High-dimensional spectrum of reliable individual differences in visual cortex

Chihye Han (chan21@jhu.edu)
Department of Cognitive Science, 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

Human visual cortex represents sensory stimuli in population codes spanning thousands of latent dimensions. Although the variance along these dimensions is largely shared across individuals, a portion appears to reflect individual differences. We sought to determine if these individual differences are reliable and, if so, across how many dimensions. We examined a large-scale fMRI dataset of movie-viewing and characterized individual differences across the full spectrum of latent dimensions in visual cortex. We detected reliable individual differences spanning the almost entire spectrum of latent dimensions. Comparisons with voxelwise inter-subject correlations showed that our procedure reveals many dimensions of individual differences that are undetectable with conventional voxelwise analyses. Together, these findings reveal the surprisingly high-dimensional nature of individual differences in visual cortex, and they demonstrate an approach for exploring dimensions of individual variability that are unreachable with conventional methods.

Keywords: dimensionality; covariance spectra; naturalistic stimuli; individual differences; visual cortex; fMRI;

Introduction

Recent work has shown that human visual cortex contains reliable stimulus-related variance that spans thousands of latent dimensions (Gauthaman, Ménard, & Bonner, 2023), contrasting with earlier estimates that suggested a much lower dimensional representational code in primate visual cortex (Lehky et al., 2014; Op de Beeck, Wagemans, & Vogels, 2001). Most of the variance along these dimensions is shared across individuals, but a portion appears to be person-specific, suggesting that there may be important individual differences in the high-dimensional population codes of visual cortex. We sought to characterize these individual differences and rigorously assess their reliability. To accomplish this, we analyzed a large-scale naturalistic fMRI dataset using a cross-decomposition approach.

Methods

We analyzed a dataset of 40 subjects viewing naturalistic movies in an fMRI scanner (Sava-Segal et al., 2023). Preprocessed data from all movie sequences were concatenated, yielding 2,178 timepoints (36 min 18 sec) per subject. We focused on the ventral visual stream by examining an anatomically defined region including occipital and temporal cortex (Fig. 1A).

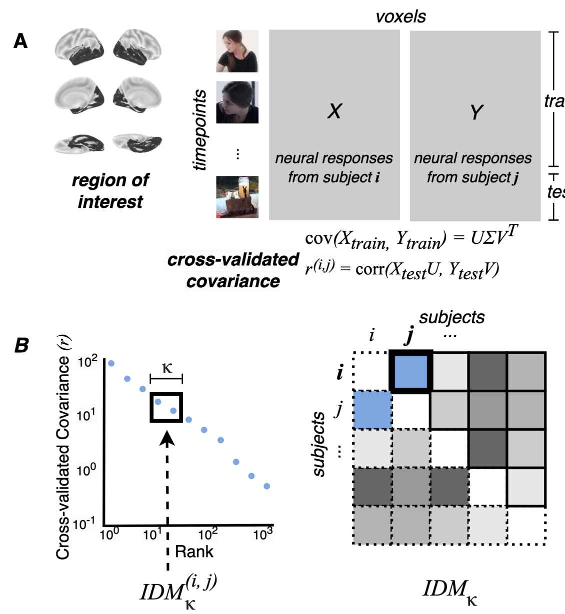

Figure 1: Illustration of our method. **(A)** Region of interest and cross-decomposition analysis. **(B)** (left) Cross-covariance spectrum of a pair of subjects, with κ indicating a segment of dimensions; (right) Individual differences matrix IDM_κ, constructed based on pairwise comparisons of subjects, reflects individual differences in representations across a bin of latent dimensions.

We characterized the spectrum of latent dimensions in these data using a cross-decomposition procedure described in previous work (Gauthaman et al., 2023). This procedure aligns the representations of different individuals along shared latent dimensions sorted by variance, and tests the reliability of these latent dimensions in held-out test data. Specifically, given two data matrices $X, Y \in \mathbb{R}^{t \times v}$ from a pair of subjects i and j with t timepoints and v voxels, we computed the singular value decomposition of their cross-covariance in a training set:

