

Statistical inference on representational geometries

Heiko H. Schütt (hs3110@columbia.edu)¹

Zuckerman Institute, Columbia University, New York City, NY 10027, USA

Alexander D. Kipnis (alexander.kipnis@tue.mpg.de)²

Zuckerman Institute, Columbia University, New York City, NY 10027, USA

Jörn Diedrichsen (joern.diedrichsen@gmail.com)

Western University, London, Ontario, Canada

Nikolaus Kriegeskorte (nk2765@columbia.edu)

Zuckerman Institute, Columbia University, New York City, NY 10027, USA

Abstract

Representational similarity analysis is a versatile method for comparing high dimensional models and neural recording data to each other. Here, we introduce a comprehensive new set of methods for statistical model comparison based on predictions of representational geometries. The inference can handle flexible parametrized models and can treat both subjects and conditions as random effects, such that conclusions generalize to the respective populations of subjects and conditions. With crossvalidated representational distance estimators and metric whitened model evaluators, the power for model comparisons approximates that of likelihood-based inference, but rank-based model evaluation is also supported. We validate the inference methods using extensive simulations with deep neural networks and resampling of calcium imaging and functional MRI data. Results demonstrate that the methods are valid and conclusions generalize correctly. These data analysis methods are available in an open-source Python toolbox.

Keywords: Representational Similarity Analysis; RSA;
Methods: Statistical Inference

Introduction

A popular method for analyzing representational geometries (Chung & Abbott, 2021; Kriegeskorte & Kievit, 2013; Kriegeskorte & Wei, 2021; Shepard & Chipman, 1970) is representational similarity analysis (RSA; Kriegeskorte & Diedrichsen, 2019; Kriegeskorte, Mur, & Bandettini, 2008; Nili et al., 2014; Storrs, Kietzmann, Walther, Mehrer, & Kriegeskorte, 2021; Walther et al., 2016). RSA is a two step process (Fig. 1a, b): In the first step, RSA characterizes the representational geometry by estimating the dissimilarity of each pair of experimental conditions (e.g. different stimuli), and assembles these in a representational dissimilarity matrix (RDM). An RDM is computed for the neural population in our brain region of interest and for each model representation. In the second

step, each model is evaluated by the accuracy of its prediction of the data RDM.

Here, we introduce a comprehensive methodology for statistical inference on models that predict representational geometries (Fig. 1b). We propose bootstrapping methods that can statistically support generalization to new subjects, new conditions, or both simultaneously, as required to support the theoretical claims researchers wish to make. We extend these bootstrapping methods with crossvalidation to enable inference on flexible models, i.e. models with parameters fitted to the data. The methodology is available in a new open-source RSA toolbox in Python (<https://github.com/rsagroup/rsatoolbox>).

Methods

Our general approach is bootstrap-wrapped cross-validation, i.e. we use bootstrapping to estimate the variability of model evaluations that may be cross-validated to prevent overfitting (Fig. 1b). For both the cross-validation and the bootstrap, we can choose whether we apply them across subjects, conditions, or both. Our choice for the cross-validations determines across which factor(s) the model needs to generalize with the same parameters. Our choice for the bootstrap determines how far our statistical inference is meant to generalize.

When both subjects and conditions are bootstrap-resampled simultaneously, the raw variance estimate is much higher than the true variance over repeated experiments with new samples of both subjects and conditions (Fig. 1d, e). We introduce a correction based on bootstrapping of each factor separately and both together. The variance decomposition enables us to estimate the true variance accurately.

Additionally, the choice of cross-validation folds adds variance to the evaluations, which include cross-validation. One could remove this variance by performing many cross-validations, but doing this for each random sample of the bootstrap would be prohibitively expensive. Instead, we found a correction formula that allows us to remove this excess variance accurately with only two cross-validation fold-assignments for each bootstrap sample.

Applied together, these corrected methods yield accurate estimates of the (co-)variances of model evaluations. Based on this (co-)variance estimate we can perform all statistical

¹Also at: Center for Neural Science, New York University, New York, USA

²present address: Max Planck Institute for Biological Cybernetics, 72076 Tuebingen, Germany



