PCA, Dimensionality, and NSD

Thomas Gebhart, Asha Ramalaxmi, Kendrick Kay, and NSD Core Contributors

Outline

Outline

- SVD / PCA Intro
- Low-dimensional principal component structure in NSD
- High-dimensional principal component structure in NSD (Stringer et al.)
- Future Work

Stringer, Carsen, et al. "High-dimensional geometry of population responses in visual cortex." *Nature* 571.7765 (2019): 361-365

Singular Value Decomposition (SVD)

- Given data matrix $X \in \mathbb{R}^{n \times p}$, we can factorize $X = U \Sigma V^\intercal$.
- U, V are orthogonal.
- $\Sigma \in \mathbb{R}^{n \times p}$ is a diagonal matrix of singular values $\sigma_i = \Sigma_{i,i}$.
- Factorizes a linear transformation into a rotation/reflection, a scaling, and another rotation/reflection.

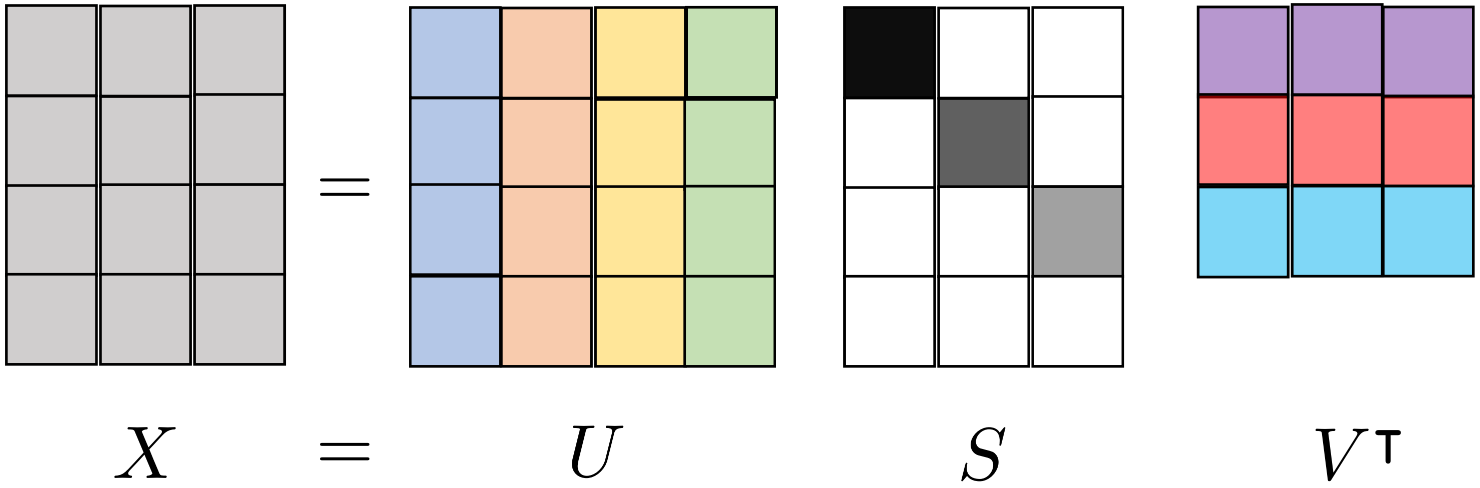

Lucas Vieira https://commons.wikimedia.org/w/index.php?curid=11416486

Principal Component Analysis (PCA)

- Let X be mean-centered column-wise (or z-scored).
- $X^\intercal X$ is then the empirical covariance matrix.
- We can factorize the covariance $X^\intercal X = V\Sigma^\intercal U^\intercal U\Sigma V^\intercal = V\Sigma^2 V^\intercal$
- V and Σ can be used to find directions of maximal variance.
- Can also be used to project into lower-dimensional space.
- Fitting an ellipsoid, where each axis represents a principal component.

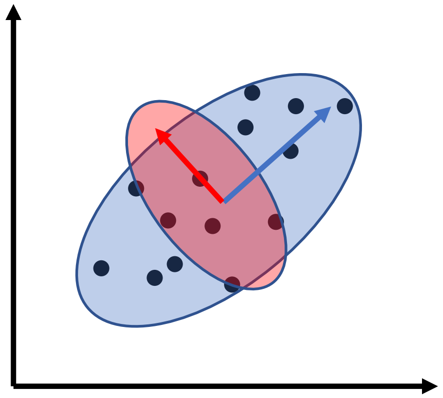

(looking at) Low-Dimensional Structure in NSD

Focus on VTC

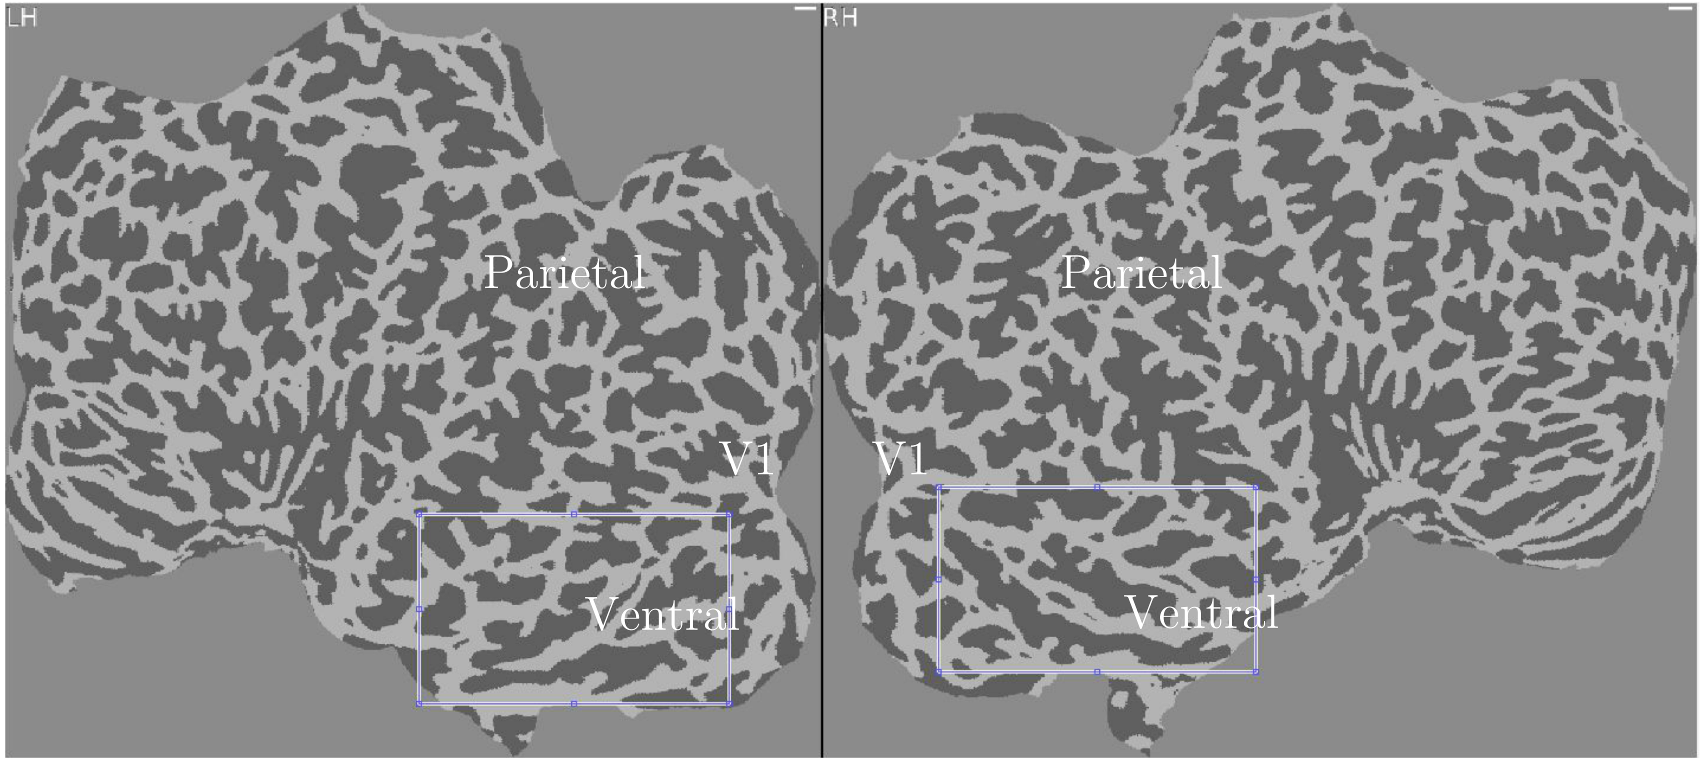

Choose Vertices

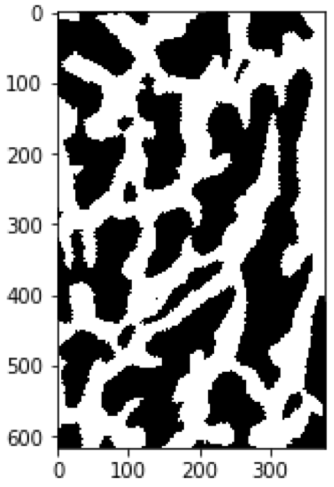

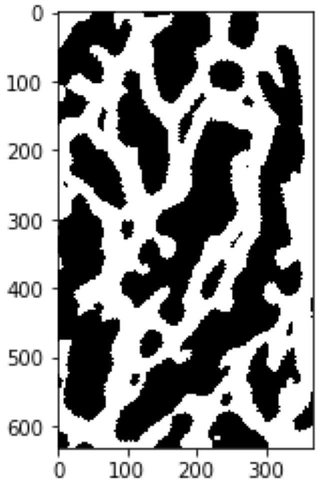

Index
Voxels

Map voxel BOLD responses to vertices via nearest-neighbor, averaging across cortical depth.

