Session VI: Representational Similarity Analysis

Martin Hebart

Laboratory of Brain and Cognition

NIMH

Goals of this Presentation

Basics of Representational Similarity Analysis

- What is RSA and why do it?

How does RSA work?

- Representational Geometry and Construction of Dissimilarity Matrices

General Considerations for RSA

- Distance metrics Euclidean vs. Correlation
- Distance Estimates are Biased
- Are Dissimilarities Ratio Scale?
- Multivariate Noise Normalization
- Noise Ceilings

What is Representational Similarity Analysis?

A multivariate pattern analysis method to investigate the *content* and *format* of representations

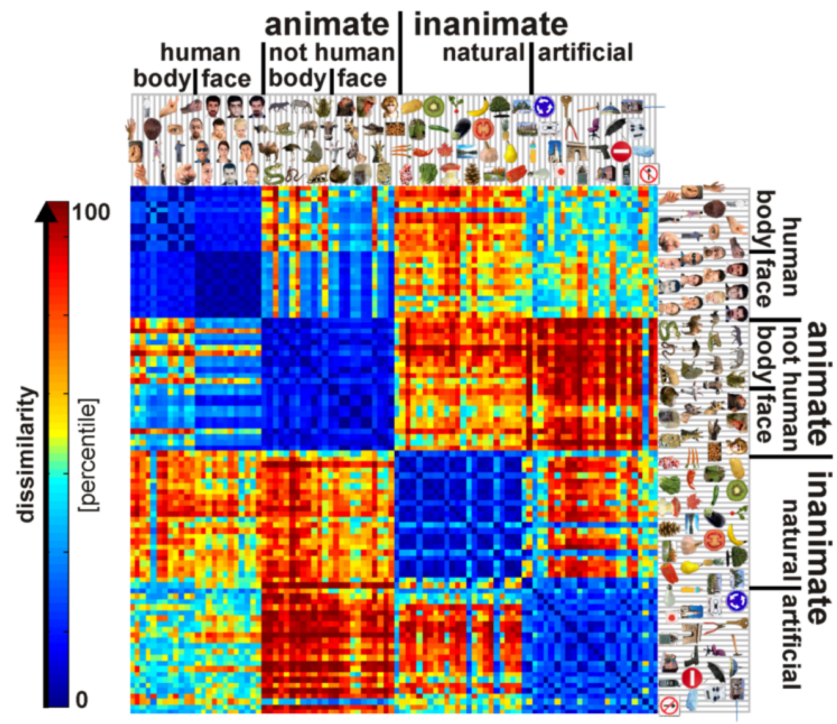

Kriegeskorte et al (2008) – Neuron; Kriegeskorte & Kievit (2013) – TICS

Why do Representational Similarity Analysis?

1. Simple exploratory approach to characterize **multidimensional representations**

Example: How does human IT represent object categories?

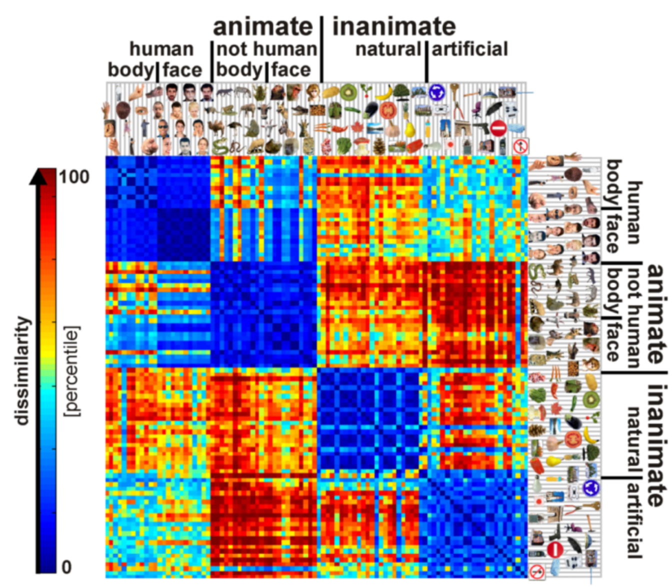

Kriegeskorte et al (2008) – Neuron; Kriegeskorte & Kievit (2013) – TICS

Why do Representational Similarity Analysis?

2. Representational (dis)similarity matrices can be seen as a **common language to study representations** across methods (MEG, fMRI, cell recording, …), brain regions, humans and species

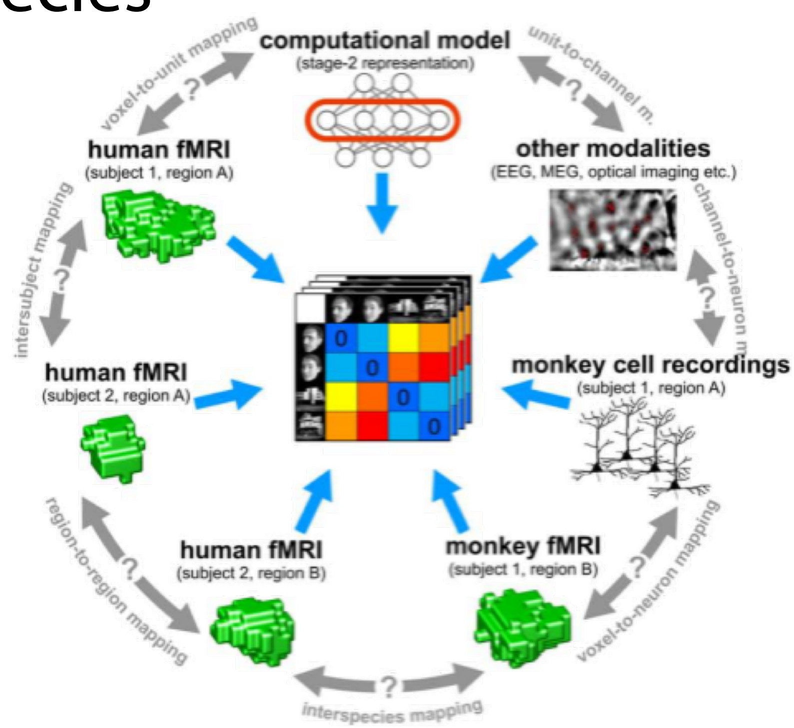

Kriegeskorte et al (2008) – Frontiers; Kriegeskorte & Kievit (2013) – TICS

Why do Representational Similarity Analysis?

3. Representational similarities can be used for **testing models of cognition**

Example: Which facial features does a brain region represent?

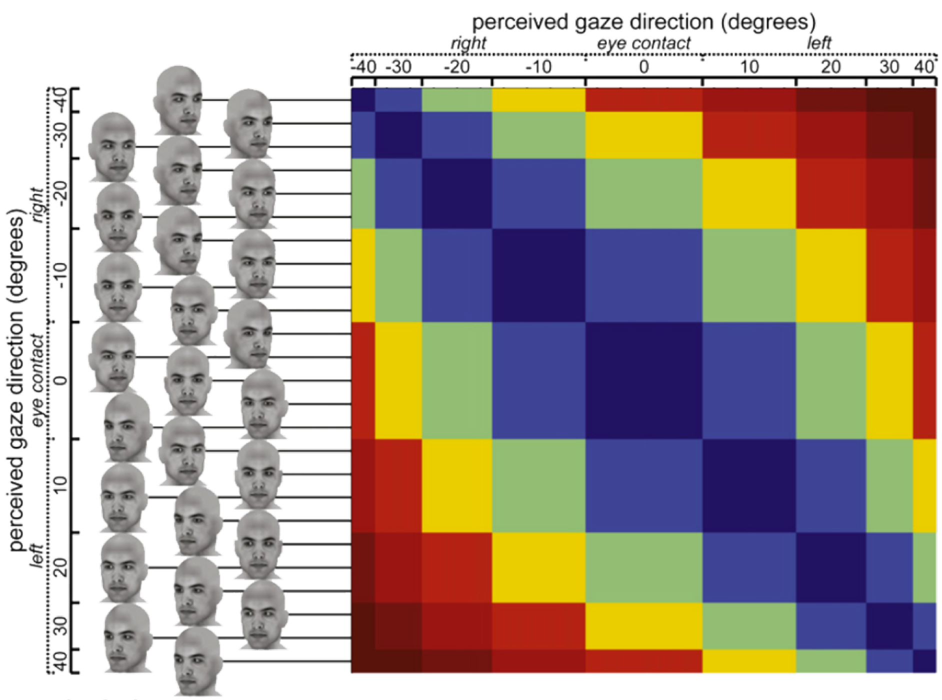

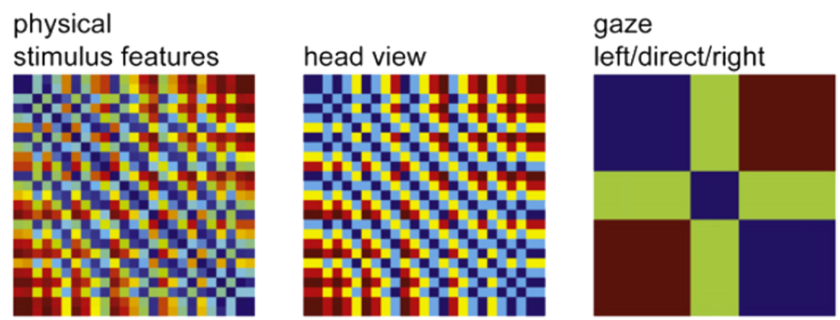

Carlin et al (2011) – Curr Biol

Why do Representational Similarity Analysis?

3. Representational similarities can be used for **testing models of cognition**

Example: Which computational model best explains responses in human IT?

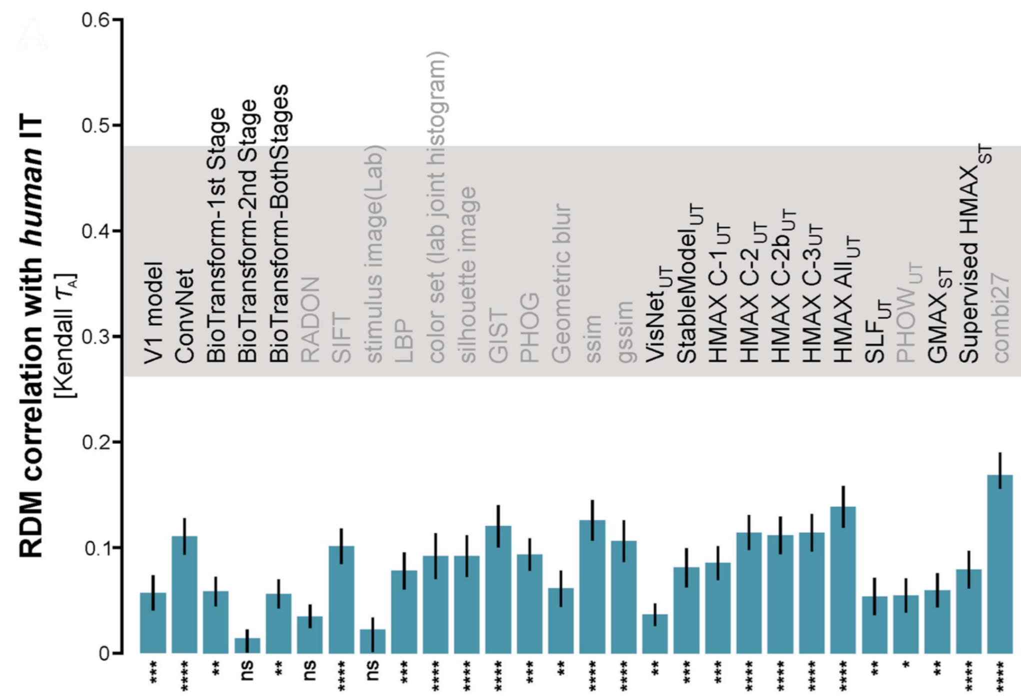

Khaligh-Razavi et al (2014) – PLoS Comput Biol

HOW DOES REPRESENTATIONAL SIMILARITY ANALYSIS WORK?

