

Potential cortical and computational biases in representational similarity analysis

Daniel D Leeds (dleeds@fordham.edu)

David Shutov (dshutov@fordham.edu)

Fordham University, 441 E Fordham Road
Bronx, NY 10458 USA

Abstract:

Representational similarity analysis (RSA) has become a valuable and common tool in the understanding of cortical representations across diverse cognitive arenas. However, RSA typically employs assumptions that may bias model comparisons. Our present work identifies common statistics of cortical responses in object perception and finds that these responses may support inflated model comparison results with unusual resistance to noise. Similarly, we find differing constructions of permutation tests alter perceived significance of model-cortical matches. We employ an fMRI voxel searchlight method to compare local cortical responses to sixty objects, with 218 diverse candidate semantic groupings of the same objects. We find semantic properties with the highest cortical correlations are high skew distance matrices, while the lowest cortical correlations are often low skew. We also find additional restrictions on “randomized” permutations may be required for more accurate assessment of statistically significant matches in RSA.

Keywords: RSA, object perception, fMRI

Background

Representational similarity analysis (RSA) is a valuable tool to observe and model complex patterns in cortical information processing (Kriegeskorte 2008). For a selected brain region or computational model, a pairwise distance matrix can be computed to reflect what stimuli are grouped together or set apart. These groupings may be compared to identify the best model matching cortical behavior. The approach of matching distance matrices allows us to judge a model’s descriptive power without requiring an exact (likely non-linear) mapping between model and cortical responses to a base set of stimuli. Accordingly, the RSA approach has gained substantial traction, e.g., in studying the link between computer vision models and biological vision. (e.g., Leeds 2013, Khaligh-Razavi 2014) RSA has been used successfully across species, recording modalities, and cortical regions. (Kriegeskorte 2008,

Leeds 2013, Devereaux 2013) Our present work benefits from RSA in identifying a common class of semantic representations during visual object perception. However, we also find evidence that the unifying statistic of these semantic models may have an undue advantage in apparent strong fMRI representational correlations. We also find the typical assessment of significance through permutation test requires additional constraints when applied to RSA.

Methods

fMRI data collection

We study fMRI BOLD data recorded by Leeds (2013), obtained from three subjects recruited from the Carnegie Mellon University Community. Subjects view 60 visually and semantically diverse object stimuli - including mammals (e.g., dog, bear), vehicle (e.g., car, plane), tools (e.g., hammer, spoon), dwelling places (e.g., house, apartment building), etc.. Each stimulus was represented through a single corresponding photograph and corresponding word label - displayed during separate, disjoint trials. Stimuli are viewed passively during a fixation onset task.

Representational similarity analysis

Cortical representations of objects are defined with respect to voxel searchlights, centered at each location in the cortex. At each location (x,y,z) , a pairwise representational distance matrix (RDM) is defined as one minus the Spearman correlation between the voxel responses for stimuli s^i and s^j , or

$$D_{x,y,z}^{\text{searchlight}}(s^i, s^j) = 1 - r(v(s^i), v(s^j)) \quad (\text{Eqn. 1})$$

The vector $v(s^i)$ represents the voxel responses for stimulus s^i (Leeds 2013, Kriegeskorte 2008). RDMs are constructed based on searchlights of radius 3 voxels.

Model semantic representations of objects are constructed for comparison with cortical representations. Semantic models are derived from each of 218 questions, spanning from diverse sensory and conceptual topics such as identity ("Is it a tool?") to component identities ("Does it have legs?") to size/weight ("Is it bigger than a loaf of bread?") to and emotion ("Is it scary?"). For each question, subjects provide ratings for each object on a scale from 1 to 5 (definitely no to definitely yes) through Amazon Mechanical Turk (Sudre 2012). At least three subjects provided a rating for each object.

Cortical and semantic-model representations were compared by converting the lower triangle of each corresponding 60 x 60 RDM into a 1770 x 1 vector and measuring Spearman correlations between the searchlight and model matrices.

RSA permutation testing

Permutation tests are further performed for each computed voxel searchlight-model correlation. Ordinarily, the values from the 1770 entries in the lower triangle of the distance matrix would be randomly permuted, preserving the frequency of each measured "distance", but not its place in the matrix. However, a randomly permuted matrix of this form will not guarantee fundamental distance properties, e.g., $d(a,c) \leq d(a,b) + d(b,c)$. In the present study, we instead begin with the ratings of each of the sixty objects, considered across all 218 semantic questions. For each object, we randomly select a rating from one of the 218 questions. After all ratings are randomly selected, we construct a new pairwise distance matrix and compute the correlation at each voxel sphere location. The randomized distance matrix construction and correlation process was repeated 50 times. The permutation test process first was repeated 500 times for five selected questions; Z-score results for 50 and 500 randomized matrices were found to be comparable.} After permutation tests are completed, Z-scores are calculated.