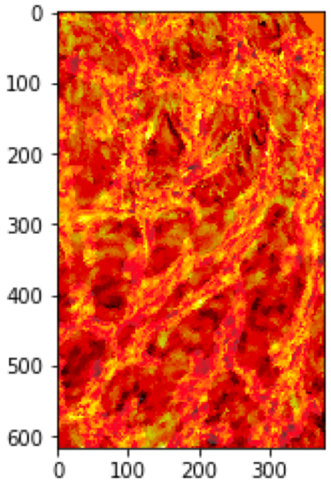

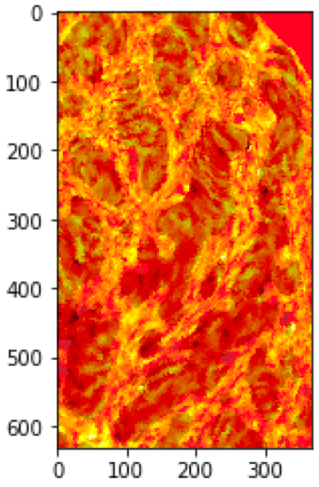

Dimensionality Reduction: Vertex Loadings on Image PCs

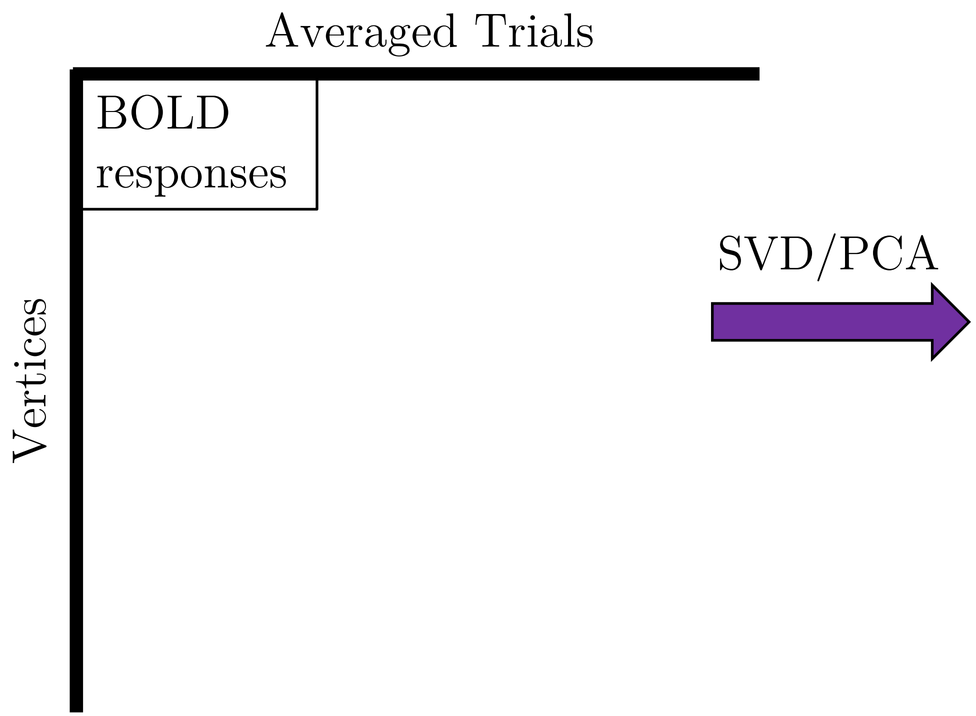

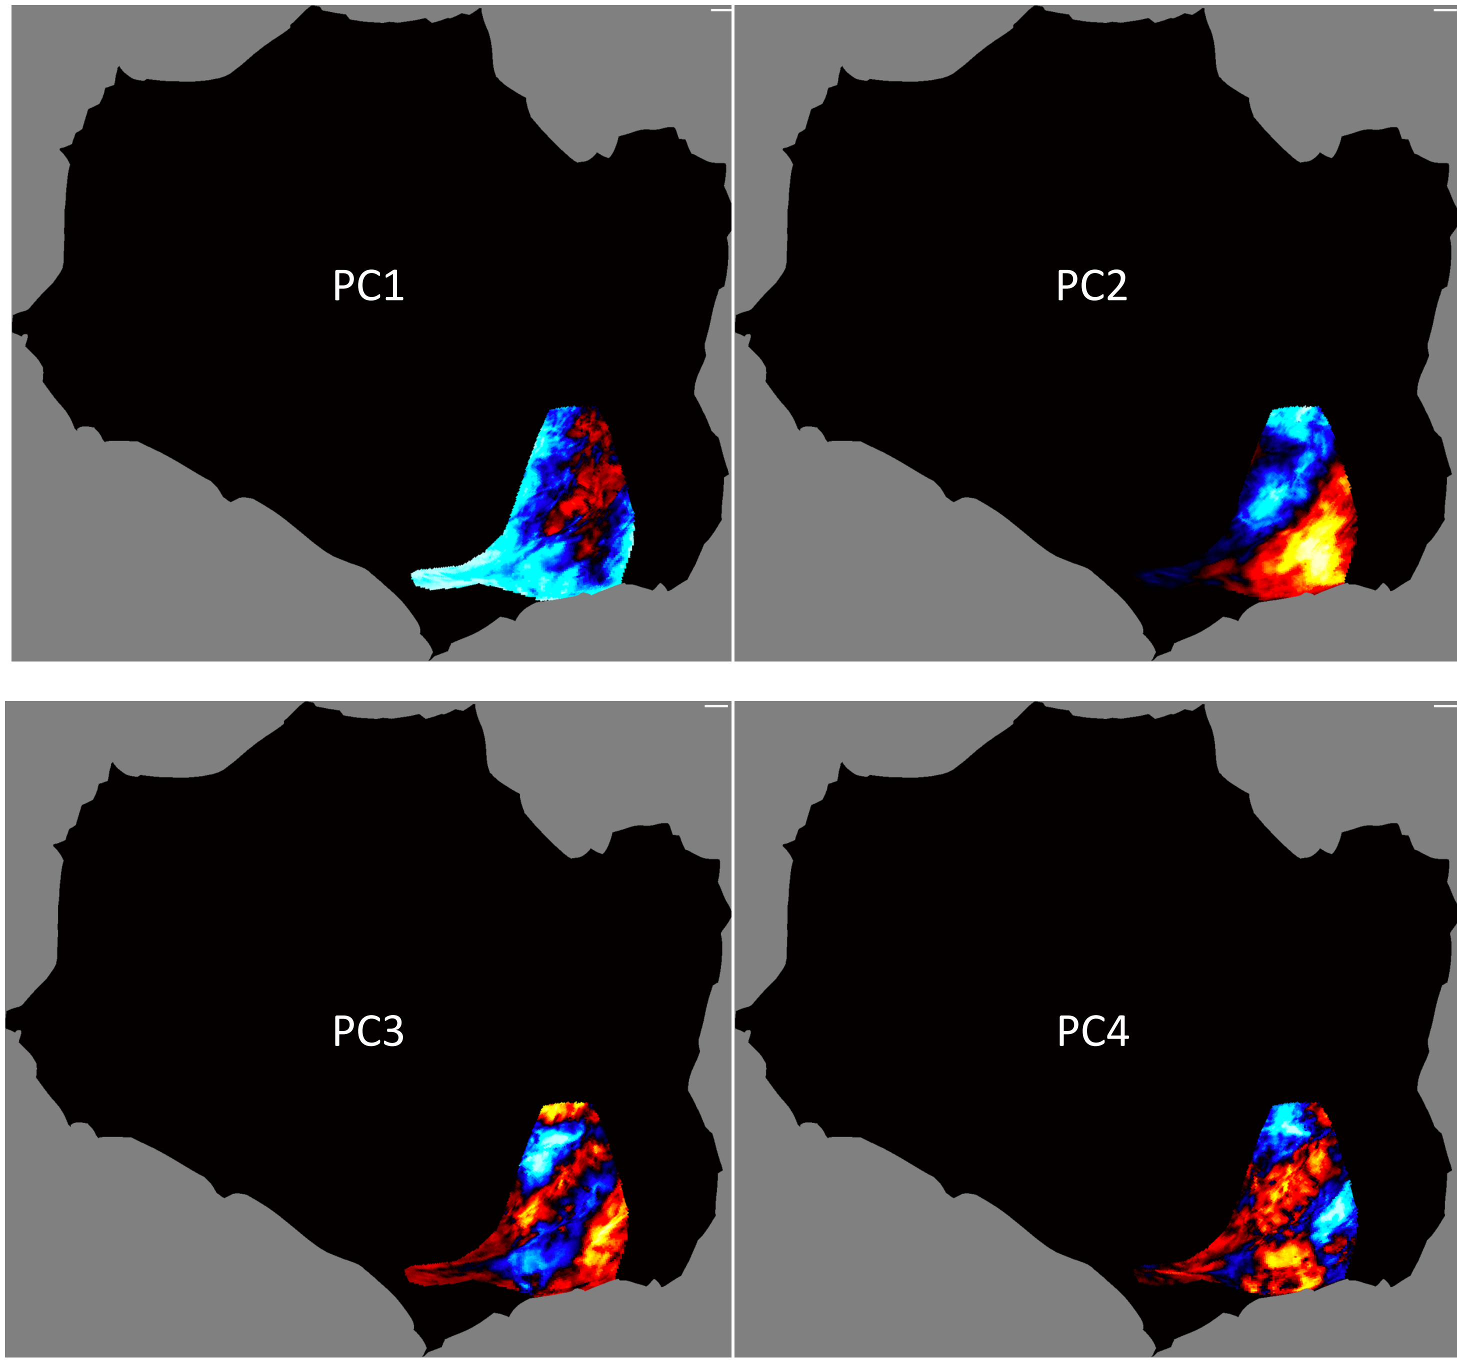

Dimensionality Reduction: Vertices in Image Space

Certainly non-Gaussian (at least in 1-3 PCs)

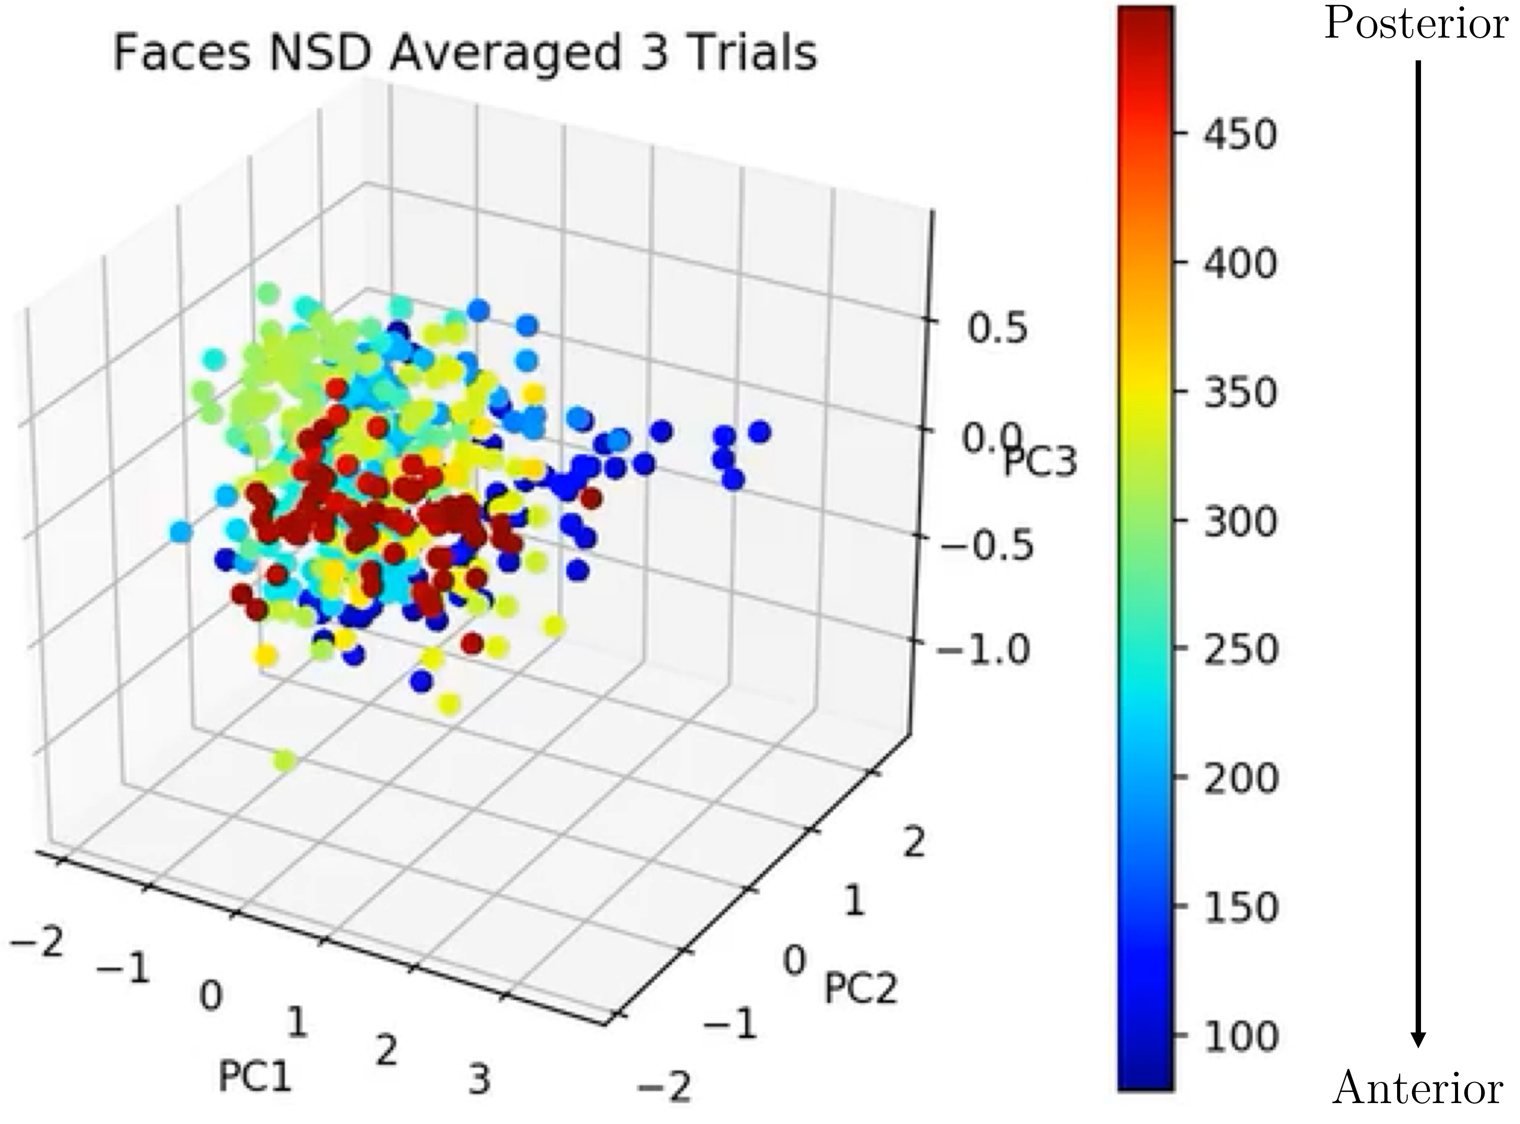

Correlate Responses with COCO Categories

COCO Supercategory

Trial	appliance	sports	person	furniture	electronic	vehicle	animal	outdoor	accessory	indoor	kitchen	food
0	0.0	0.0	0.0	0.0	0.0	0.0	14.0	0.0	0.0	0.0	0.0	0.0
1	0.0	13.0	9.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	0.0	1.0	6.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	1.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