RSA: Linking Data at the Representational Level

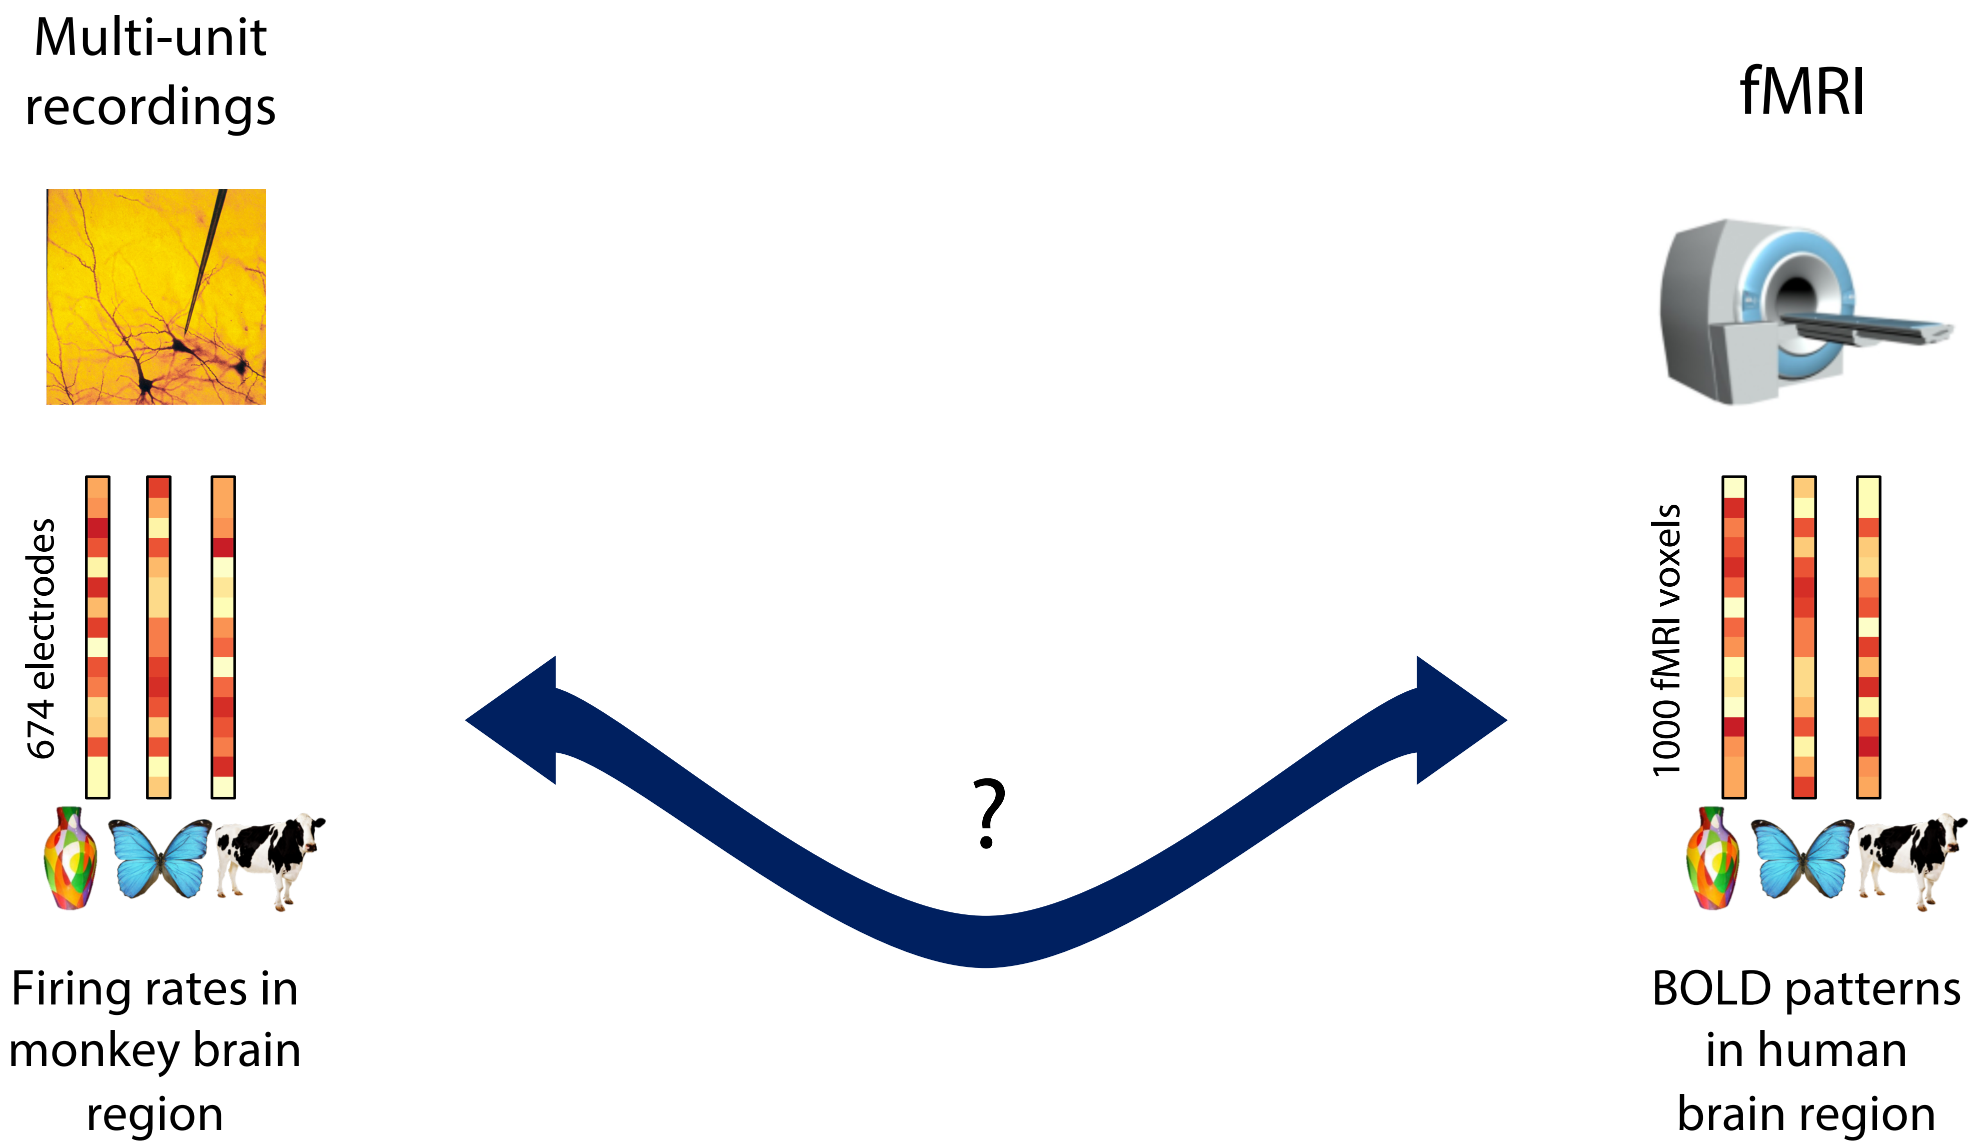

RSA: Linking Data at the Representational Level

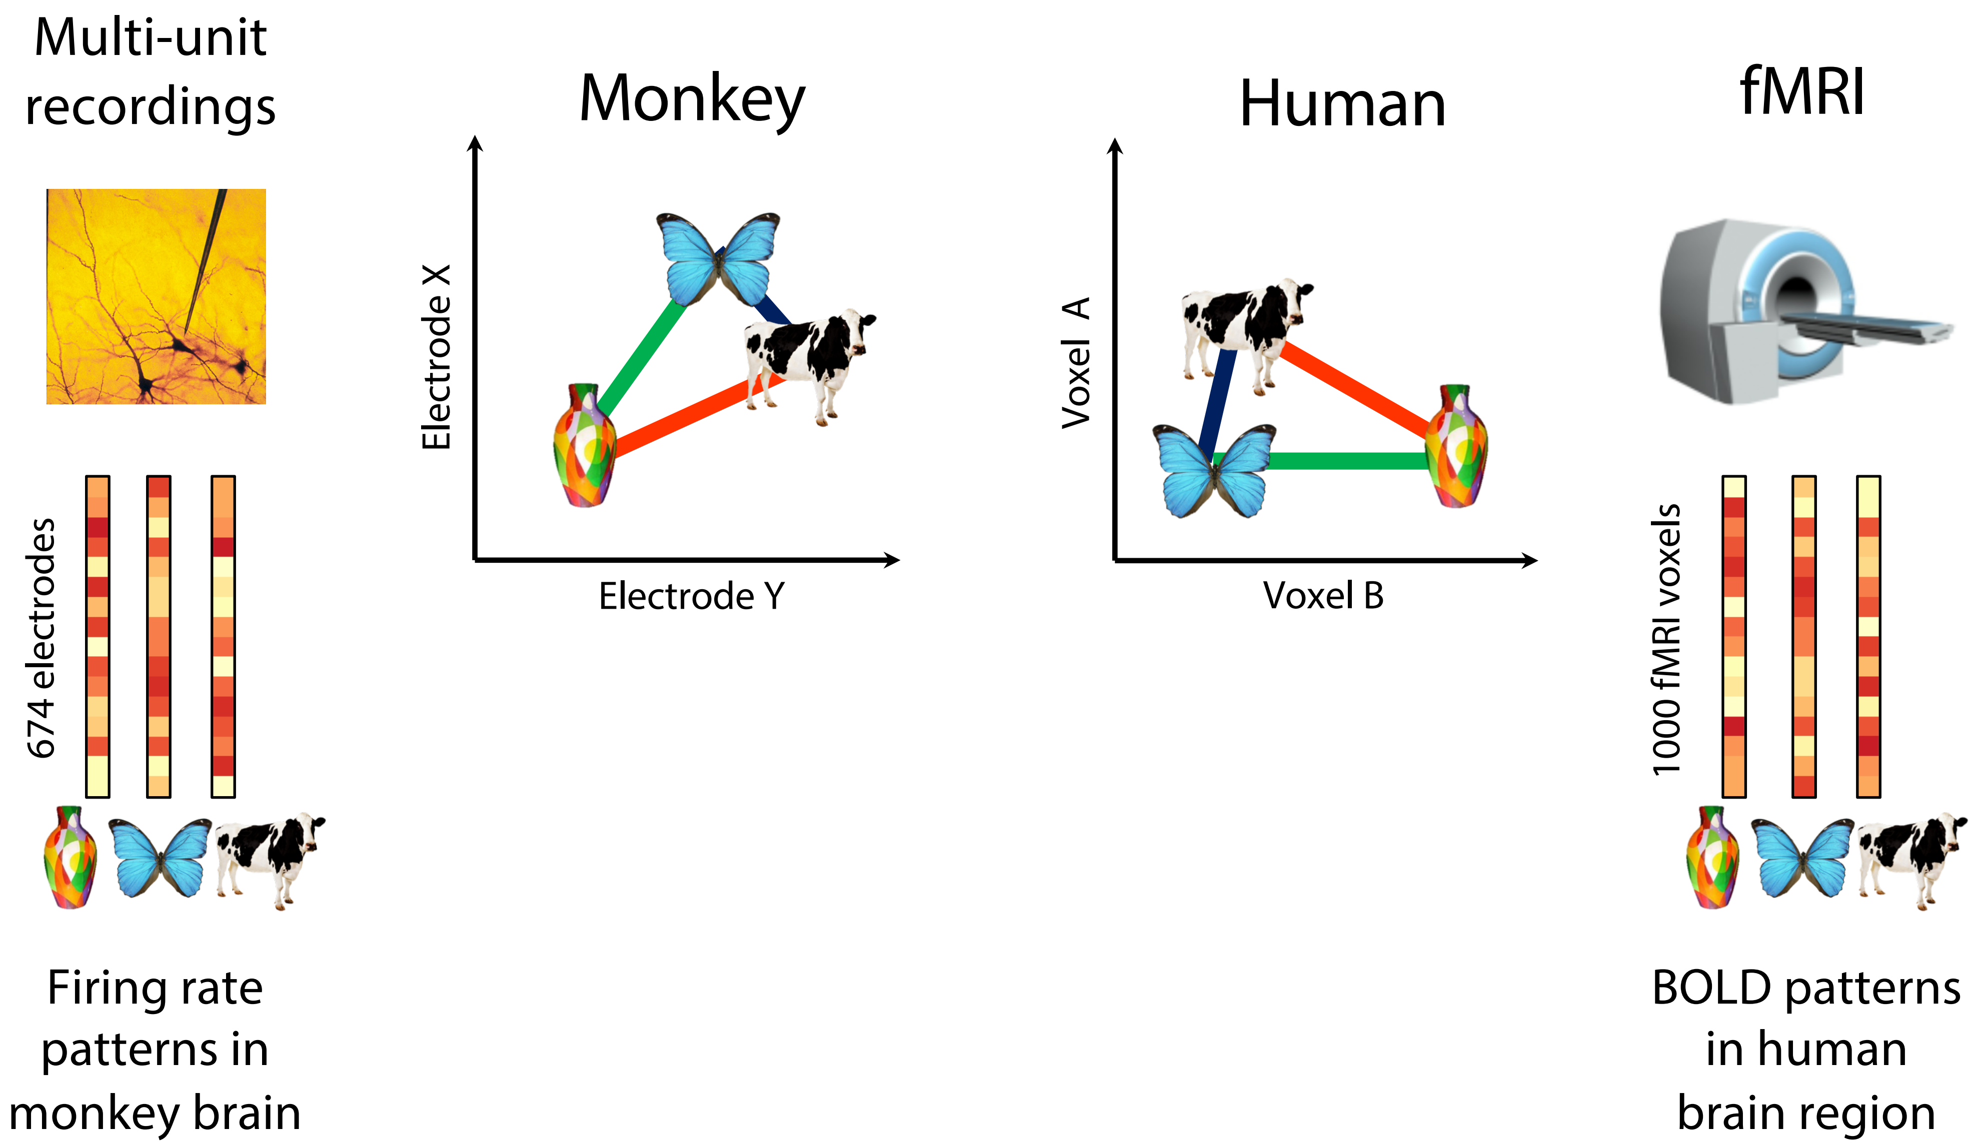

RSA: Linking Data at the Representational Level

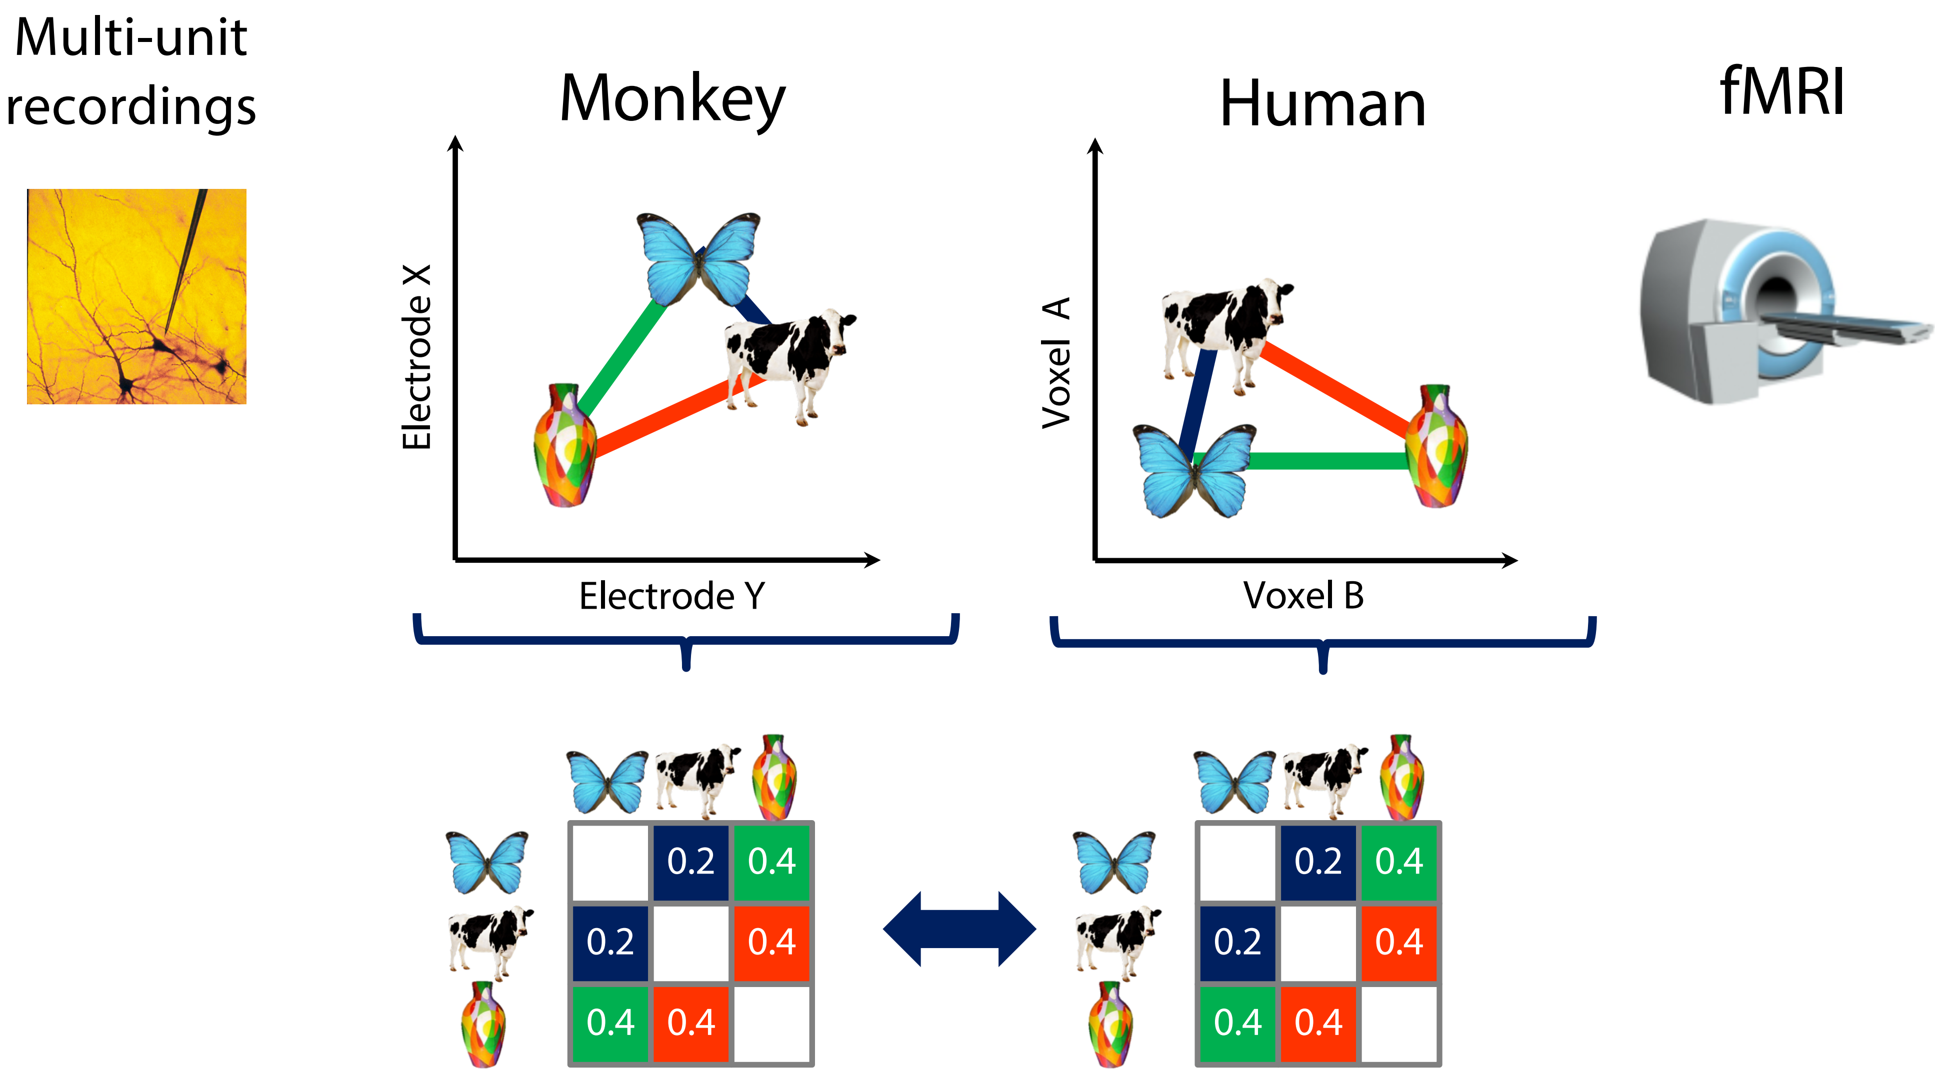

Representational Geometry (1)

- Representations can geometrically be interpreted as being embedded in a multi-dimensional space
- One particular representation is a point in this space and is a combination of unknown representational features