Results

We observe searchlight-model RDM correlations to be highest in cortical regions typically associated with mid-level visual perception - consistent with past cortical-semantic results (e.g., Sudre 2012, Devereaux 2013). We further observe strength of searchlight-

model correlations varies based on the skew of the semantic model tested. High skew in representations are common across early and mid-level vision, while low-skew elsewhere in the brain is more susceptible to degraded cortical-model correlation with noise. We study these cortical statistics and their interactions with RSA in the sections below.

High and Low-Skew Semantic Rankings

We observed a connection between the statistics of a semantic model's distance matrix and the magnitude of its correlation with voxel representations. The 218 semantic features adapted from Sudre (2012) divide in skew of answers --- e.g., "Is it an insect?" is a high skew question as most objects are not machine-like and a few are machines (most objects rated 1, a few rated higher), but "Is it found in school" is a low-skew question as roughly equal number of objects are hot, cold, and in-between (evenly rating between 1, 2, 3, 4, and 5; **Figure 1**). These skews remain when answers are recorded for a larger set of 1000 objects.

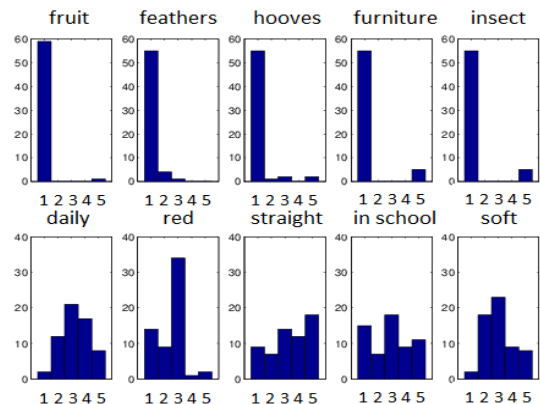


Figure 1: Distribution of Mechanical Turk answers for five questions with skew greater than 1 (top) and less than 1 (bottom).

We find semantic features/models with the highest skew answers also have higher cortical correlations on average; models with the lower skew have lower cortical correlations on average (**Figure 2a**). (Note: we measure the absolute value of the skew for each model.) Conversely, models with the highest correlations largely have the highest skews (**Figure 2b**). Given the semantic diversity of stimuli, our results indicate a preference for single-object-category activation rather than a representation of a continuum of property scales.




Figure 2: (a) Distribution of maximum cortical correlations for models with skew greater than 1 (top) and less than 1 (bottom) (b) Distribution of skews for models with maximum correlation greater than 0.3 (top) and less than 0.3 (bottom)).

Ignoring candidate semantic models, analysis of voxel searchlights across the brains of all subjects reveals high-skew representations in early and mid-level vision (shown in red for subject S1 in **Figure 3**) and lower skew representations across the rest of the brain (areas of no color overlay in **Figure 3**). While low skew is common across higher-level vision and non-vision areas, correlation with low-skew semantic models is relatively weak.




Figure 3: Distribution of skews greater than 0.5 across ventral slices of brain for subject S1.

Noise effects on RSA correlations

We explore whether distance matrices from high skew responses may have unduly inflated correlations with high skew cortical data. Specifically, we select the sixty object ratings for a high-skew semantic model (“has paws?”) and a low-skew semantic model (“hard inside?”) as ground truths and also create 100 copies of the object ratings for each model perturbed by Gaussian noise. We generate the resulting RDMs (e.g., **Figure 4a**) and compare ground-truth matrices

to noise matrices. We observe substantially higher post-noise correlations when the ground-truth representation has higher skew (**Figure 4b**). Intuitively, for high-skew models, the few objects with the unusually high (or unusually low) response will contribute the largest values in the distance matrix. Gaussian perturbations may bring high responses slight lower and low responses slightly higher, but the few “different” objects will still stand out. In contrast, low-skew models are dominated by a wealth of smaller inter-object distances; enough smaller perturbations in these differences would lower the post-noise correlations.