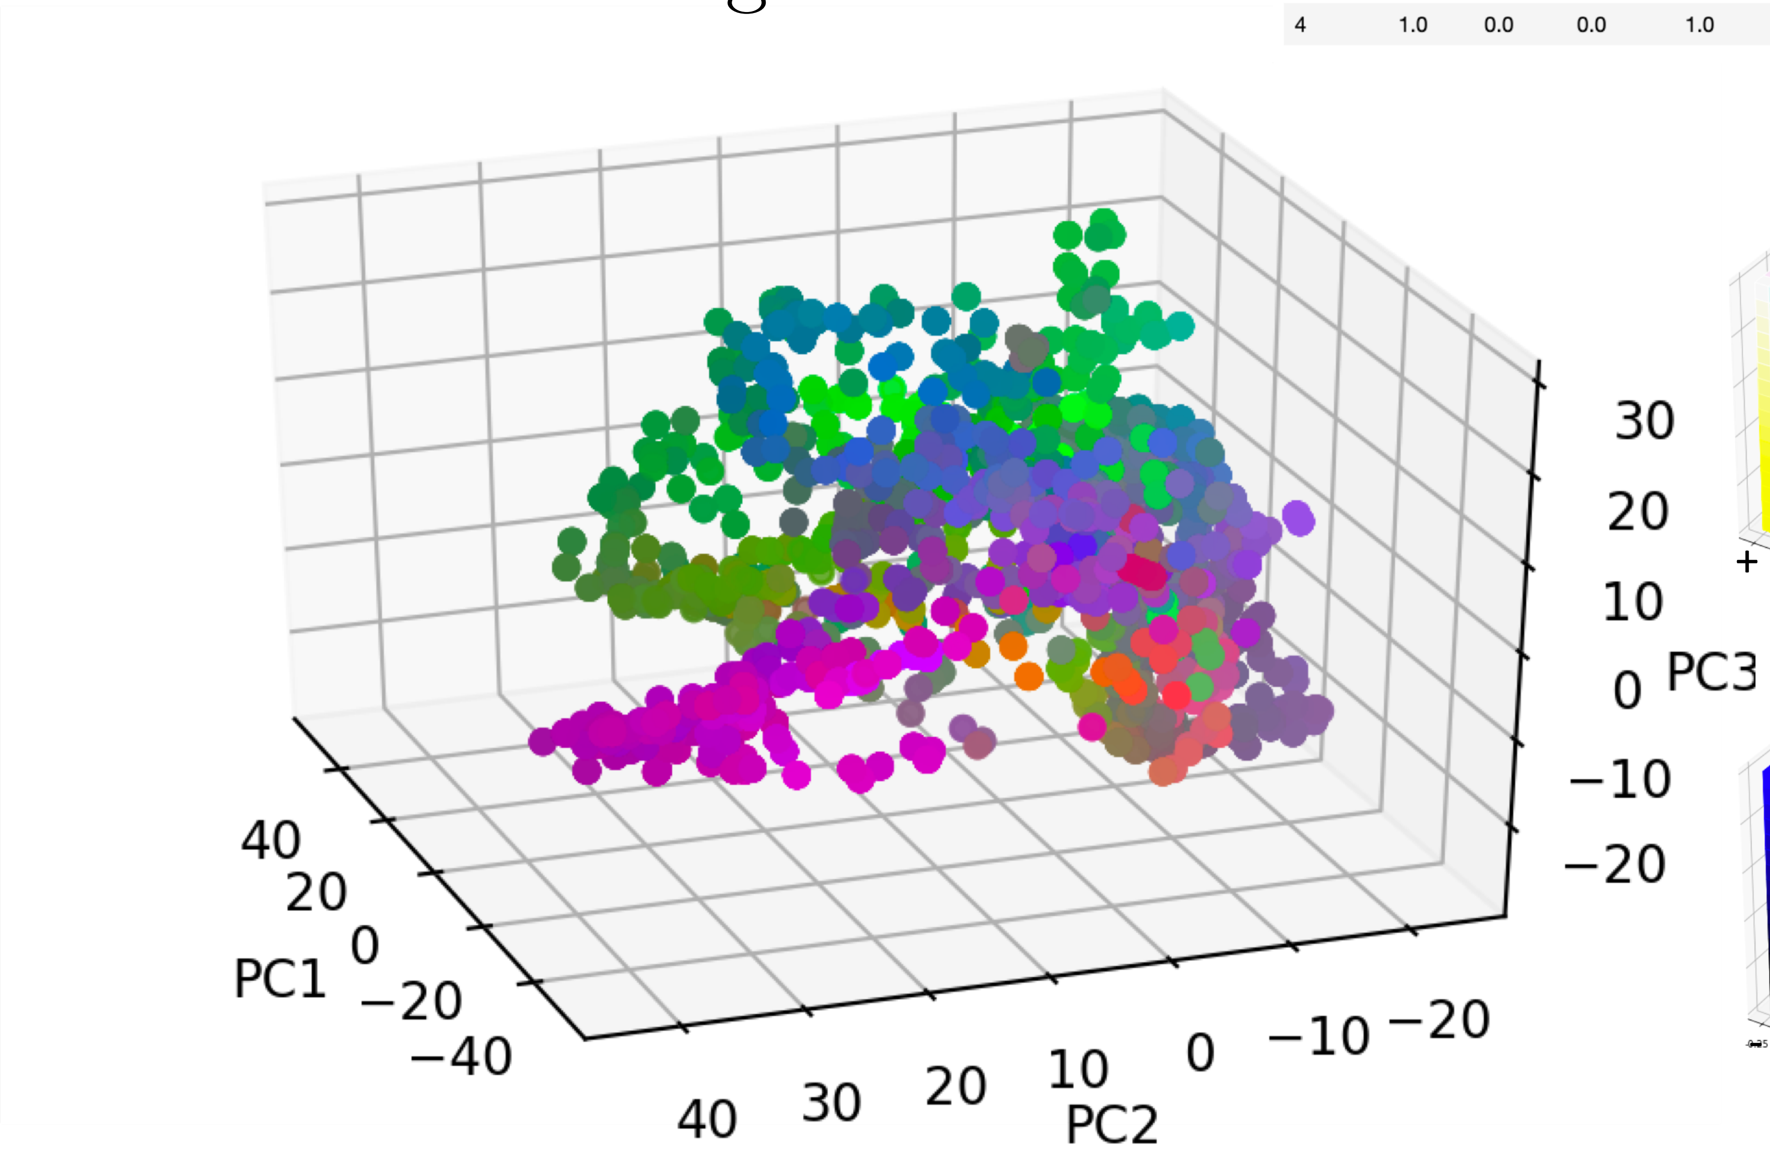

Correlate with Vertex x Image selectivity

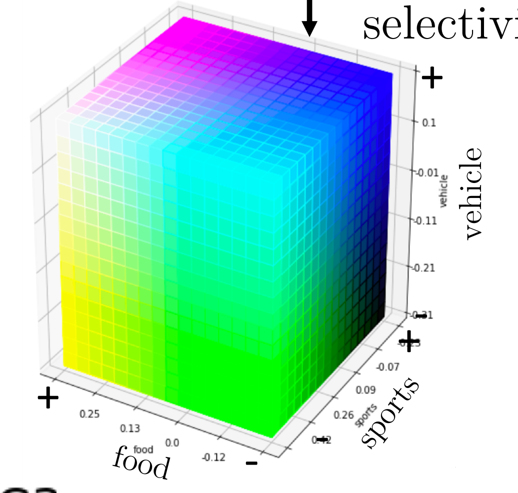

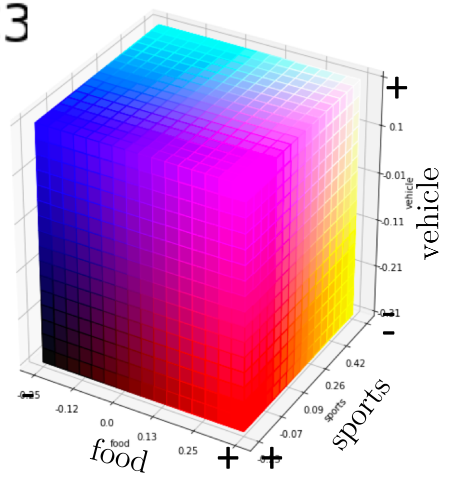

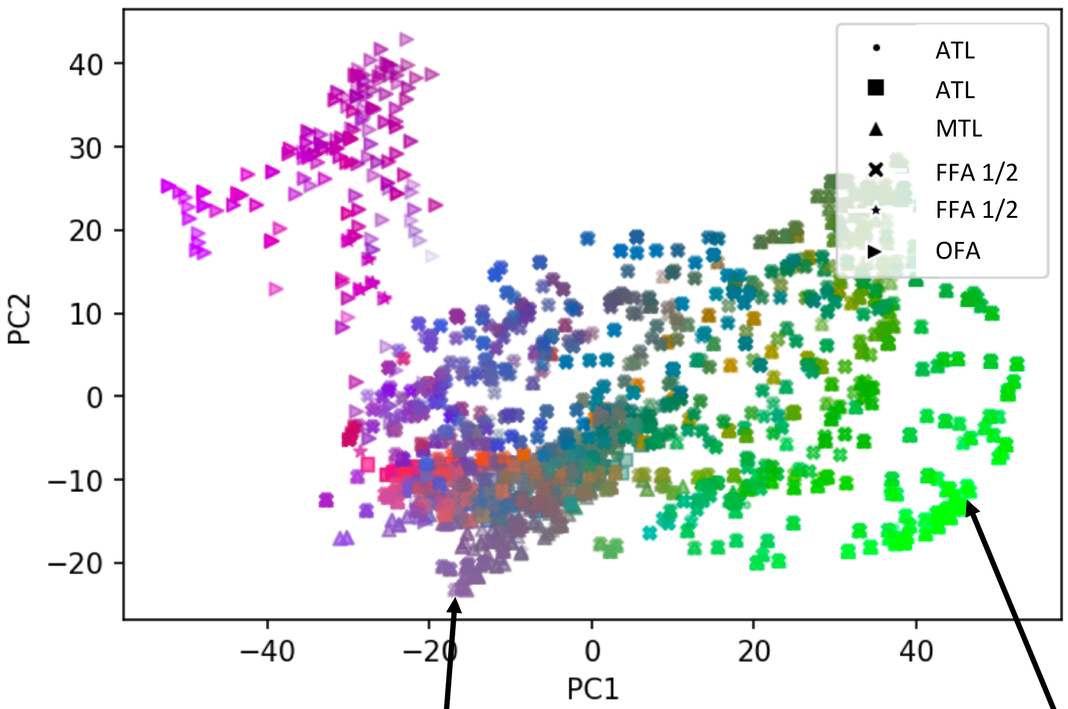

Top Stimuli Loadings on Vertices

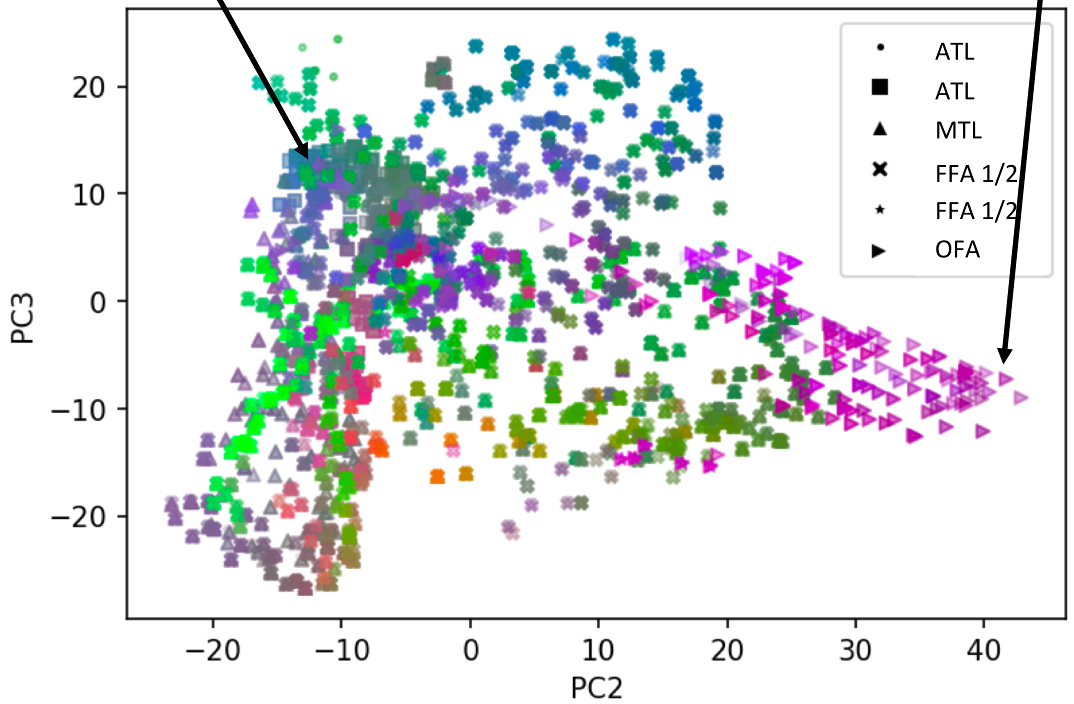

(looking at) High-Dimensional Structure in NSD

A Higher-level View

ARTICLE

https://doi.org/10.1038/s41586-019-1346-5

High-dimensional geometry of population responses in visual cortex

Carsen Stringer[1,2,6]*, Marius Pachitariu[1,3,6]*, Nicholas Steinmetz[3,5], Matteo Carandini[4,7] & Kenneth D. Harris[3,7]*

A neuronal population encodes information most efficiently when its stimulus responses are high-dimensional and uncorrelated, and most robustly when they are lower-dimensional and correlated. Here we analysed the dimensionality of the encoding of natural images by large populations of neurons in the visual cortex of awake mice. The evoked population activity was high-dimensional, and correlations obeyed an unexpected power law: the *n*th principal component variance scaled as $1/n$. This scaling was not inherited from the power law spectrum of natural images, because it persisted after stimulus whitening. We proved mathematically that if the variance spectrum was to decay more slowly then the population code could not be smooth, allowing small changes in input to dominate population activity. The theory also predicts larger power-law exponents for lower-dimensional stimulus ensembles, which we validated experimentally. These results suggest that coding smoothness may represent a fundamental constraint that determines correlations in neural population codes.