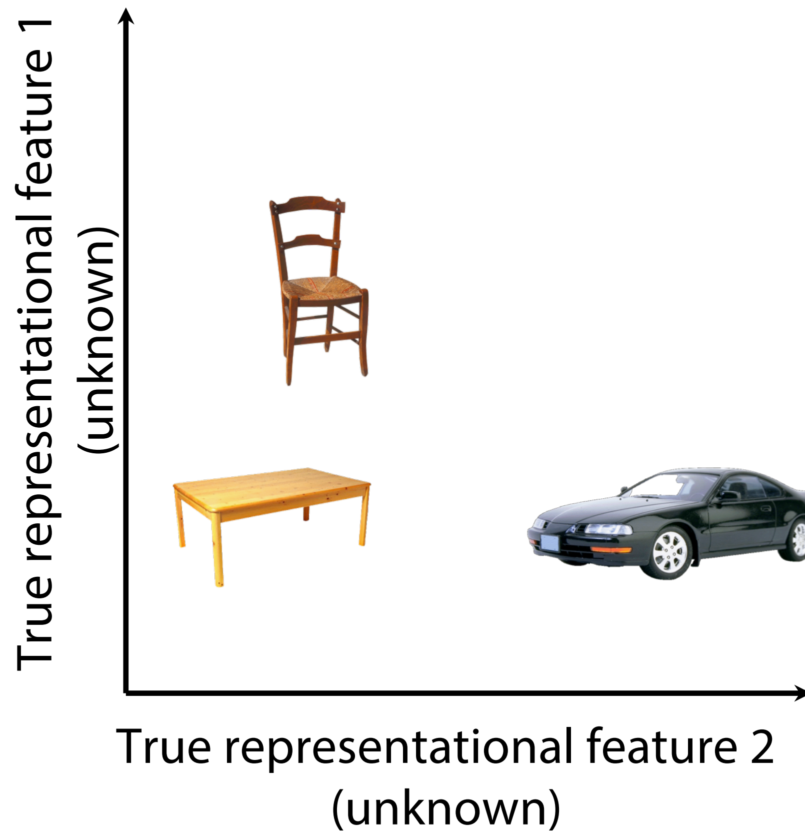

Representational Geometry (2)

- We don't know the true high-dimensional representational geometry, because we don't know the representational features
- We can describe the geometry by the relative distance between each pair (e.g. chair is closer to table than to car)
- We can measure this geometry with our recordings (= slices through representational space)

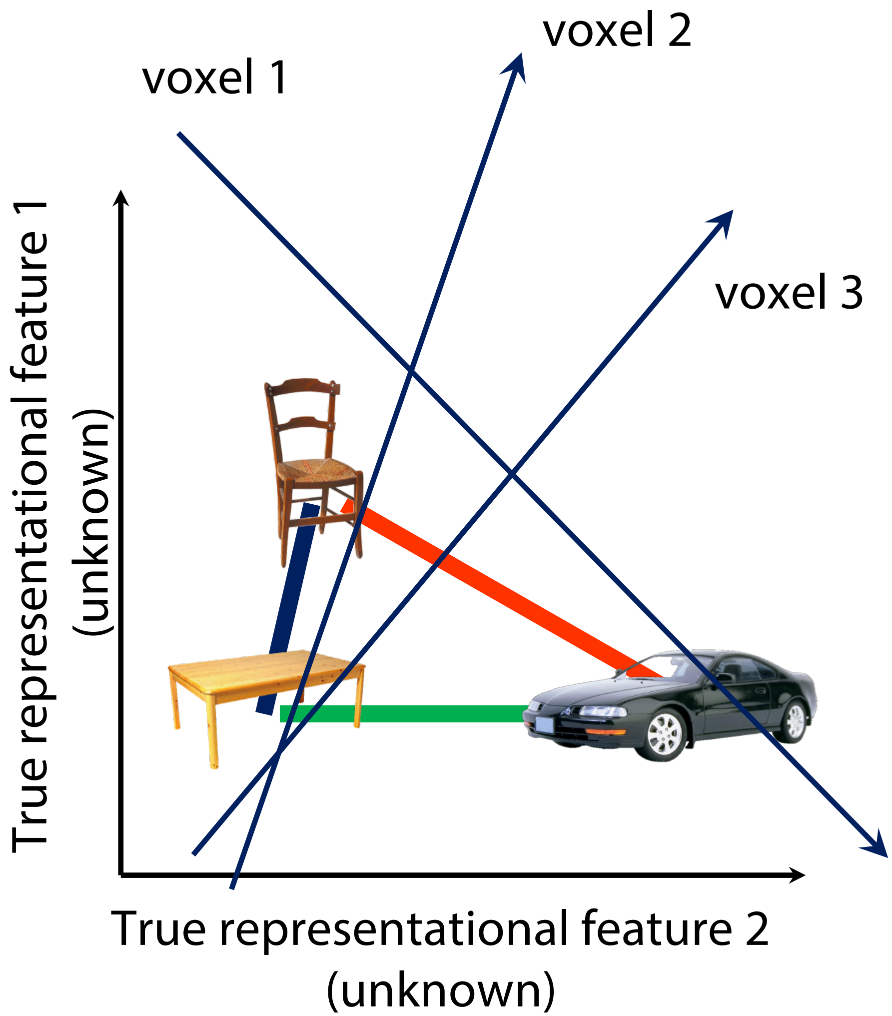

How does RSA work?