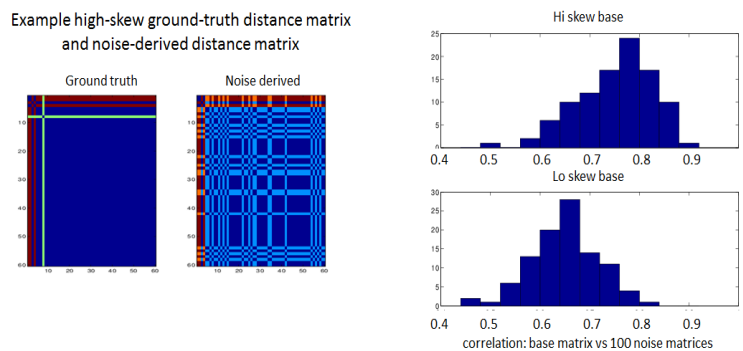


Figure 4: (a) Example of high-skew ground-truth and Gaussian-perturbed distance matrix (RDM) for sixty objects. (b) Distribution of correlations for noisy versus ground-truth representations of sixty objects based on high skew ground-truth object ratings (top) and low skew rating (bottom)..

Evaluating assumptions of the permutation test

Beyond initial correlation, we consider the approximated significance of model-cortical matches through Z scores derived by permutation testing. We use both permutation of pairwise distance entries and permutation of single-object ratings, comparing the resulting computed Z scores. The variance of these Z-scores substantially increases with the magnitude of semantic-cortical correlations (Figure 5). Notably, a higher variance is found between correlation and Z-scores when using traditional permutation testing, in which model matrix entries are randomly permuted without regard to distance matrix structure. The incorporation of inherent constraints on distance matrices produced a narrower spread of correlations for a given Z-score.

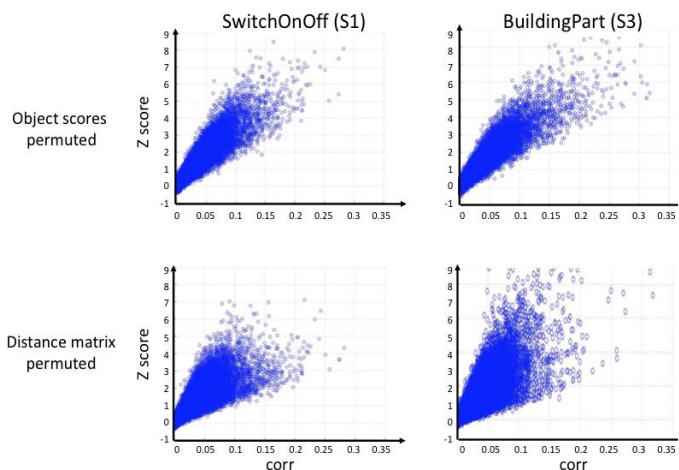


Figure 5: Comparison of positive correlation values and Z scores computed based on random permutation of object scores (**top**) and on random permutations on entries of distance matrix (**bottom**). Correlation-Z score spreads shown for two high-correlation semantic features. Constructing distance matrices after permuting object scores results in tighter distribution of Z scores across voxel with a given model correlation.

Discussion

Representational similarity analysis study of candidate semantic models underlying visual object perception shows a preference for multiple skewed groupings of a semantically diverse set of sixty objects in early to mid-level visual regions. Our results indicate a cortical preference for single object categories in early to mid-level vision. Semantic representations in additional cortical regions do not appear to be as strongly modeled by the simple single-question models adapted we adapted from Sudre (2012).

Additional study on the statistics of distance matrices indicates a possible “unfair advantage” for skewed distributions during RSA model comparisons, which may affect the results of our study as well as other ongoing studies. Skewed distributions show greater robustness to Gaussian noise, commonly expected while studying neuroimaging data. Significance analysis of distance matrix correlations also may be inflated or deflated without proper consideration of matrix statistics during the popular permutation test approach.

Our work stresses the importance of carefully framing future analyses to properly incorporate RSA’s

sensitivity to model statistics **and** to intrinsic distance matrix structure.

References

- Devereux, B. J., Clarke, A., Marouchos, A., & Tyler, L. K. (2013). Representational Similarity Analysis Reveals Commonalities and Differences in the Semantic Processing of Words and Objects. *Journal of Neuroscience*, 33(48). doi:10.1523/jneurosci.3809-13.2013
- Khaligh-Razavi, S., & Kriegeskorte, N. (2014). Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation. *PLoS Computational Biology*, 10(11). doi:10.1371/journal.pcbi.1003915
- Kriegeskorte, N., Mur M., Bandettini P. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4). doi:10.3389/neuro.06.004.2008
- Leeds, D. D., Seibert, D. A., Pyles, J. A., & Tarr, M. J. (2013). Comparing visual representations across human fMRI and computational vision. *Journal of Vision*, 13(13), 25-25. doi:10.1167/13.13.25
- Sudre, G., Pomerleau, D., Palatucci, M., Wehbe, L., Fyshe, A., Salmelin, R., & Mitchell, T. (2012). Tracking neural coding of perceptual and semantic features of concrete nouns. *NeuroImage*, 62(1), 451-463. doi:10.1016/j.neuroimage.2012.04.048