The visual cortex contains millions of neurons, and the patterns of activity that images evoke in these neurons form a 'population code'. The structure of this code is largely unknown, due to the lack of techniques that are able to record from large populations. Nonetheless, the population code is the subject of long-standing theories.

Simultaneous recordings of over 10,000 neurons

To obtain simultaneous recordings of approximately 10,000 cells from mouse V1, we used resonance-scanning two-photon calcium microscopy, using 11 imaging planes spaced at 35 μm (Fig. 1a). The slow time course of the GCaMP6s sensor enabled activity to be detected at

Stringer, Carsen, et al. "High-dimensional geometry of population responses in visual cortex." *Nature* 571.7765 (2019): 361-365

A Higher-level View

- Basic idea: noise upwardly biases estimation of variance, so compute covariance for PCA from different trials.
- Compute variance explained by each PC dimension.
- Fit line to eigenspectrum decay in log space.

$$\frac{\Sigma}{\mathbf{1}^\intercal \Sigma \mathbf{1}}$$

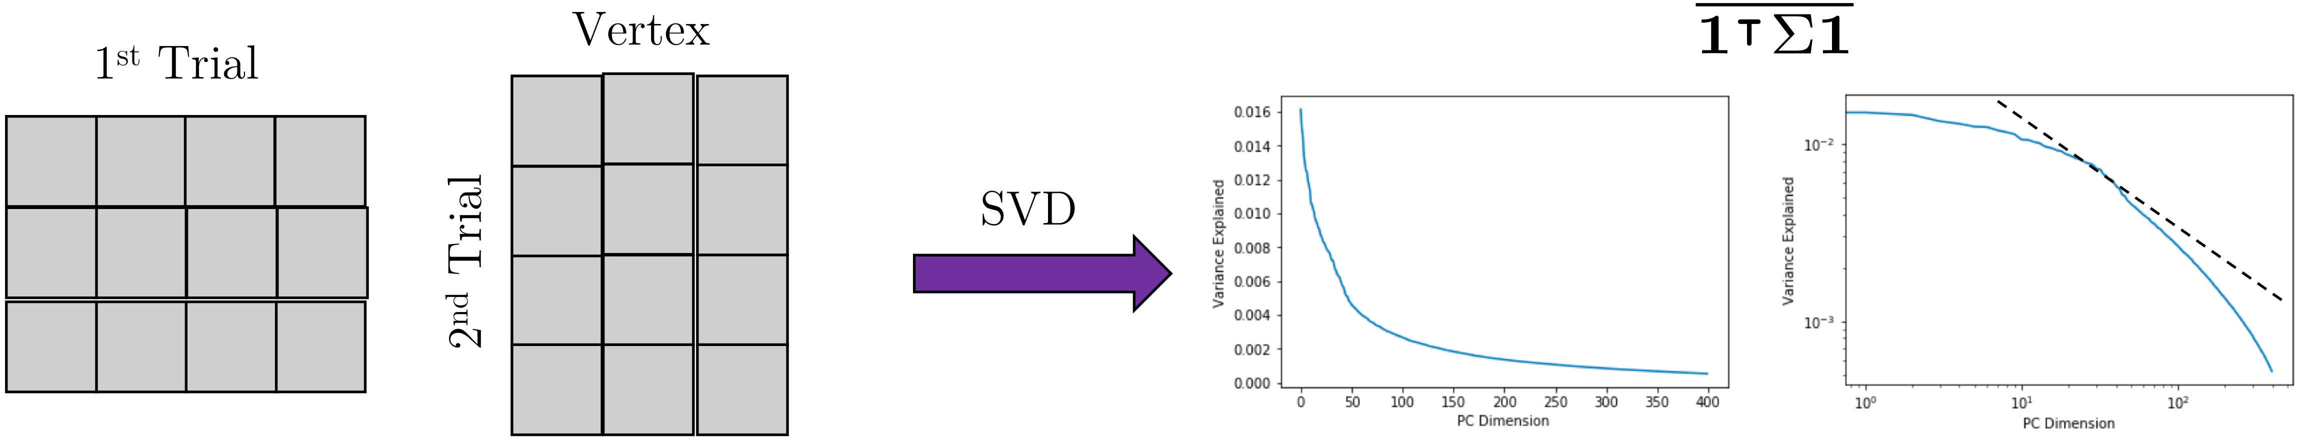

$$X_1^\intercal \quad \cdot \quad X_2 \qquad V\Sigma^2 V^\intercal$$

A Higher-level View

- Higher slope implies lower-dimensional, correlated coding.