- In RSA, we take our multivariate patterns (e.g. voxels) and calculate pairwise dissimilarities (e.g. Euclidean distance or 1 – Pearson's *r*)

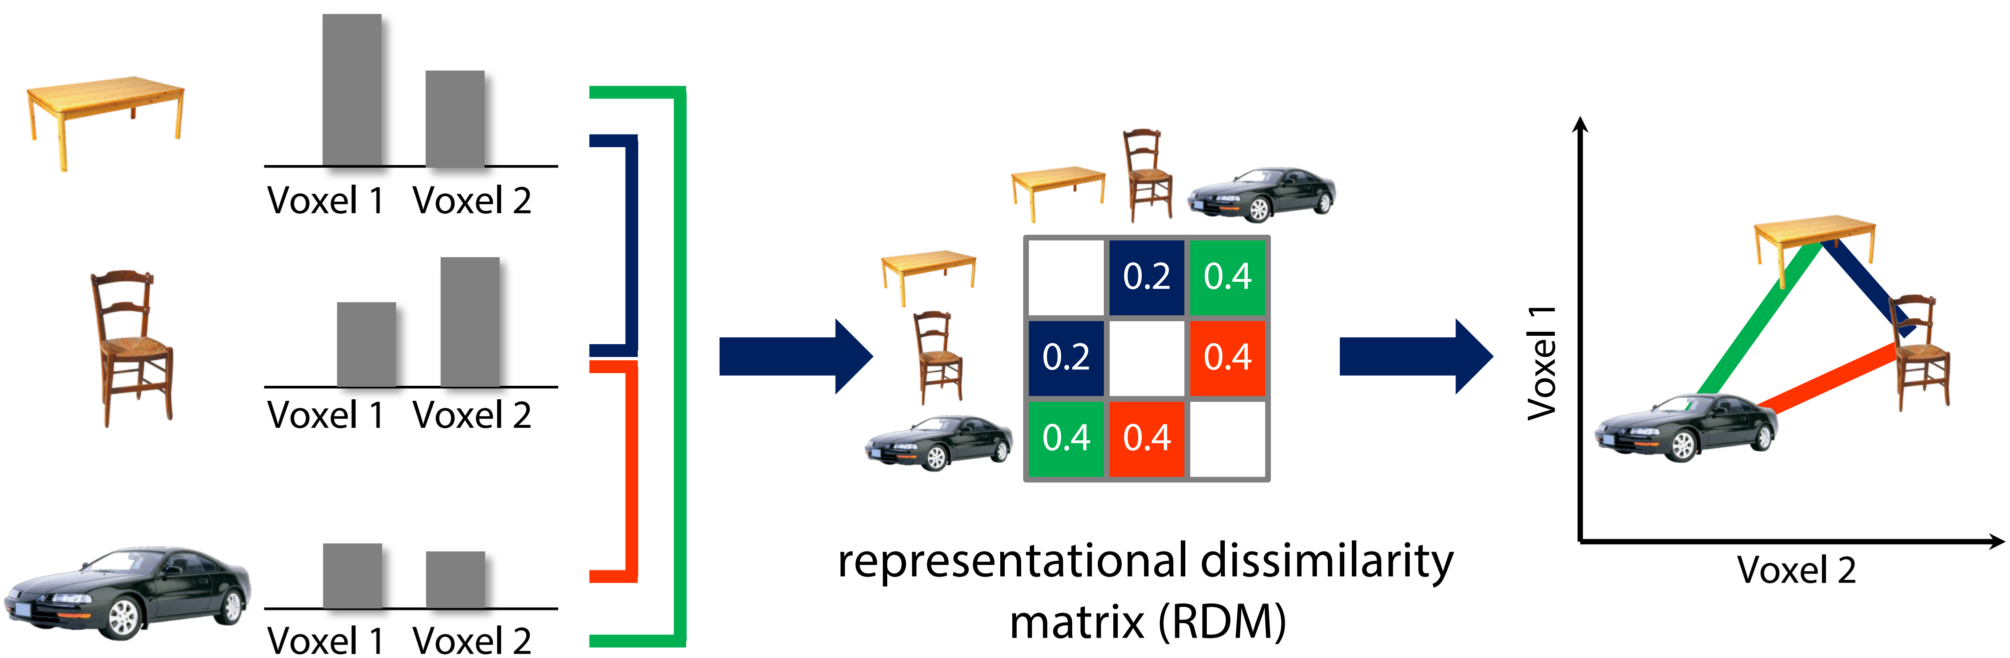

How does RSA work?

- In RSA, we take our multivariate patterns (e.g. voxels) and calculate pairwise dissimilarities (e.g. Euclidean distance or 1 – Pearson's r)

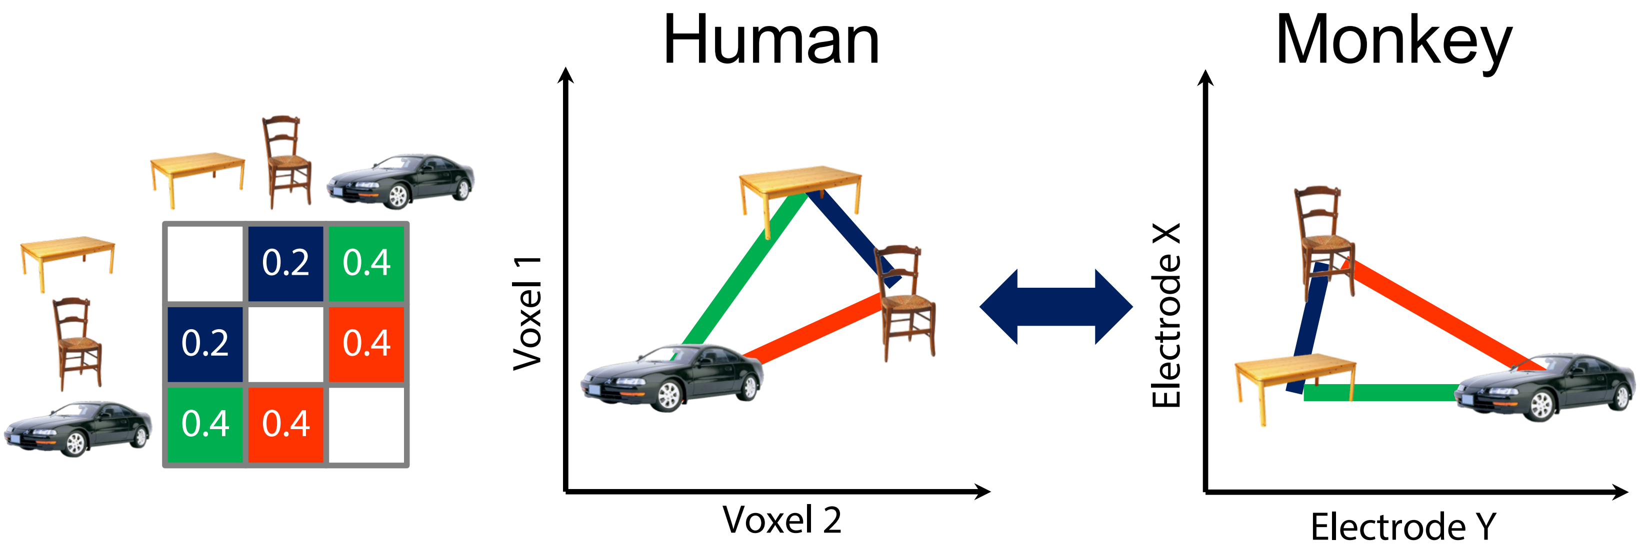

How does RSA work?

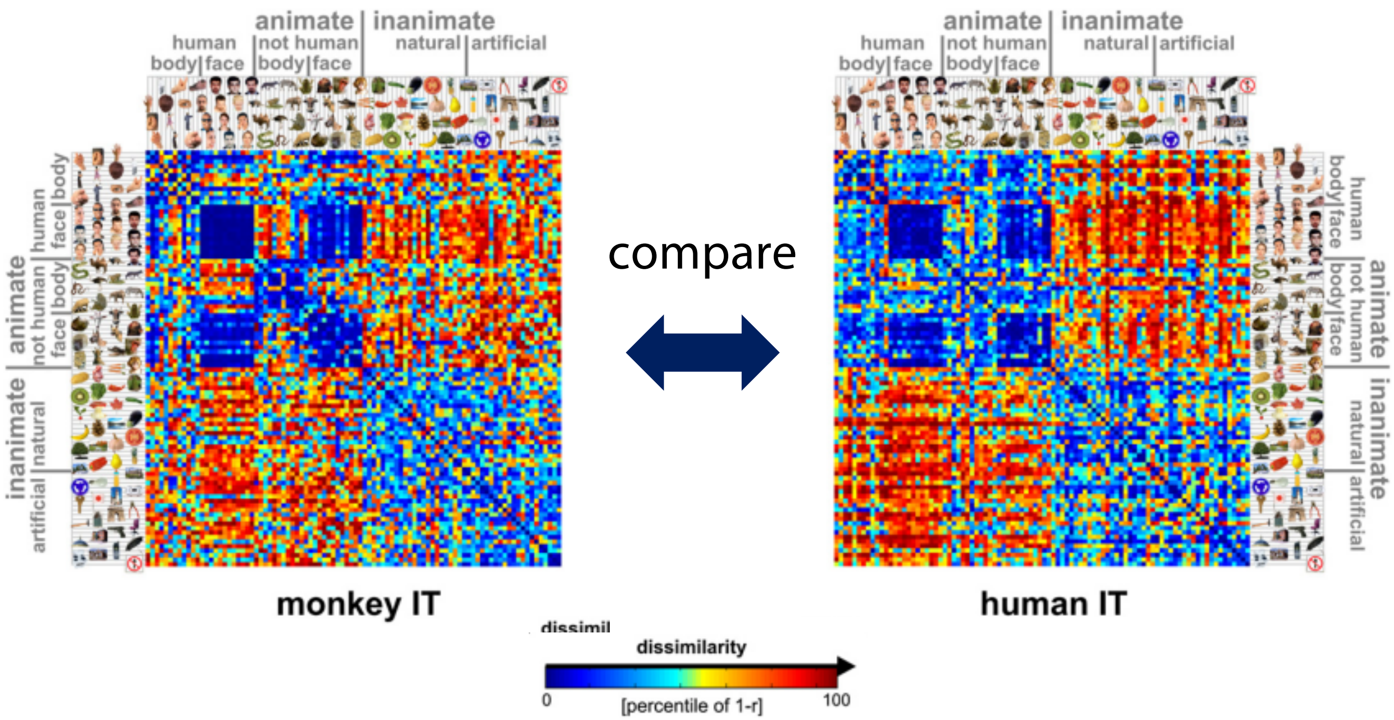

→ Representational spaces serve as common language between different methods, brain regions, individuals and species

Kriegeskorte et al (2008) – Neuron

When is RSA useful / when is it not?

