly Cloudy). This revealed that the association of ISC with ID strengthens across childhood ($F(5,146) = 5.53,\ p = 0.0001$). Lastly, we validated that this effect extends beyond the visual domain by replicating the association across adults listening to four audio-only stories (all $p < 0.001$).

After establishing that ID becomes task-selective over development, we contrasted ID in task-free versus naturalistic tasks. Across groups, ID during sleep/rest was higher overall than during movie viewing ($F(1,67) = 21.7,\ p < 0.001$). During sleep/rest, adults still showed greater dimensionality and more variability across the brain than infants ($t(30) = 3.0$, $p = 0.003$), as they did during movie viewing (Fig. 2A,B).

Finally, we asked how dimensionality shifts as a function of task engagement. To do this, we considered where ID changed between task-free and naturalistic task scans. In adults, the decrease in dimensionality during movie-watching relative to rest was specific to brain regions less engaged by tasks. That is, non-selective (low-ISC) regions showed compressed dimensionality during movie-watching relative to rest, while selective regions retained high dimensionality from rest to task. In infants, dimensionality decreased uniformly across the brain from sleep to task, suggesting that the infant brain does not yet dynamically allocate processing capacity to task-relevant regions (Fig. 2C).

In sum, we find that ID of task-driven activity increases over development, culminating in a reliable convergence between high-dimensional and task-selective activity in the mature brain. Without task state driving activity, ID is high-dimensional and unconstrained across the cortex regardless of age. Regardless of age, participants show compressed activity across the brain during tasks relative to task-free states. In adults, compression is specific to task-irrelevant regions, suggesting that tasks selectively dampen the dimensionality of other regions that are high-dimensional at rest. These findings indicate that the flexible expansion and compression of activity across the brain follows a protracted developmental trajectory. In the future, we plan to investigate how dimensionality collapse from rest to task emerges across childhood, to more deeply characterize how the dynamic allocation of dimensionality relates to cognitive processing.

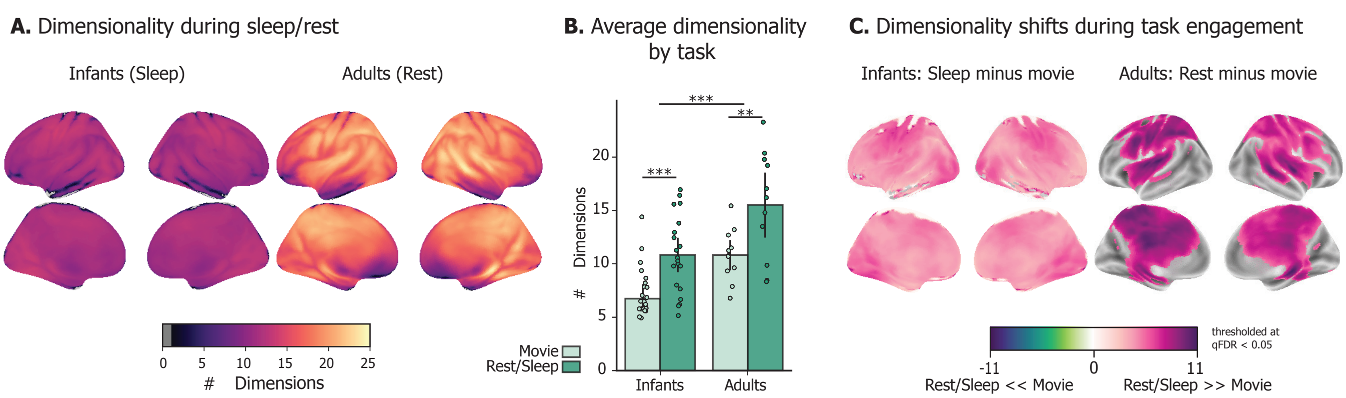

Figure 2: Differences in intrinsic dimensionality between task-free (sleep/rest) and task scans. (A) Maps of ID averaged across participants for sleeping infants and resting adults. (B) Average ID across the brain for movie-watching and sleep/rest. Dots are participants, bars are average across participants. (C) Difference between task-free (sleep/rest) and movie-watching ID estimates at each searchlight. Statistical significance was assessed with FSL's randomise with cluster-based correction. Visualization thresholded at $qFDR < 0.05$. $**p < 0.01$, $***p < 0.001$

References

Altan, E., Solla, S. A., Miller, L. E., & Perreault, E. J. (2021). Estimating the dimensionality of the manifold underlying multi-electrode neural recordings. *PLoS computational biology*, *17*(11), e1008591.

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*, *3*(3), 240–253.

Cunningham, J. P., & Yu, B. M. (2014, November). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509. Retrieved 2020-08-23, from https://www.nature.com/articles/nn.3776 (Number: 11 Publisher: Nature Publishing Group) doi: 10.1038/nn.3776

Jazayeri, M., & Ostojic, S. (2021). Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. *Current opinion in neurobiology*, *70*, 113–120.

Mazzucato, L., Fontanini, A., & La Camera, G. (2016). Stimuli reduce the dimensionality of cortical activity. *Frontiers in systems neuroscience*, *10*, 11.

Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., ... Krishnaswamy, S. (2019, December). Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, *37*(12), 1482–1492. Retrieved 2020-09-23, from https://www.nature.com/articles/s41587-019-0336-3 (Number: 12 Publisher: Nature Publishing Group) 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* (Vol. 14) (No. 6). Oxford University Press.

Nastase, S. A., Liu, Y.-F., Hillman, H., Zadbood, A., Hasenfratz, L., Keshavarzian, N., ... Hasson, U. (2021, September). The "Narratives" fMRI dataset for evaluating models of naturalistic language comprehension. *Scientific Data*, *8*(1), 250. Retrieved 2023-08-17, from https://www.nature.com/articles/s41597-021-01033-3 (Number: 1 Publisher: Nature Publishing Group) doi: 10.1038/s41597-021-01033-3

Richardson, H., Lisandrelli, G., Riobueno-Naylor, A., & Saxe, R. (2018). Development of the social brain from age three to twelve years. *Nature communications*, *9*(1), 1027.

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, *364*(6437), eaav7893.

Yates, T. S., Ellis, C. T., & Turk-Browne, N. B. (2023). Functional networks in the infant brain during sleep and wake states. *Cerebral Cortex*, *33*(21), 10820–10835.