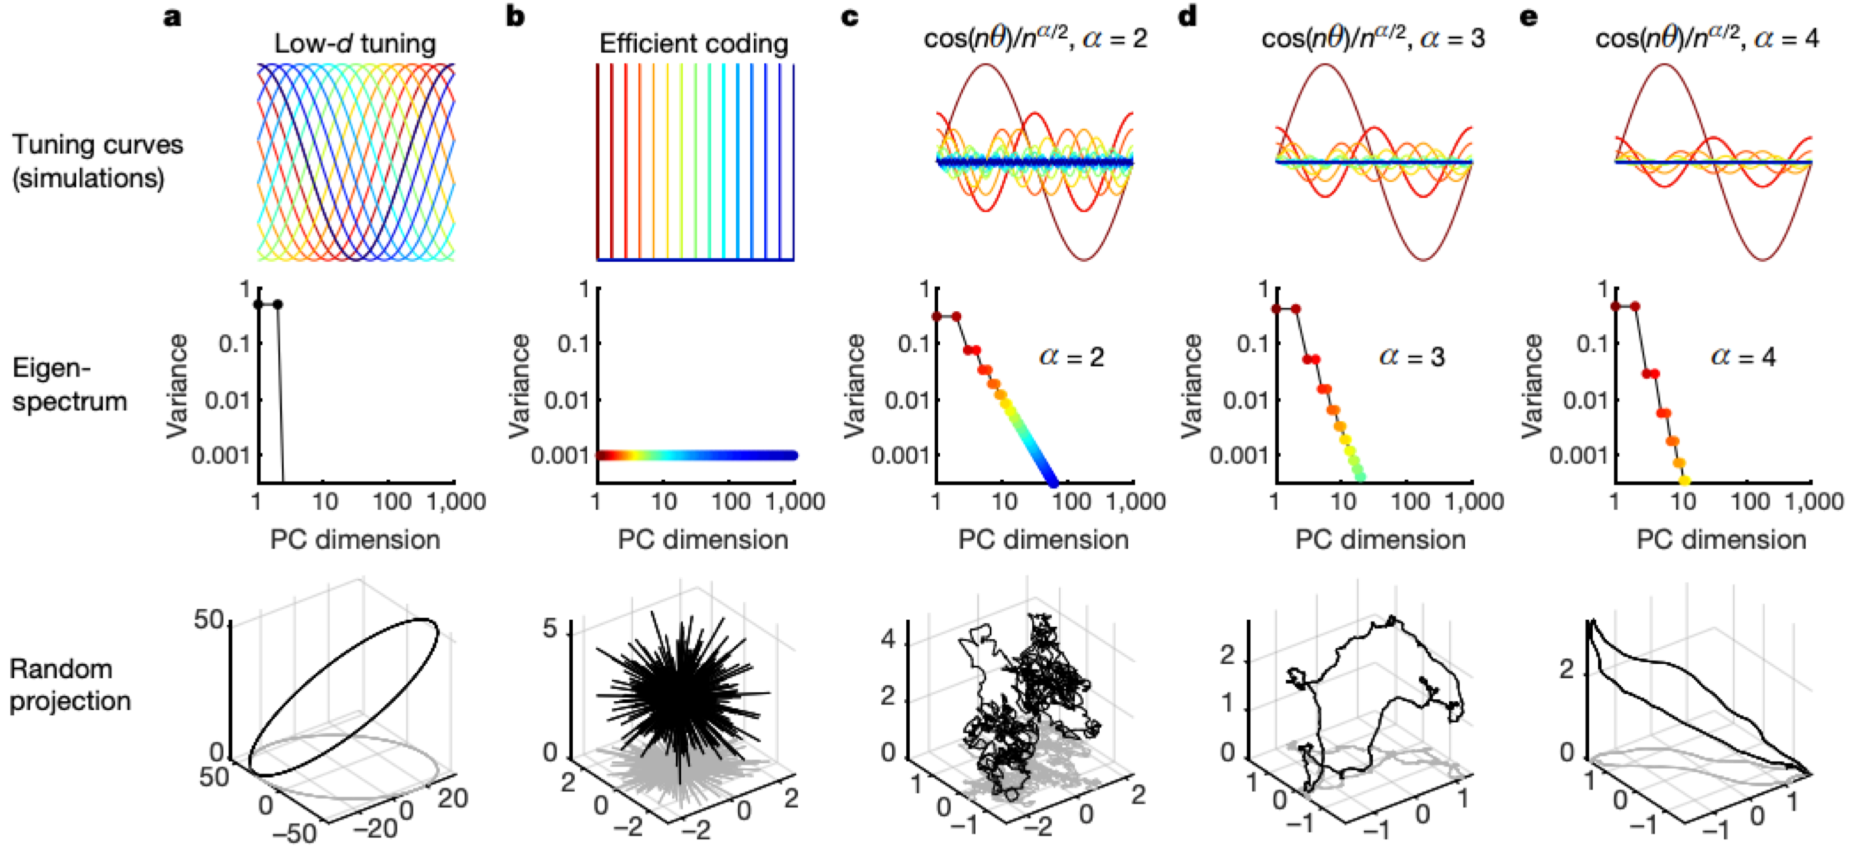

Stringer, Carsen, et al. "High-dimensional geometry of population responses in visual cortex." *Nature* 571.7765 (2019): 361-365

A Higher-Level View

- Can we do the same analysis with NSD?
- Many differences between datasets:
  - Mouse calcium microscopy vs. human fMRI.
  - Individual neurons vs. voxels.
  - Stimulus dimensionality not controlled (some NSD synthetic).
  - Source and characteristics of noise differ.
  - V1 vs. whole brain.
- But these differences could lead to insightful generalizations.

NSD Eigenspectrum

- Sample “searchlights” of vertices, map to native surface.
- Compute eigenspectrum decay of each searchlight.