True, high-dimensional representational space

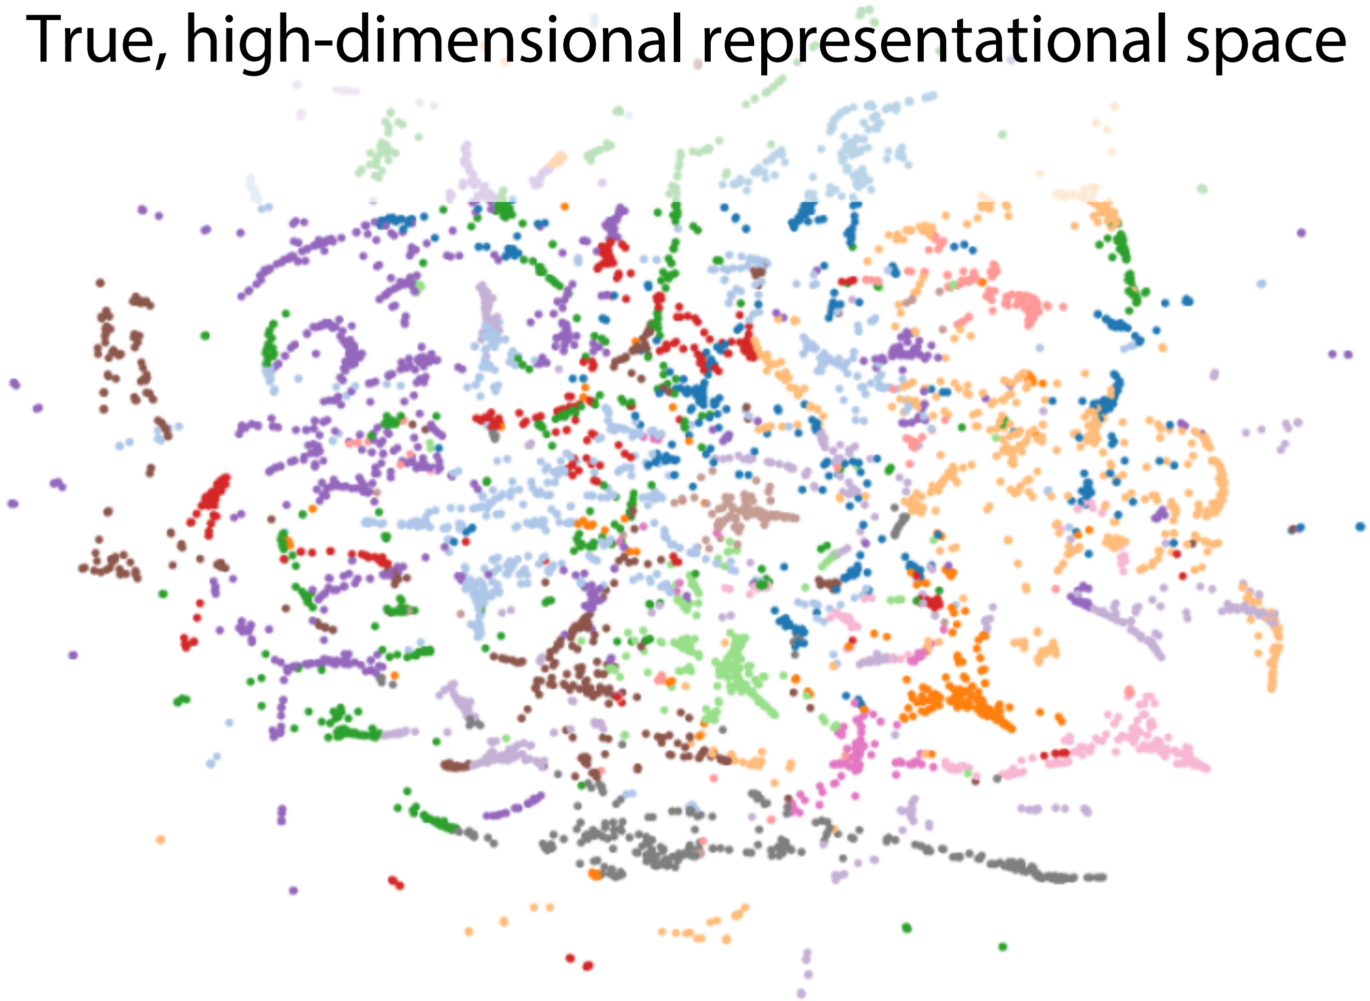

When is RSA useful / when is it not?

Problem: Unknown representations

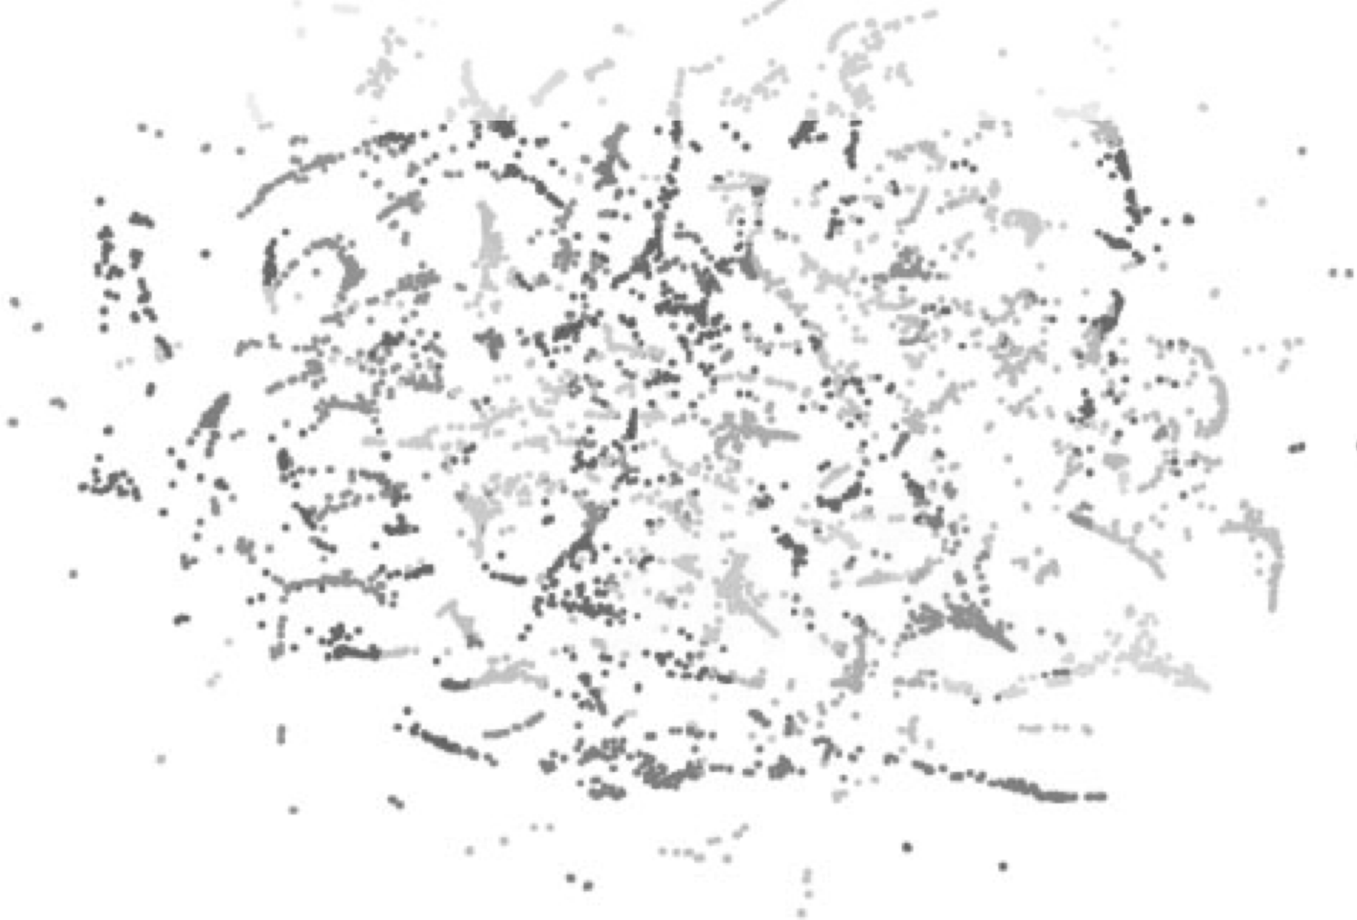

When is RSA useful / when is it not?

Classical experiment
(e.g. stimulus type A vs. stimulus type B)

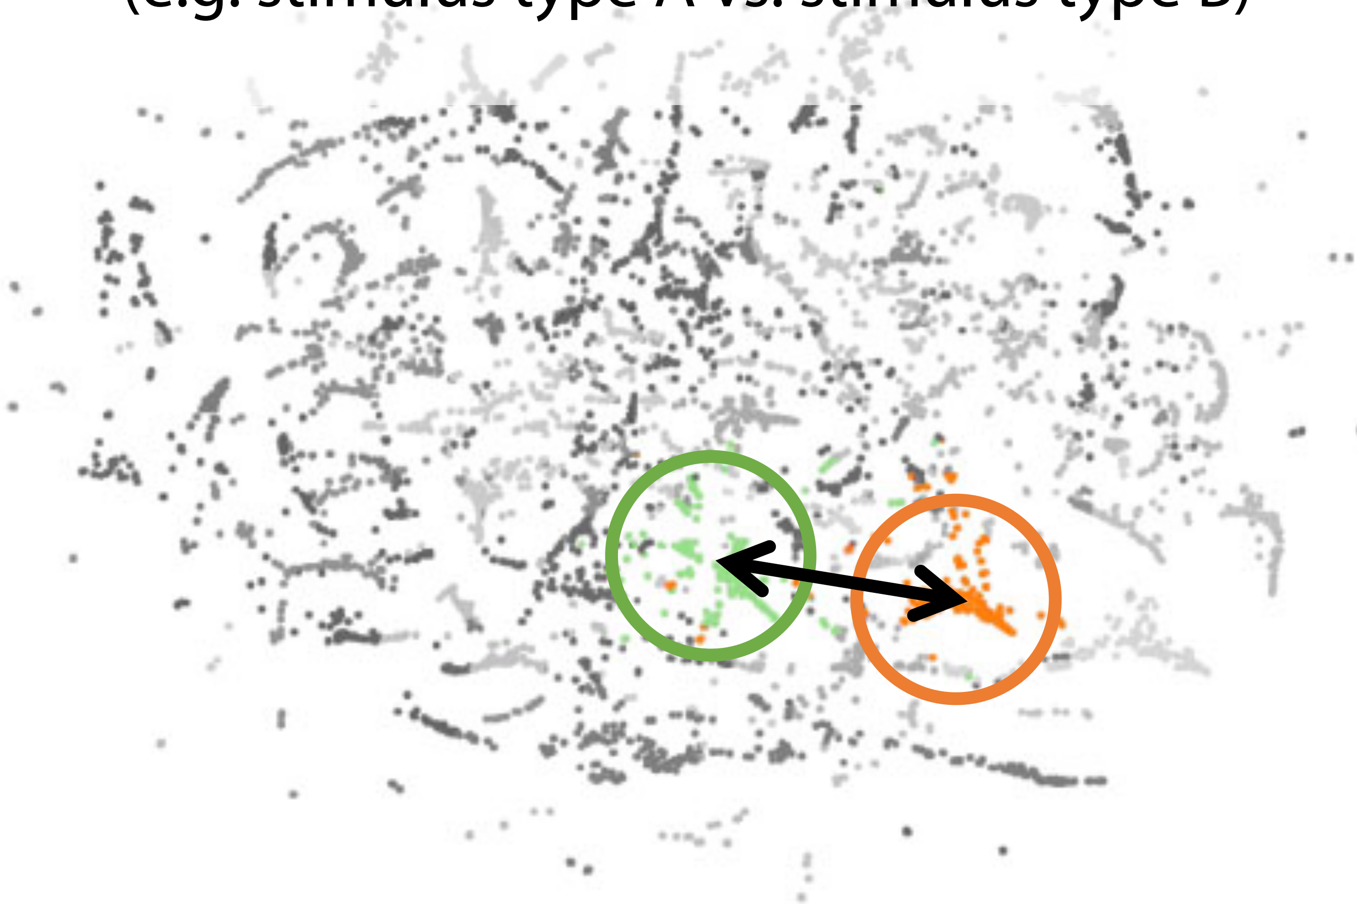

When is RSA useful / when is it not?