$$\text{cov}(X_{train}, Y_{train}) = U \Sigma V^T \quad (1)$$

The spectrum of reliable variance was then computed in the

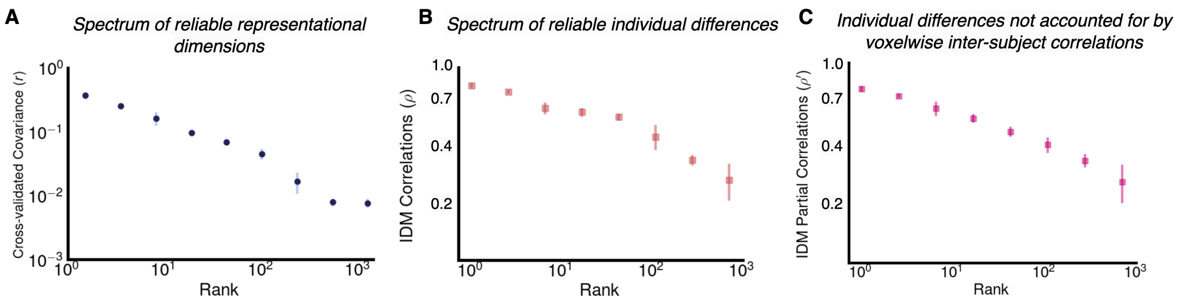

Figure 2: **A)** Cross-covariance spectrum observed for all subject pairs. Error bars indicate standard deviations of covariance spectra across subject pairs. **B)** Reliable individual differences measured by correlations between IDMs. **C)** Partial correlations between IDMs after removing the effect of voxelwise inter-subject correlations.

test set as:

$$r^{(i,j)} = \text{corr}\left(X_{test}U, Y_{test}V\right) \tag{2}$$

Each r_n is the Pearson correlation of the n-th dimension aligned across subjects. This procedure was applied using 12-fold cross-validation across all pairs of subjects (Fig. 2A).

Next, we asked whether there is reliable individual *variability* along the ranks of this high-dimensional spectrum. We constructed individual differences matrices (IDMs) that capture the pairwise similarities of subjects (Feilong et al., 2018). To examine individual differences at different ranks of latent dimensions, we constructed IDMs for bins of latent dimensions. For a given bin κ spanning dimensions l to u, each cell of an IDM was computed from a pair of subjects i and j as:

$$IDM_\kappa(i,j) = \sum_{n=l}^{u} r_n^{(i,j)} \tag{3}$$

An IDM for each bin, IDM_κ, aggregated these pairwise statistics, reflecting individual differences in representations across a segment of latent dimensions. We measured the reliability of these differences by first averaging the IDMs from even and odd test folds and then computing the Spearman correlation between these averaged matrices:

$$\rho_\kappa = \text{corr}\left(IDM_\kappa^{even}, IDM_\kappa^{odd}\right) \tag{4}$$

For comparison, we performed a parallel analysis of individual differences using a conventional method of computing voxelwise inter-subject correlations (ISCs) in anatomically aligned data. Specifically, we constructed IDMs by taking the mean voxelwise correlation of timepoints for each pair of subjects. We then examined whether this voxelwise-correlation IDM, IDM_{ISC}, could account for the individual differences that were reliably detected using our cross-decomposition approach. To do this, we computed partial correlations of the IDM reliability scores after partialling out the variance that could be accounted for by the voxelwise-correlation IDM:

$$\rho'_\kappa = \text{corr}\left(IDM_\kappa^{even}, IDM_\kappa^{odd} \mid IDM_{ISC}\right) \tag{5}$$

Results & Conclusion

As shown in Fig. 2B, we found that the IDM correlations were reliable across multiple orders of magnitude of latent dimensions (all $p < 0.0001$; permutation tests n=1000). These findings suggest that individual differences are not isolated to specific subspaces of visual representation but, instead, span the entire representational code. In addition, our cross-decomposition procedure reveals reliable variance across many latent dimensions that cannot be detected with a conventional voxelwise-correlation approach (Fig. 2C).

In sum, our findings reveal a high-dimensional spectrum of reliable individual differences in the cortical processing of natural visual stimuli. This work demonstrates that functional alignment techniques like cross-decomposition combined with large-scale, naturalistic data can allow neuroscientists to explore uncharted dimensions of individual variability in cortical representations.

References

Feilong, M., Nastase, S. A., Guntupalli, J. S., & Haxby, J. V. (2018). Reliable individual differences in fine-grained cortical functional architecture. *NeuroImage*, *183*, 375–386.

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. doi: 10.32470/ccn.2023.1652-0

Lehky, S. R., Kiani, R., Esteky, H., & Tanaka, K. (2014). Dimensionality of object representations in monkey inferotemporal cortex. *Neural computation*, *26*(10), 2135–2162.

Op de Beeck, H., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nature neuroscience*, *4*(12), 1244–1252.

Sava-Segal, C., Richards, C., Leung, M., & Finn, E. S. (2023). Individual differences in neural event segmentation of continuous experiences. *Cerebral Cortex*, *33*(13), 8164–8178.