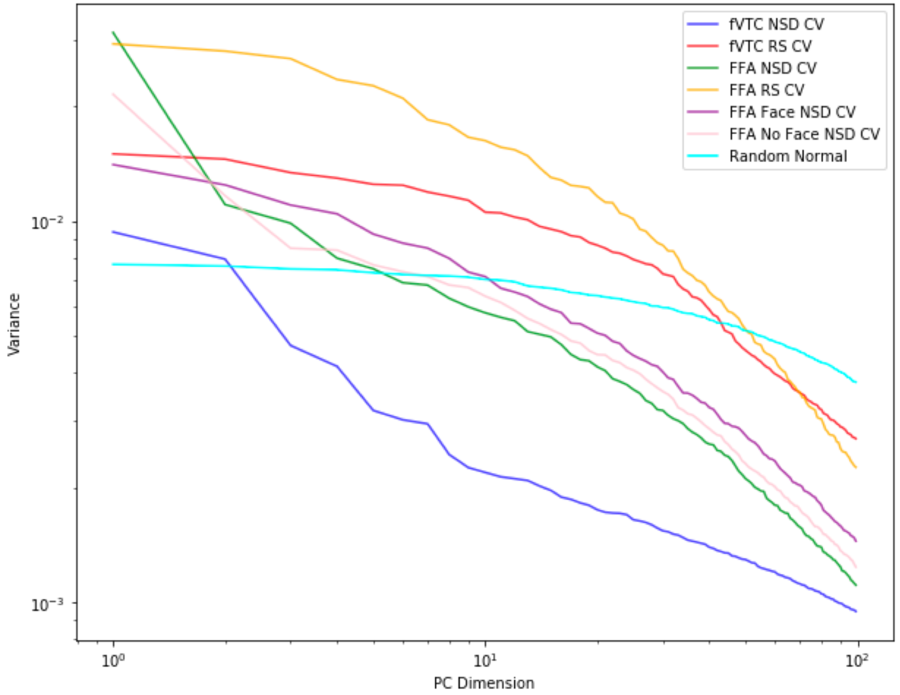

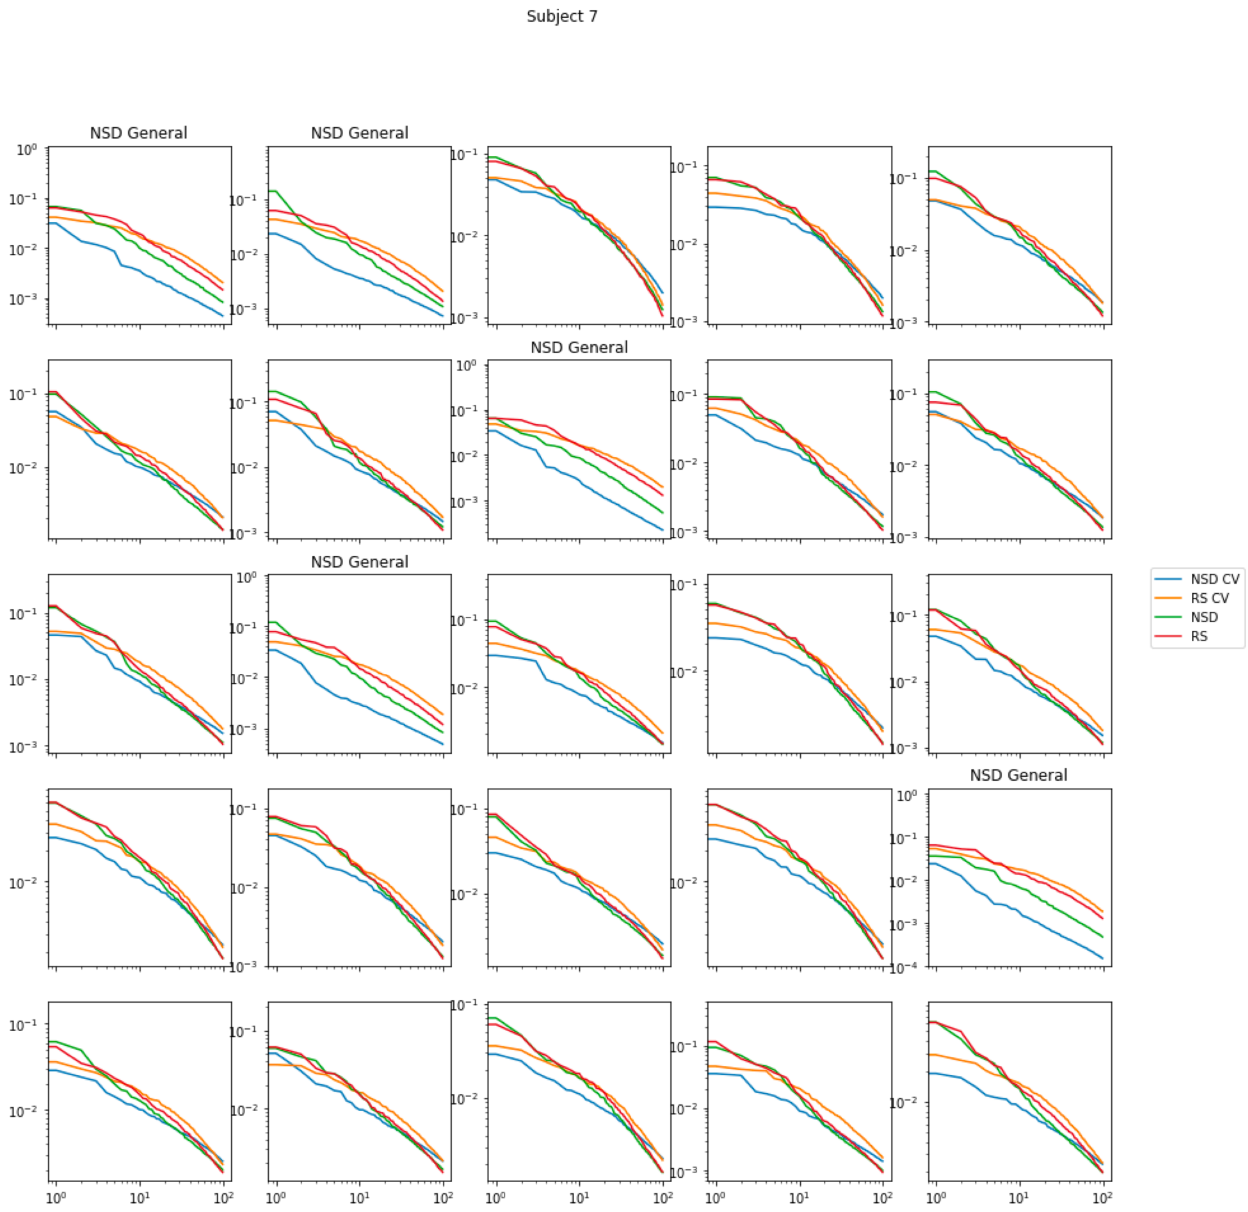

NSD Eigenspectrum

- Sample "searchlights" of vertices, map to native surface.
- Compute eigenspectrum decay of each searchlight.