Condition-rich design
(= a little bit of everything)

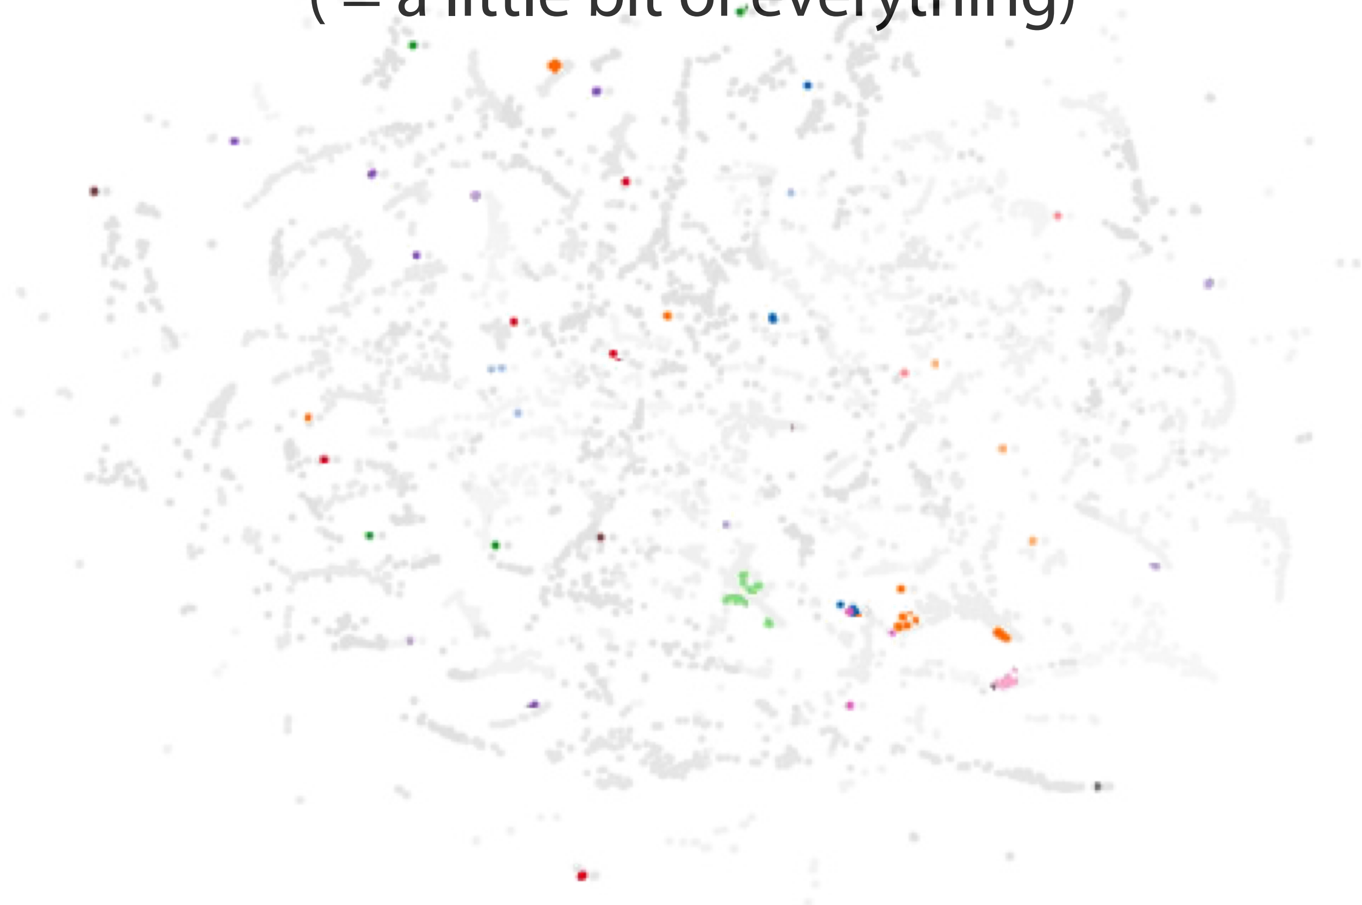

Quiz: which one is RSA? (A), (B), or both?

(A)

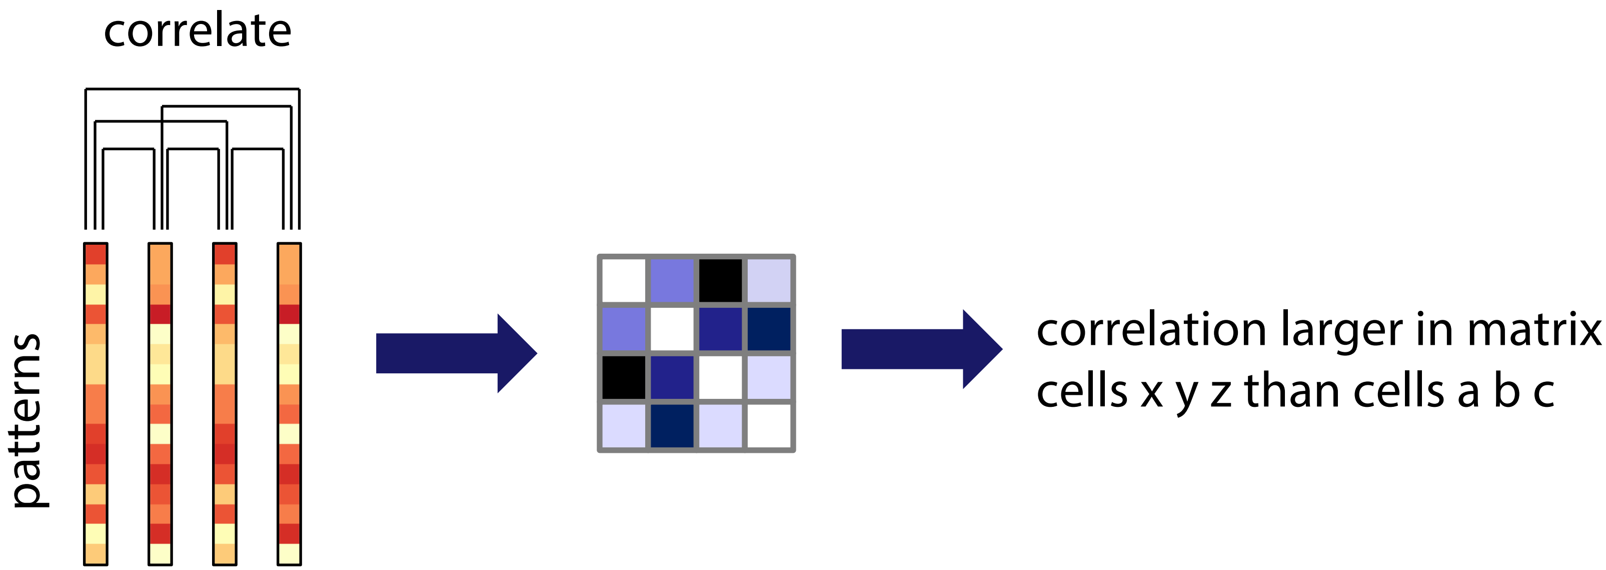

(B)

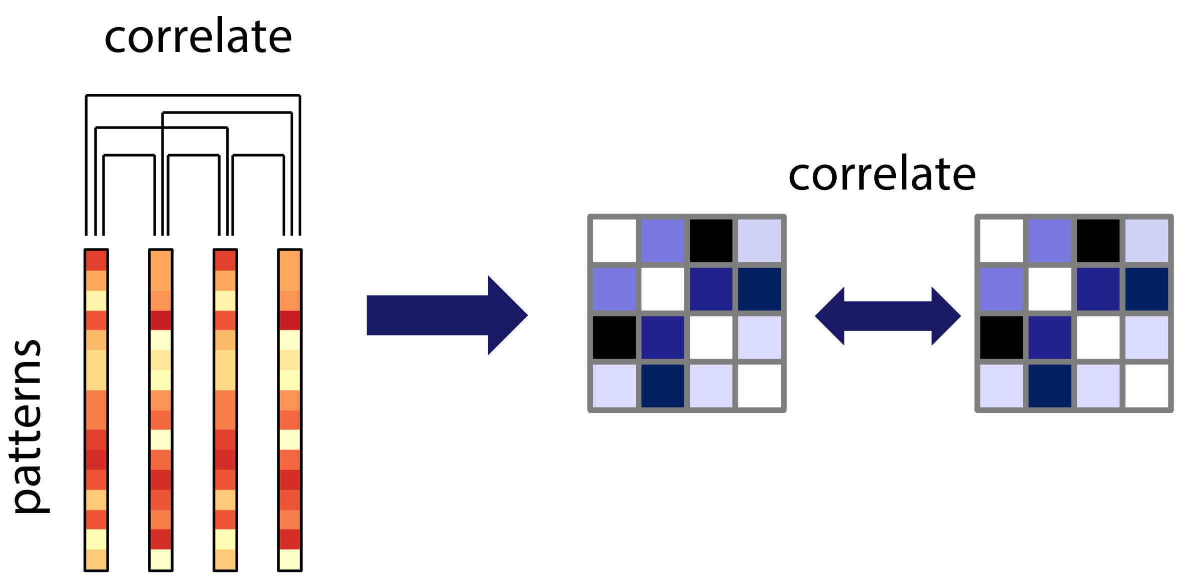

DISTANCE METRICS

Construction of RDM: Which Distance Metric to Use?

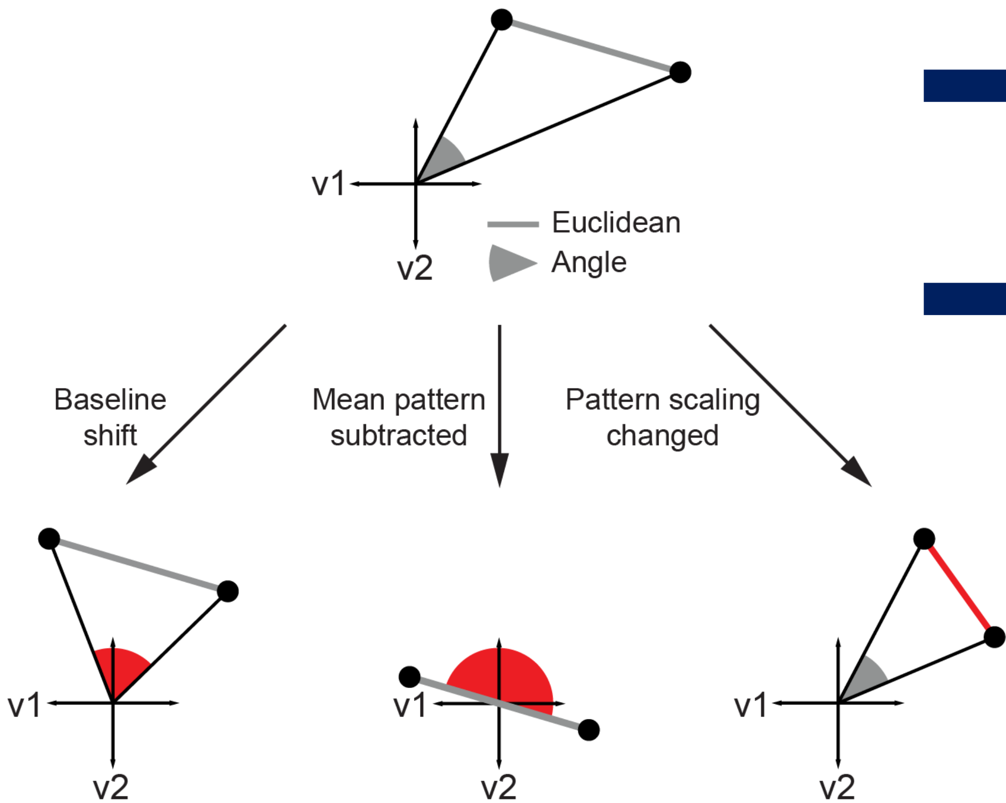

Euclidean distance is invariant to baseline shifts

Correlation distance is invariant to scaling differences

See talk by Fernando Ramirez

Garrido et al. (2013) – Front Hum Neurosci; Walther et al. (2016), Neuroimage

Distance Estimates are Biased

Bias in estimation of distances

1. Distances are always positive

2. The null hypothesis assumes true distances of zero

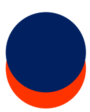

3. Measurement noise always leads to non-zero deviations between both

→ **Distances are overestimated**

→ **More generally: this also holds when there is a true distance**

→ **Solution: Cross-validation (allows negative distances)**

Distance Estimates are Biased

Do we care if distances are overestimated?

