Universality in mouse and human visual cortex: relating covariance to the spatial structure of latent dimensions

Raj Magesh Gauthaman (rgautha1@jh.edu)
Department of Cognitive Science, Johns Hopkins University
3400 N. Charles Street, Baltimore, MD 21218

Brice Ménard (menard@jhu.edu)
Department of Physics & Astronomy, Johns Hopkins University
3400 N. Charles Street, Baltimore, MD 21218

Michael F. Bonner (mfbonner@jhu.edu)
Department of Cognitive Science, Johns Hopkins University
3400 N. Charles Street, Baltimore, MD 21218

Abstract

Recent work has revealed the high-dimensional structure of visual cortex responses to natural images in both mice and humans, where stimulus-related variance is distributed as a power law over thousands of latent dimensions. Here, we characterize the covariance spectra of two datasets containing V1 responses to thousands of visual stimuli measured at two very different scales: mouse calcium imaging and human fMRI. We find that the power-law exponent α characterizing the spectral decay varies substantially across experiments, contradicting previous claims of universality and optimality in the power law exponents of visual cortex. However, we also discover a striking pattern where variance along a latent dimension is directly related to its spatial scale – a measure of how strongly neighboring neurons co-activate. When viewed through this lens, the spectra of the mouse and human neural activations show striking similarities, suggesting that both visual systems represent natural images in similar ways. Our results demonstrate that analyzing the spatial scale of latent modes of variation might be a more fundamental way to quantify the covariance structure of neural representations.

Keywords: visual cortex; dimensionality; natural scenes; human fMRI; mouse calcium imaging

Introduction

Visual cortex responses to natural images are often reported to be low-dimensional, where all the reliable variance lies along a small handful of dimensions corresponding to interpretable features. However, recent studies of mouse and human visual cortex have demonstrated that neural representations of visual stimuli have high-dimensional latent structure, with reliable stimulus-related variance spanning all latent dimensions and limited only by dataset size (Stringer, Pachitariu, Steinmetz, Carandini, & Harris, 2019; Gauthaman, Ménard, & Bonner, 2023). Specifically, the covariance spectrum of neural responses follows a power law where variance $\propto$ rank$^{\alpha}$.

The key metric characterizing this spectral decay is the power-law exponent (α) which controls the balance between expressivity and smoothness of the neural representation (Stringer et al., 2019). Early experimental and theoretical results argued that mouse visual cortex operates in a critical regime where α attains its optimal value of -1 (Stringer et al., 2019), and similar claims have been made about the power-law exponents in the internal representations of neural networks trained for visual tasks (Ghosh, Mondal, Agrawal, & Richards, 2022; Kong, Margalit, Gardner, & Norcia, 2022; Nassar, Sokol, Chung, Harris, & Park, 2020). However, more recent studies have reported a range of α values across different model organisms and computational models (Wang et al., 2023; Manley et al., 2024; Pospisil & Pillow, 2024; Gerum, Pirlot, Fyshe, & Zylberberg, 2022).

Our contribution Here, we analyze previously published mouse calcium imaging data (Stringer et al., 2019) as well as a large-scale human fMRI dataset (Allen et al., 2021) containing primary visual cortex responses to thousands of natural images. We find that α varies significantly between the mouse and human datasets. Since neighboring neurons tend to co-activate, we reasoned that the latent modes of neural variation are likely to exhibit topographic organization, and, thus, we analyzed their spatial structure. Strikingly, when the stimulus-related variance along each latent dimension is plotted against its *spatial scale* instead of its rank, we discover similar power-law exponents for both the mouse and human datasets, suggesting that both neural systems encode visual information in the same way across an immense range of spatial scales ranging from single neurons to voxels.

Methods

Large-scale datasets We analyze fMRI responses in the primary visual cortex (V1) of 8 participants from the Natural Scenes dataset who performed a continuous recognition task on 10,000 natural scene stimuli (Allen et al., 2021). Additionally, we analyze a mouse calcium imaging dataset also containing V1 responses to 2,800 natural images from 7 mice (Stringer et al., 2019).