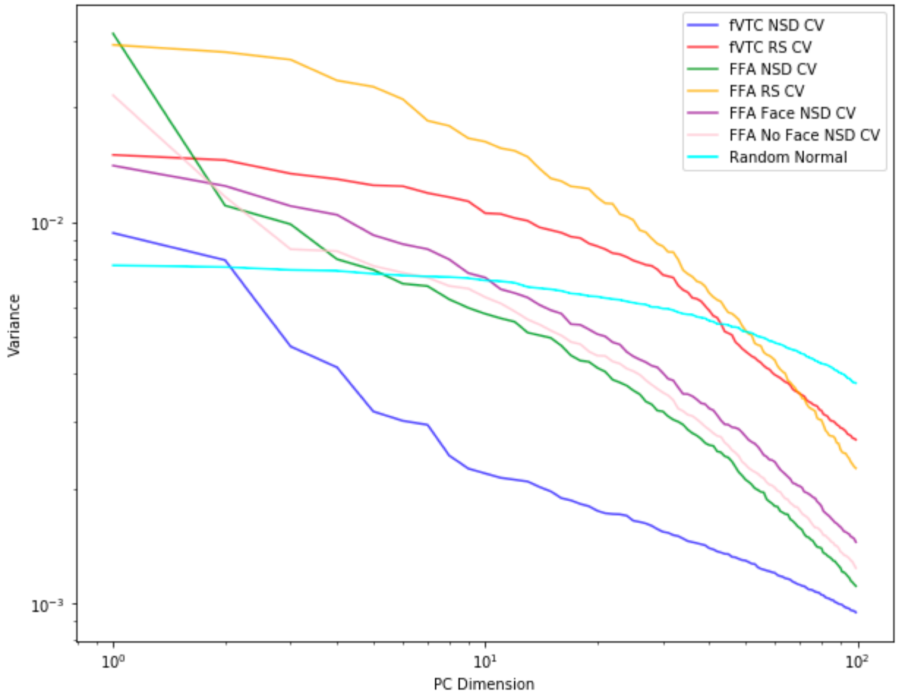

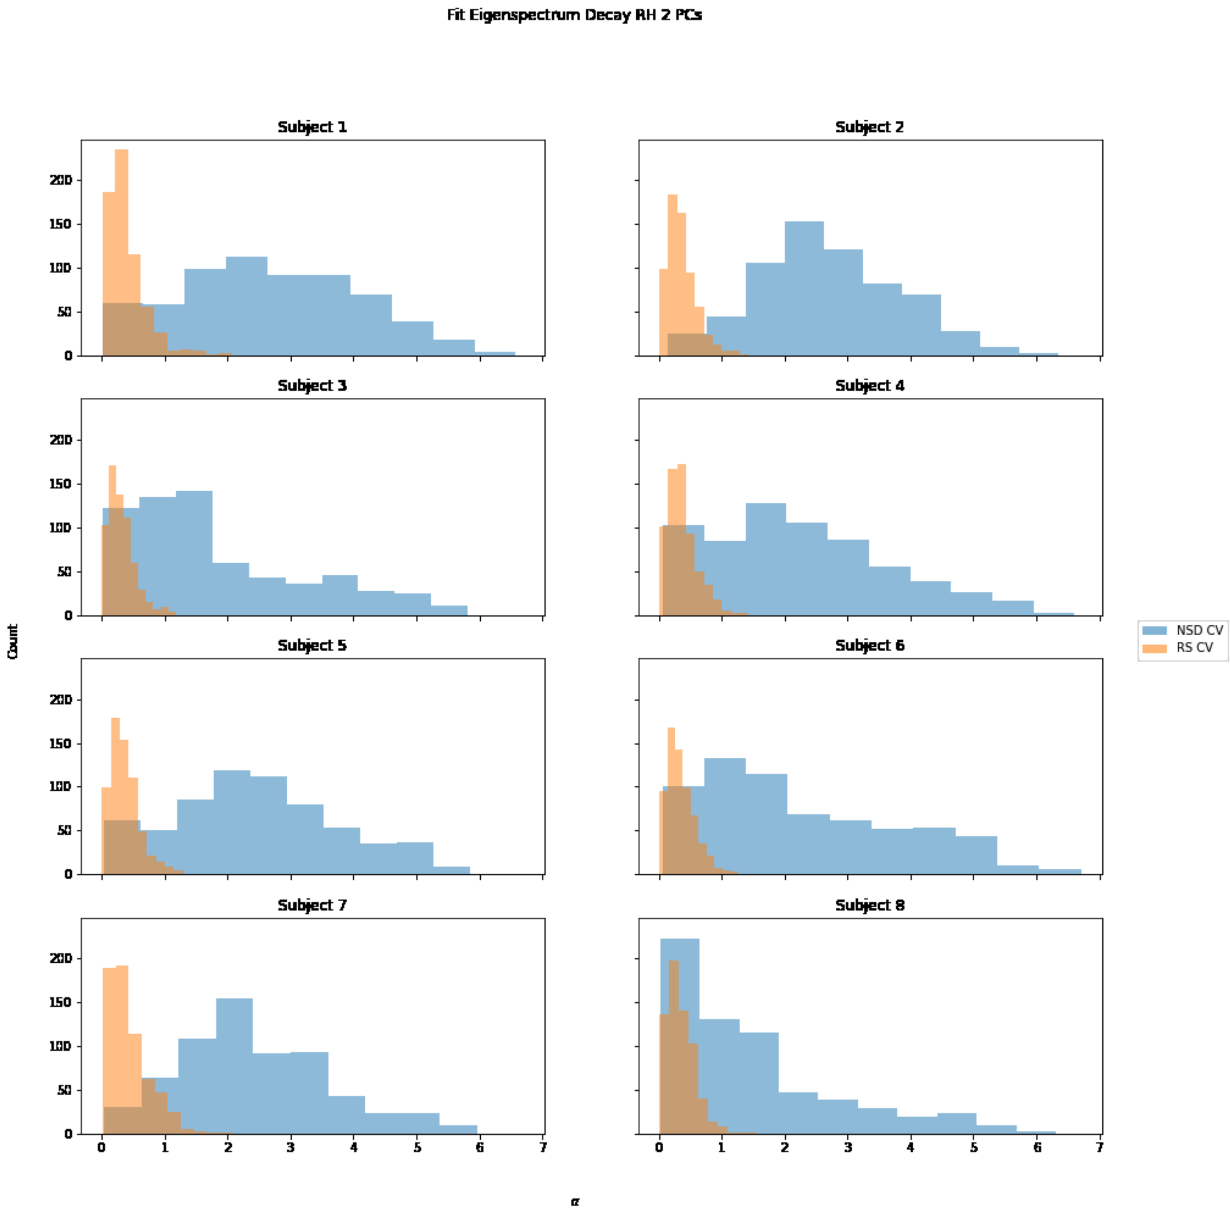

NSD Eigenspectrum

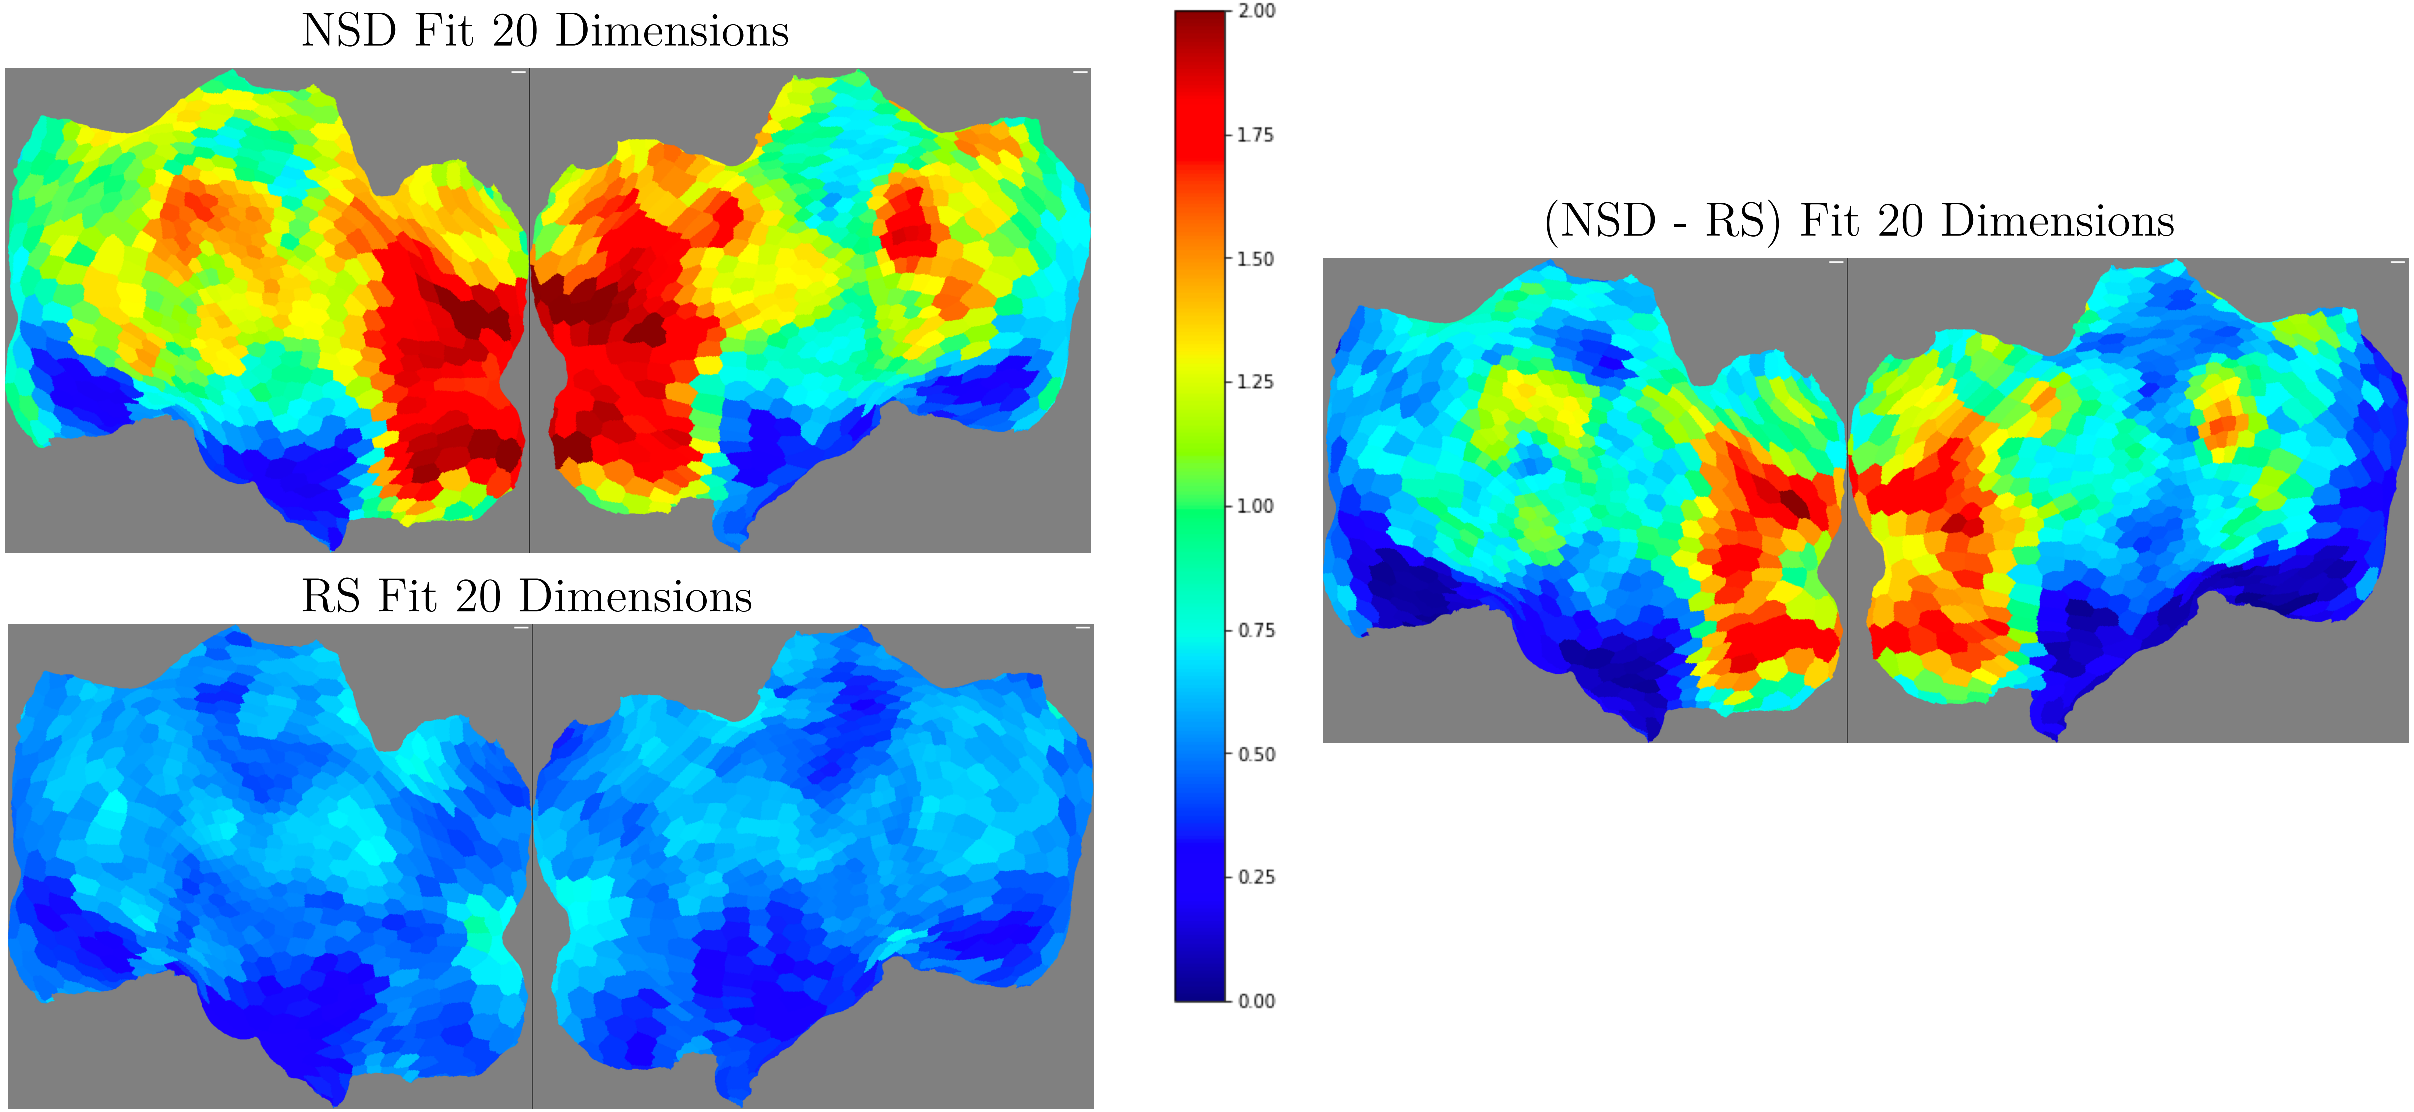

Open Questions

- Can we explain, using stimuli features, the low-dimensional activation structure seen in NSD PC maps?
- Are principal components, the orthogonal directions of maximal variation, "meaningful" in themselves?
- Is it meaningful to interpret the eigenspectral decay of fMRI data?
- What are the proper random controls?
- How might we introduce stimulus dimensionality back into such spectral analyses?