- No if we are only interested in the order (“classical” RSA)
- Yes, if we want to test against 0
- Yes if we want to interpret relative distances
  (i.e. distance A = 2x distance B)

Can we get unbiased estimates of the distance?

- Yes, using cross-validation (which allows also negative distances)

DISTANCE METRICS

Multivariate Noise Normalization

Mahalanobis distance downscales noisy voxels / voxels with high noise covariance and can improve reliability

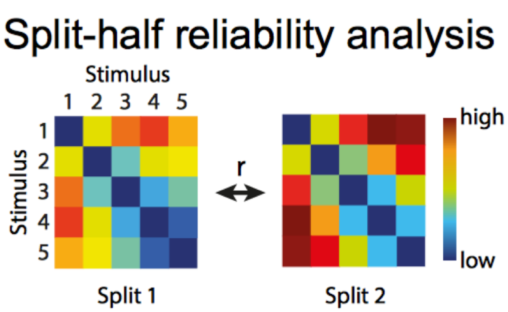

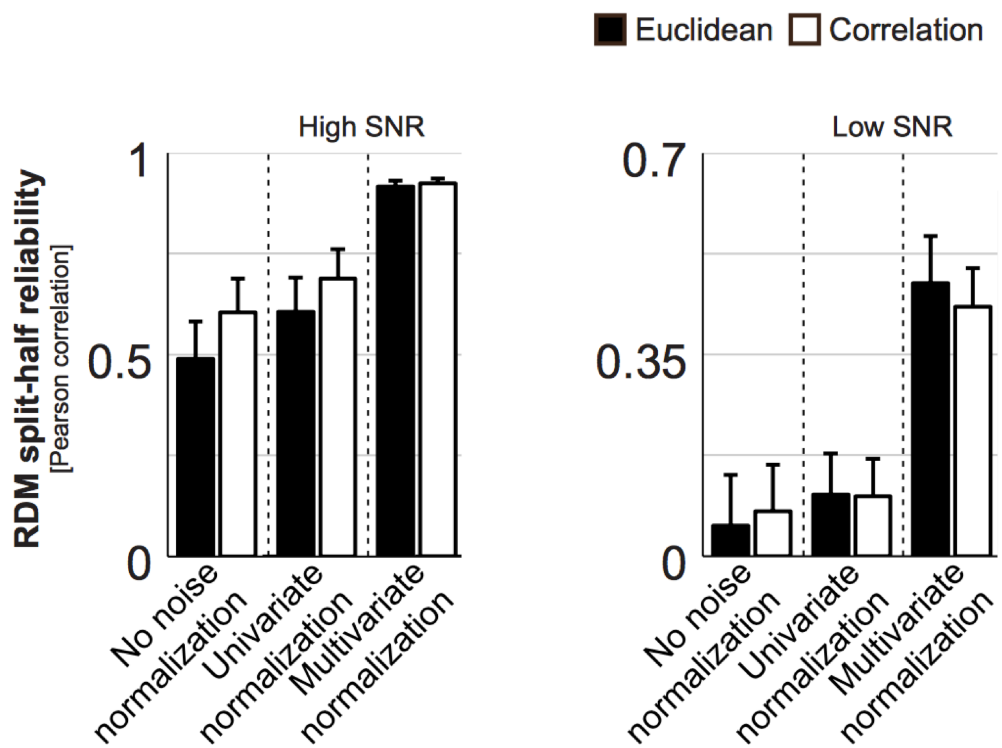

Multivariate Noise Normalization

Multivariate noise normalization recovers true space when scaled, because also noise variance is scaled

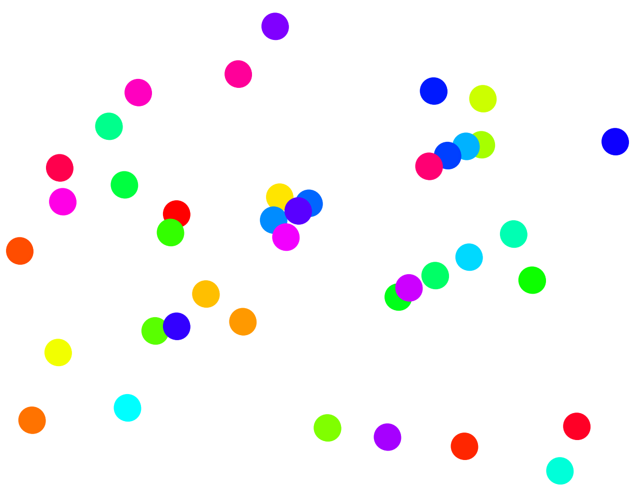

true space

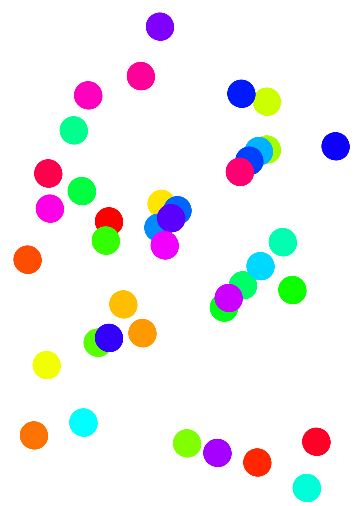

measured space (biased sampling)

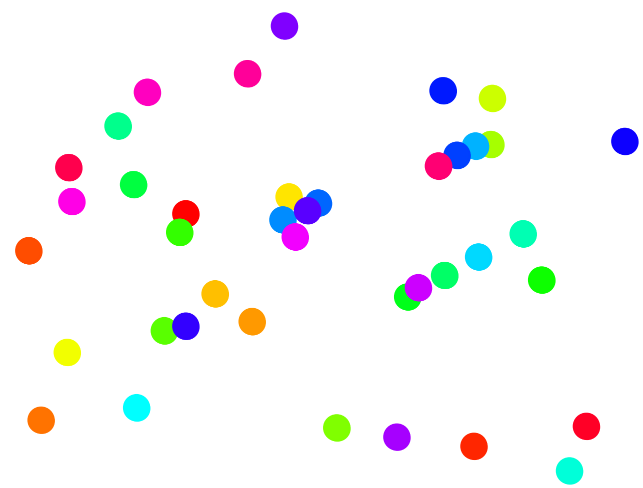

space after noise normalization

Multivariate Noise Normalization

Multivariate noise normalization recovers true space when sampling introduces covariance between voxels

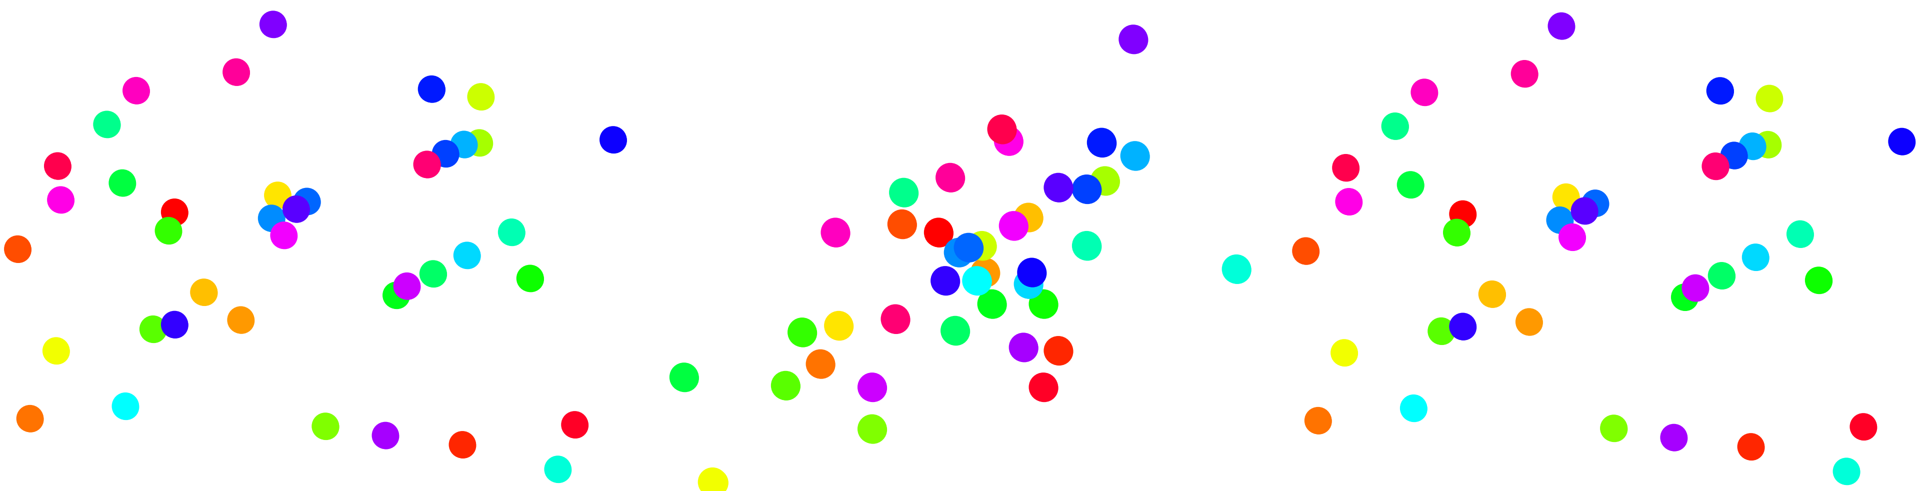

true space

measured space (correlated noise)

space after noise normalization

Multivariate Noise Normalization

Multivariate noise normalization distorts true space when representational features are really correlated

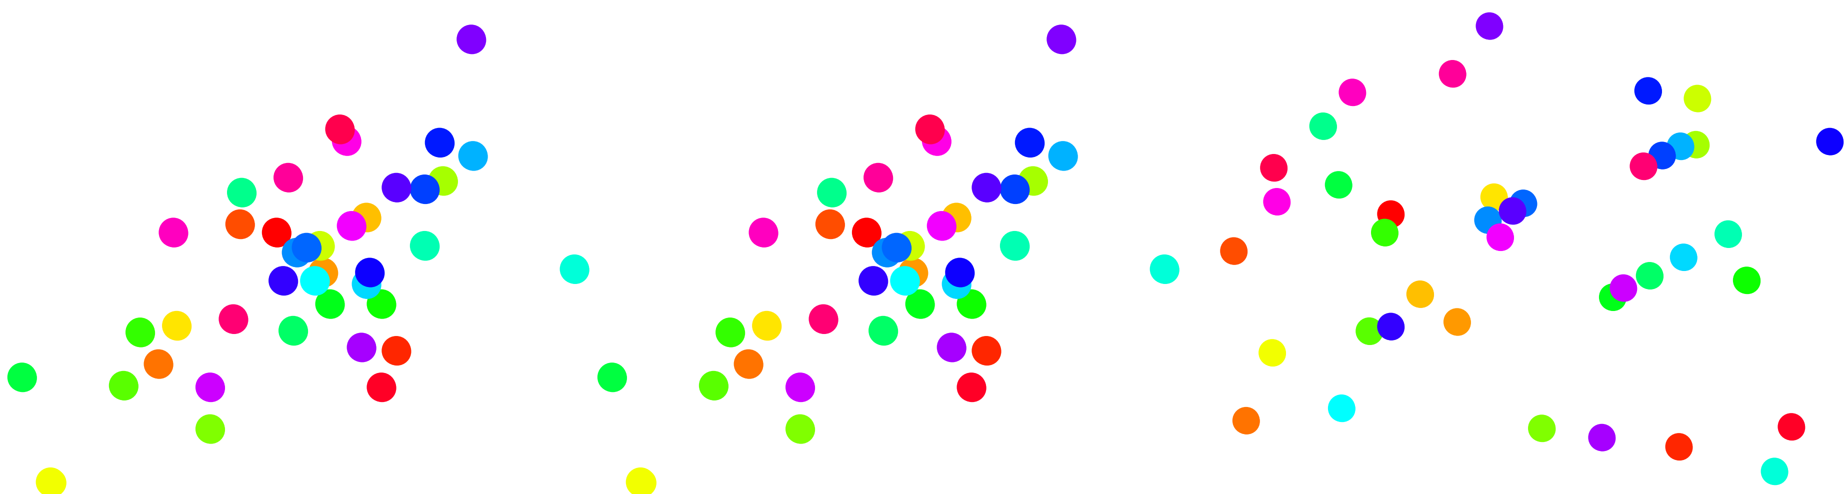

true space

measured space (veridical sampling)

space after whitening

CONSIDERATIONS FOR COMPARING REPRESENTATIONAL DISSIMILARITY MATRICES

Are Dissimilarities Ratio Scale?