Computing covariance spectra Cross-decomposition is a generalization of principal component analysis (PCA) that estimates cross-validated covariance spectra, measuring only stimulus-related variance that generalizes across different image presentations[1] (Fig. 1A). Given two sets of neural responses $X, Y \in \mathbb{R}^{n \times p}$ to the same n stimuli from p neurons (or voxels), we first identify shared $d = \min(n, p)$ latent dimensions by computing the singular value decomposition of their cross-covariance on *training* data: $\text{cov}(X_{\text{train}}, Y_{\text{train}}) = U\Sigma V^{\top}$, where $U, V \in \mathbb{R}^{p \times d}$ are orthonormal. Then, we project *test* data onto the shared latent space and compute their covariance $\hat{\Sigma} = \text{cov}(X_{\text{test}}U, Y_{\text{test}}V)$. Finally, we extract the diagonal of $\hat{\Sigma} \in \mathbb{R}^{d \times d}$, which represents the spectrum of reliable variance in the dataset that generalizes to novel stimuli across trials[2].

Estimating spatial scale Each latent dimension in the dataset is represented by a linear combination of neurons (columns of U and V) whose weights can be visualized on the cortical surface (Fig. 1C[3]). These modes of variation could have coarse- or fine-grained spatial structure depending on how strongly neighboring neurons co-activate. We define the spatial scale of the k-th latent dimension[4] (e.g.

[1] In both datasets, we focus on stimuli seen at least twice and restrict our analyses to the first two presentations of each image.

[2] In practice, we compute better estimates of the covariance spectra using 8-fold cross-validation across stimuli and normalize the spectra by the number of neurons p.

[3] For visualization purposes in Fig. 1C, we plot the latent dimensions of a large visual region in the human brain but all other results are for V1.

[4] We average the spatial scales of U_k and V_k to obtain a single estimate for each latent dimension.