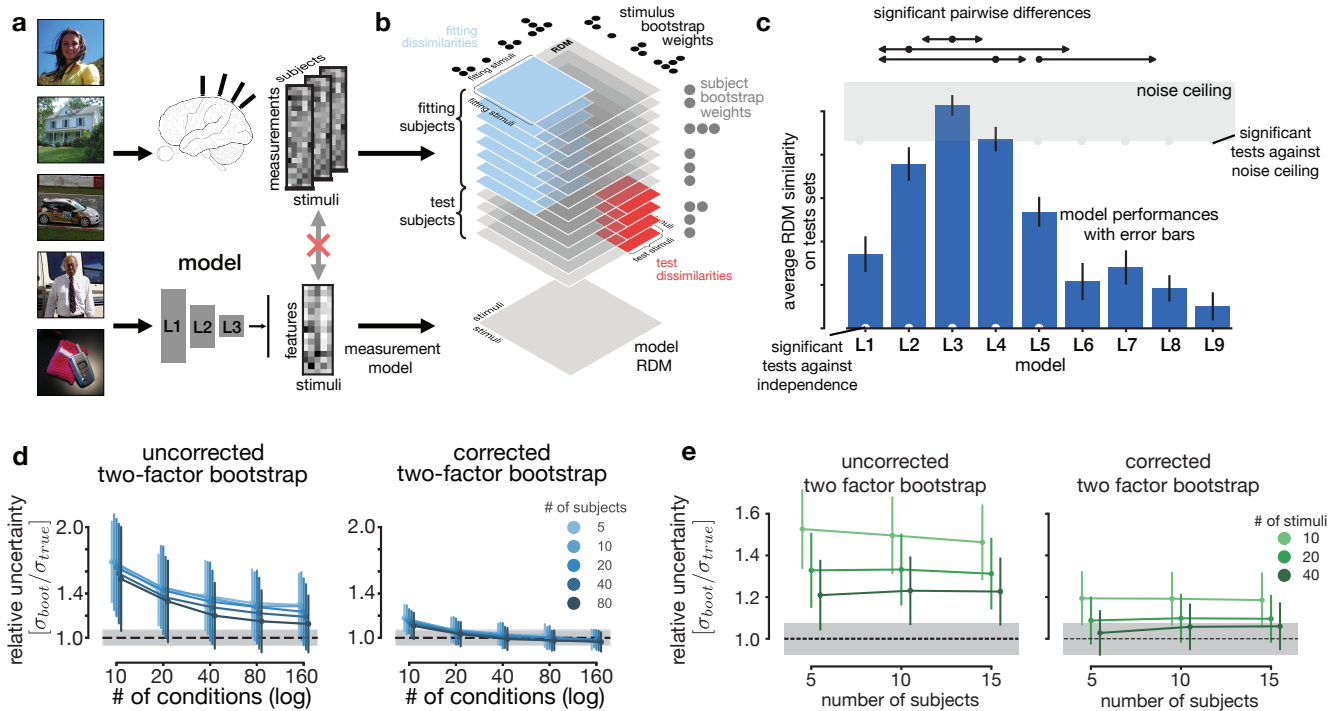


Figure 1: **New RSA methods for model evaluation.** **a:** The same set of stimuli or conditions is presented to the models and to experimental subjects. Recordings of neural activity and feature extractions from the models yield representational spaces, which are not directly comparable. **b:** The two-factor bootstrap-wrapped cross-validation. **c:** A typical model comparison result with statistical tests marked, here across layers L1-L9 of a deep neural network model. **d:** Deep neural network based simulation result: Comparison of the uncertainty estimates obtained with or without correction from the bootstrap against the true variability of results. The relative uncertainty should be close to 1, the shaded region is the range expected for a perfect method due to the limited number of simulations. **e:** As in panel d, but based on the resampling calcium imaging data.

tests we are interested in as t-test variants (Fig. 1 c).

Results

We thoroughly validate the new inference methods using simulations and neural activity data. Extensive simulations based on deep neural network models and models of the measurement process, where ground truth is known, confirm the validity of the inference procedures and their ability to generalize to the populations of subjects and/or conditions. Using real data from fMRI (human) and calcium imaging (mouse), we confirm that conclusions generalize from an experimental data set (subset of real data) to the entire data set (which serves as a stand-in for the population). In all simulations we used the *relative uncertainty*, i.e. the ratio of estimated standard deviation σ_{boot} and true standard deviation σ_{true} as the criterion, which should be close to 1.

In both DNN based simulations (Fig. 1d) and resampling based simulations (Fig 1e), we found that our corrected two-factor bootstrap method yields accurate estimates of the variance across the simulated populations of subjects and conditions when the dataset is large enough (≥ 20 subjects, ≥ 40 conditions) and the type of bootstrap matches the desired level of generalization.

To test our method for inference on flexible models, we made a variant of the deep neural network simulation in which we assume that the size of the voxels' averaging pools and the weighting of features are unknown to the researcher and need to be fitted to the data.

We found that our bootstrap-wrapped crossvalidation with the two-factor bootstrap and the excess-variance correction yielded accurate estimates of the uncertainty (not shown). The relative uncertainties were close to 1. Furthermore, flexible models that allowed fitting of the measurement model accurately selected the underlying representational model, while fixed models that ignored the unknown degree of pooling in voxels could lead to consistently incorrect inferences and/or low model discriminability.

Conclusion

We present a comprehensive new methodology for inference on models of representational geometries that enables neuroscientists to draw conclusions that generalize to new subjects and conditions, can handle flexible models, and is more powerful than previous approaches. The validity of the methods has been established through extensive simulations and using real neural data.

References

- Chung, S., & Abbott, L. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. *arXiv preprint arXiv:2104.07059*.
- Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the onion of brain representations. *Annual review of neuroscience*, *42*, 407–432.
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412.
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*.
- Kriegeskorte, N., & Wei, X.-X. (2021). Neural tuning and representational geometry. *Nature Reviews Neuroscience*, *22*(11), 703–718.
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLOS Computational Biology*, *10*(4), e1003553.
- Shepard, R. N., & Chipman, S. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive psychology*, *1*(1), 1–17.
- Storrs, K. R., Kietzmann, T. C., Walther, A., Mehrer, J., & Kriegeskorte, N. (2021). Diverse Deep Neural Networks All Predict Human Inferior Temporal Cortex Well, After Training and Fitting. *Journal of Cognitive Neuroscience*, 1-21.
- Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *NeuroImage*, *137*, 188–200.