Ratio scale: Dissimilarity of 4 is 2x as much as dissimilarity of 2

- Holds only when features add up linearly to create measured patterns
- Holds only when BOLD response is linear

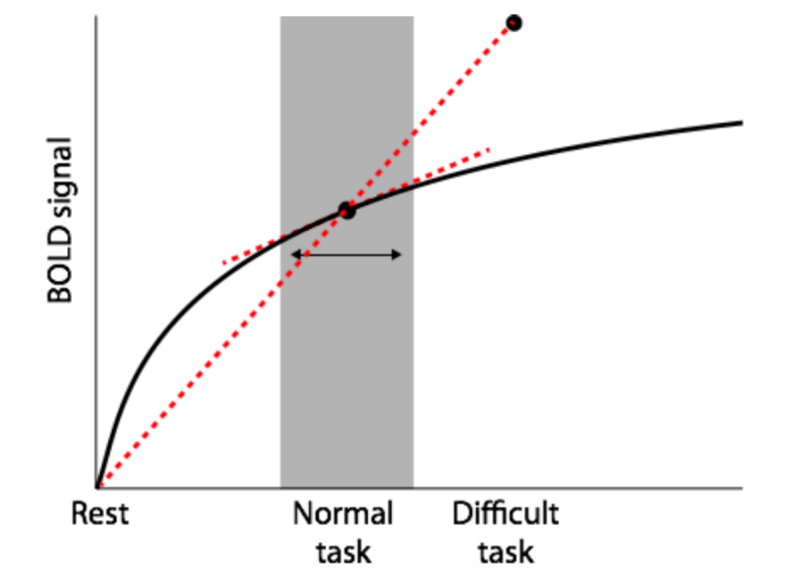

Ratio-scale justified when activation level across conditions is similar

Mean value subtraction does not fix this!

Default is to use rank order correlation coefficient when comparing RDMs (Spearman’s rho or Kendall’s tau)

Practical Considerations for RSA

When correlating symmetrical RDMs, only correlate the lower triangle, <u>always exclude diagonal</u>!

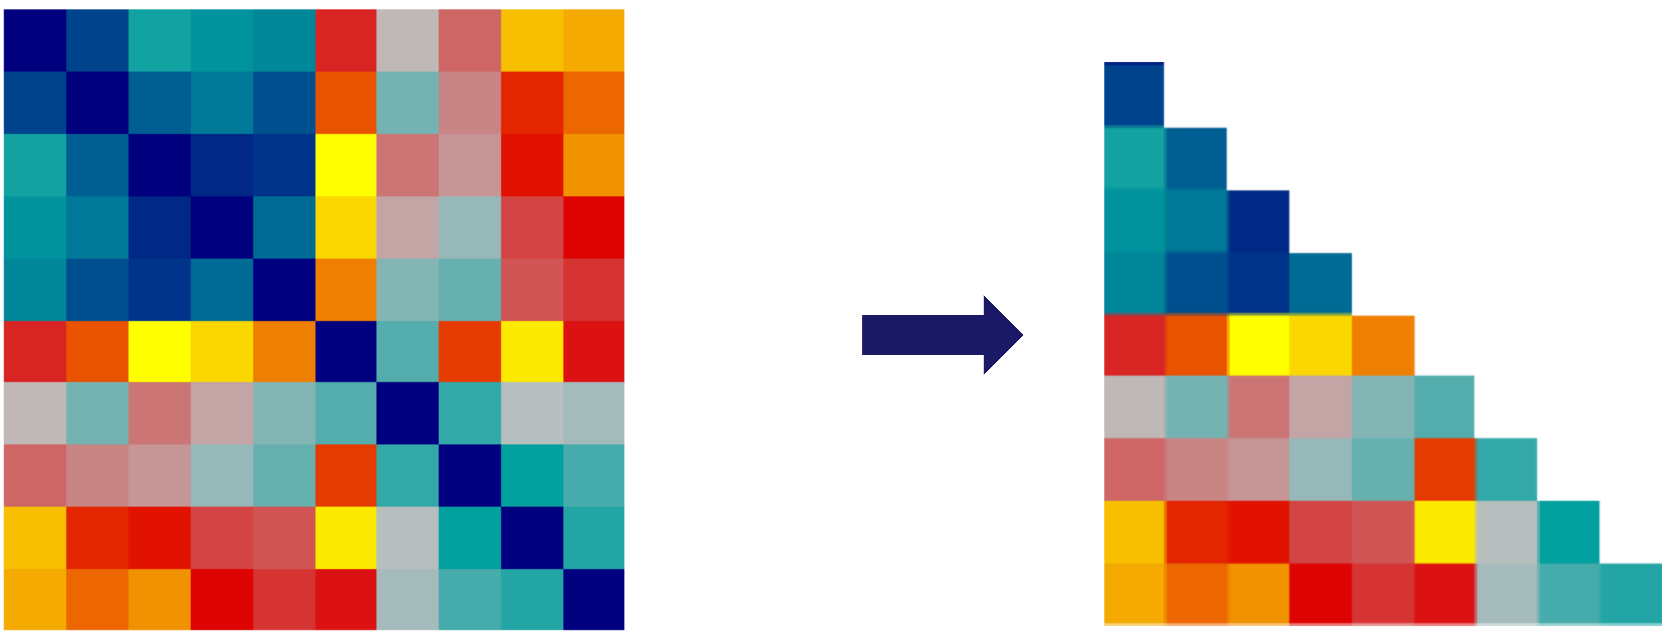

Ritchie et al (2017) – Neuroimage

Practical Considerations for RSA

When viewing RDMs , both scaled and unscaled versions might be helpful for illustrative purposes

original distances

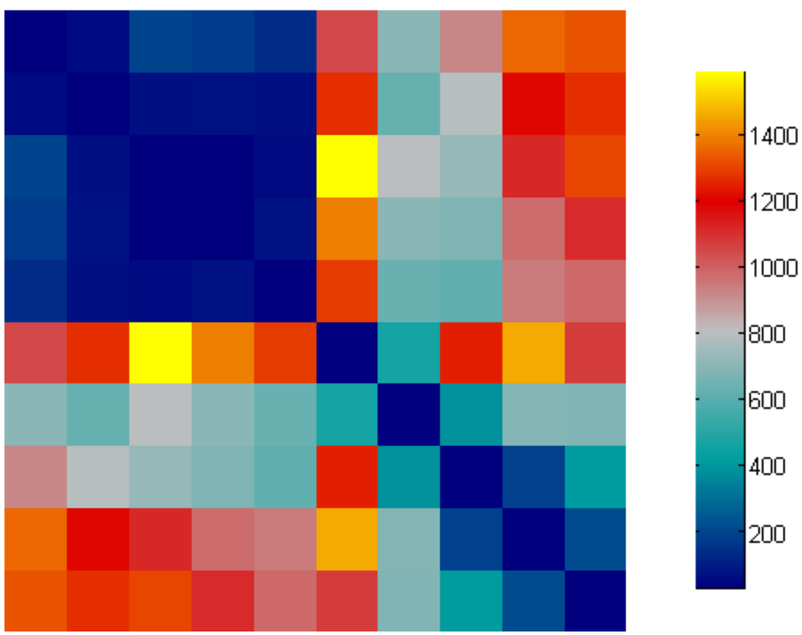

scaled ranks of distances

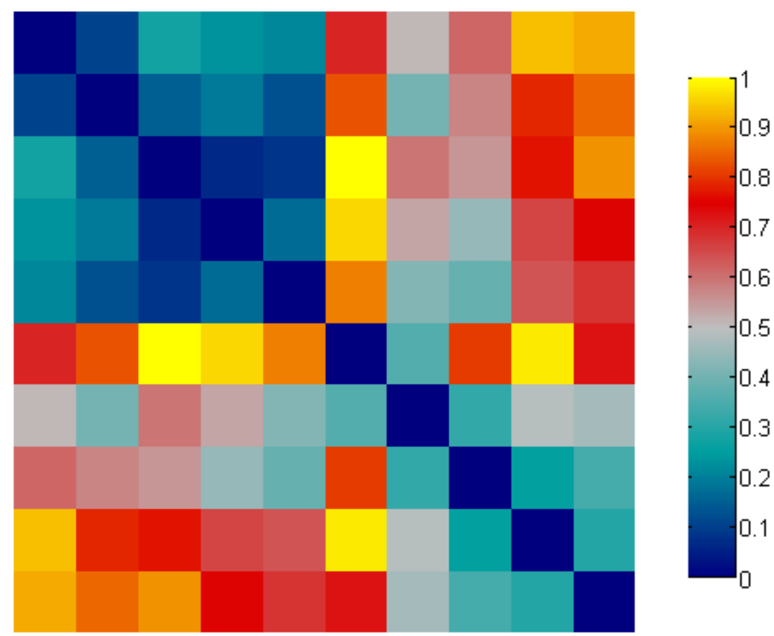

Kriegeskorte et al (2008) – Front Hum Neurosci; Nili et al (2014) – PloS Comp Biol

Noise Ceilings: Upper Limit of Model Performance

Are non-perfect model correlations related to model or are they related to noise in the data?

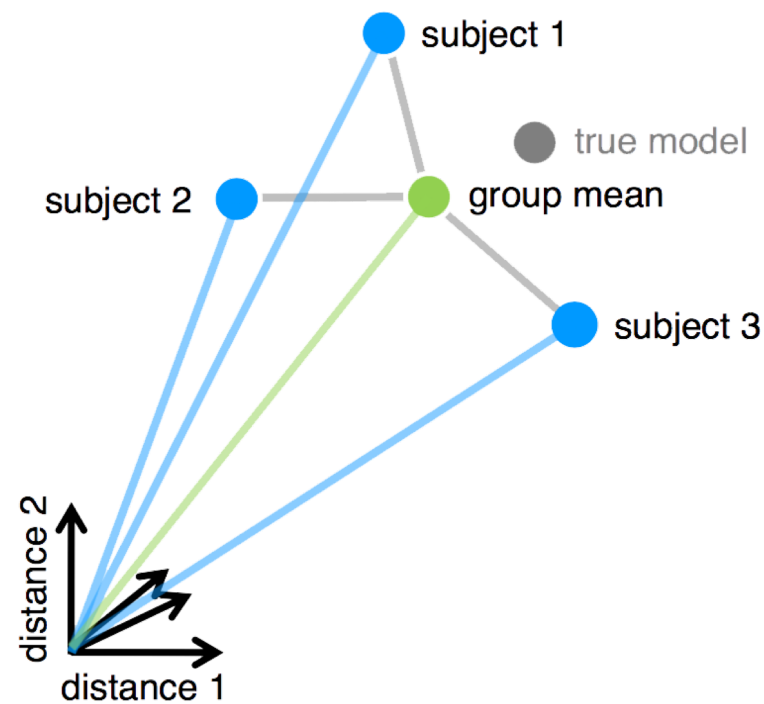

Estimate of best any model can do on the data: Group mean RDM

- Upper bound: Compare each subject to mean (positively biased because subject included in mean)
- Lower bound: Compare each subject to mean excluding that subject (negatively biased because not all data used)

Kriegeskorte et al (2008) – Front Hum Neurosci; Kriegeskorte – 2015 RSA workshop

Summary

- RSA measures the representational content and format of representations
- Representational dissimilarity matrices can be compared between brain regions, individuals, species and measurement modalities
- RSA can be used to test (computational) models of cognition
- Noise normaization can improve the reliability of RDMs
- Noise ceilings assess whether the model sufficiently accounts for the data or whether more data is needed

Study Questions

- How does RSA relate to multivariate decoding?

- Imagine you carry out RSA between two regions A and B and you get a correlation of those two RDMs. Now repeat this analysis for all comparisons of 10 different regions. Can you come up with ideas what you can do with these second-order comparisons?