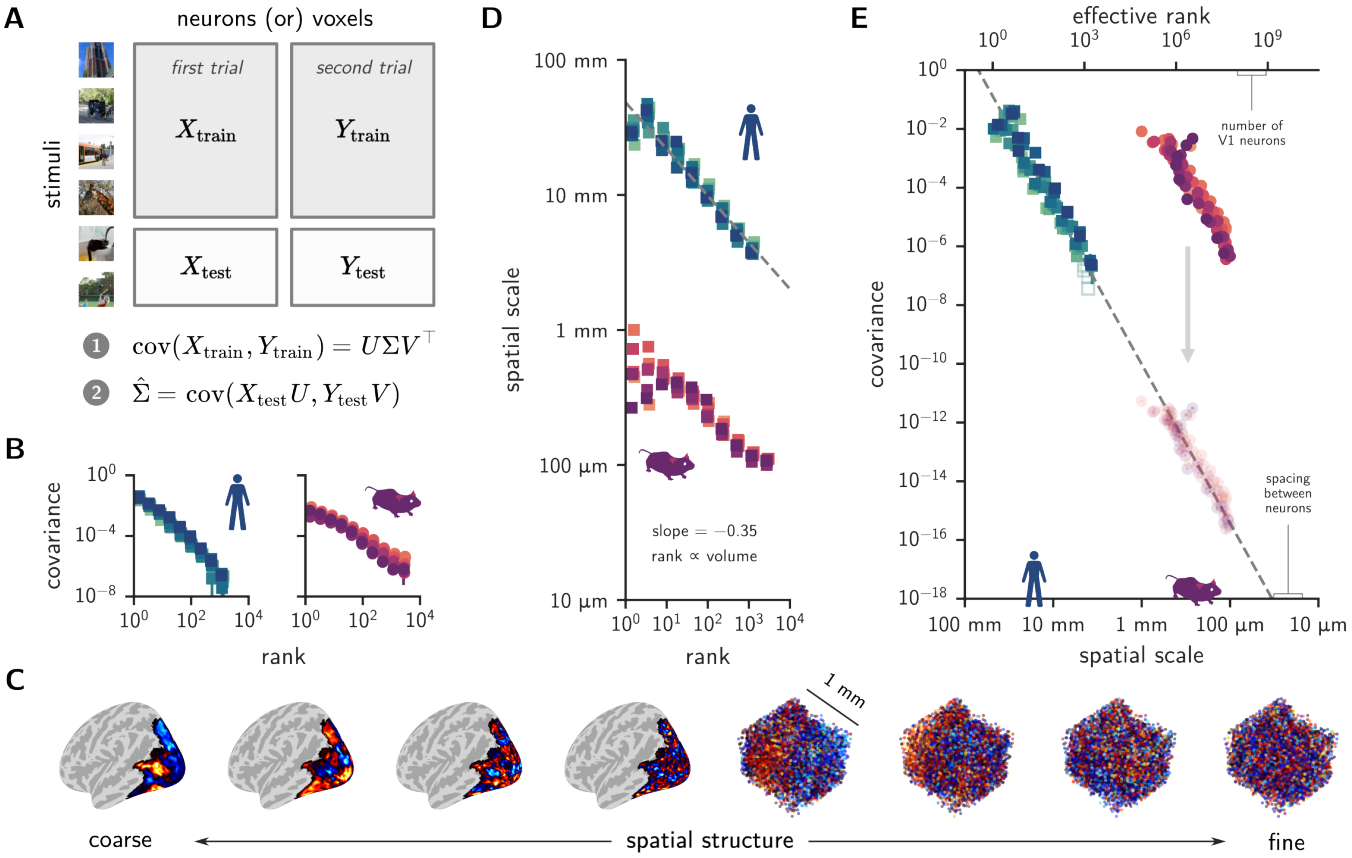

Figure 1: **(A)** Schematic of cross-decomposition procedure. Neural responses to the same stimuli on two different trials are used to identify shared latent dimensions and estimate reliable covariance spectra $\hat{\Sigma}$. **(B)** The covariance spectra for human (blue-green) and mouse (purple-orange) V1 responses to natural images have characteristic power-law structure that is consistent within individuals in each dataset but different across datasets. **(C)** Sample latent dimensions in humans (projected onto the cortical surface) and mice (in a 1 mm cube of cortex) display spatial structure at various scales. **(D)** The spatial scale of the latent dimensions decays with rank as a power law with exponent -0.35, suggesting volumetric scaling. **(E)** The reliable variance along a latent dimension scales with its spatial scale as a power law whose exponent is consistent both *within* and *across* the datasets. All spectra are averaged within bins of exponentially increasing width to extract reliable signal even in noisy regimes (Lin & Newberry, 2023). Error bars denote standard deviations across 8 folds of cross-validation.

$U_k \in \mathbb{R}^p$) as the width σ of the Gaussian kernel G_σ required to spatially smooth its weights for their variance to decay to $||G_\sigma(U_k)||^2 = 10^{-3}$.

Results & Discussion

Both human fMRI and mouse calcium imaging responses to natural images in primary visual cortex display characteristic power-law covariance spectra (variance $\propto$ rank$^\alpha$) over multiple orders of magnitude (Fig. 1B). However, while the power-law exponents α are consistent across individuals within each dataset, they differ significantly across datasets ($\alpha_{\text{mouse}} = -1.19 \pm 0.11$; $\alpha_{\text{human}} = -1.85 \pm 0.11$, mean $\pm$ sd).

Additionally, we observe that the spatial scale of these latent dimensions also decays with rank as a power law (Fig. 1D). In the human data, spatial scale $\propto$ rank$^{-0.35}$, which implies that the number of latent dimensions approximately scales with volume (i.e., rank $\propto$ spatial scale^3 = volume). This suggests that visual information is densely distributed across all voxels.

Surprisingly, plotting the stimulus-related variance along each latent dimension against its spatial scale (Fig. 1E) reveals universal power-law scaling across both humans and mice (variance $\propto$ spatial scale$^\beta$), with similar indices within and across datasets ($\beta_{\text{human}} = -5.36 \pm 0.32$; $\beta_{\text{mouse}} = -5.27 \pm 1.14$). We also note that when the spatial scale approaches the typical separation between neurons ($\approx 20\,\mu\text{m}$), the effective rank of the spectrum approximates the number of neurons in human V1, again suggesting that the neural representation of natural images is *dense* in the available neurons.

Together, our results suggest that analyzing the spatial structure of latent dimensions might reveal fundamental principles of cortical organization that are shared across species and detectable at scales ranging from single neurons to voxels comprising millions of neurons.

References

Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., ... Kay, K. (2021, December). A massive 7t fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, *25*(1), 116–126. Retrieved from https://doi.org/10.1038/s41593-021-00962-x doi: 10.1038/s41593-021-00962-x

Gauthaman, R. M., Ménard, B., & Bonner, M. (2023). Revealing the high-dimensional latent structure in visual cortical representations. In *2023 conference on cognitive computational neuroscience.* Cognitive Computational Neuroscience. Retrieved from http://dx.doi.org/10.32470/CCN.2023.1652-0 doi: 10.32470/ccn.2023.1652-0

Gerum, R. C., Pirlot, C., Fyshe, A., & Zylberberg, J. (2022). *Different spectral representations in optimized artificial neural networks and brains.* arXiv. Retrieved from https://arxiv.org/abs/2208.10576 doi: 10.48550/ARXIV.2208.10576

Ghosh, A., Mondal, A. K., Agrawal, K. K., & Richards, B. (2022). *Investigating power laws in deep representation learning.* arXiv. Retrieved from https://arxiv.org/abs/2202.05808 doi: 10.48550/ARXIV.2202.05808

Kong, N. C. L., Margalit, E., Gardner, J. L., & Norcia, A. M. (2022, January). Increasing neural network robustness improves match to macaque v1 eigenspectrum, spatial frequency preference and predictivity. *PLOS Computational Biology*, *18*(1), e1009739. Retrieved from https://doi.org/10.1371/journal.pcbi.1009739 doi: 10.1371/journal.pcbi.1009739

Lin, Q., & Newberry, M. (2023, August). Seeing through noise in power laws. *Journal of The Royal Society Interface*, *20*(205). Retrieved from http://dx.doi.org/10.1098/rsif.2023.0310 doi: 10.1098/rsif.2023.0310

Manley, J., Lu, S., Barber, K., Demas, J., Kim, H., Meyer, D., ... Vaziri, A. (2024, March). Simultaneous, cortex-wide dynamics of up to 1 million neurons reveal unbounded scaling of dimensionality with neuron number. *Neuron*. Retrieved from http://dx.doi.org/10.1016/j.neuron.2024.02.011 doi: 10.1016/j.neuron.2024.02.011

Nassar, J., Sokol, P., Chung, S., Harris, K. D., & Park, I. M. (2020). On 1/n neural representation and robustness. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 6211–6222). Curran Associates, Inc.

Pospisil, D. A., & Pillow, J. W. (2024, February). Revisiting the high-dimensional geometry of population responses in visual cortex. Retrieved from http://dx.doi.org/10.1101/2024.02.16.580726 doi: 10.1101/2024.02.16.580726

Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M., & Harris, K. D. (2019, June). High-dimensional geometry of population responses in visual cortex. *Nature*, *571*(7765), 361–365. Retrieved from https://doi.org/10.1038/s41586-019-1346-5 doi: 10.1038/s41586-019-1346-5

Wang, Z., Mai, W., Chai, Y., Qi, K., Ren, H., Shen, C., ... Wen, Q. (2023, February). The scale-invariant covariance spectrum of brain-wide activity. Retrieved from http://dx.doi.org/10.1101/2023.02.23.529673 doi: 10.1101/2023.02.23.529673