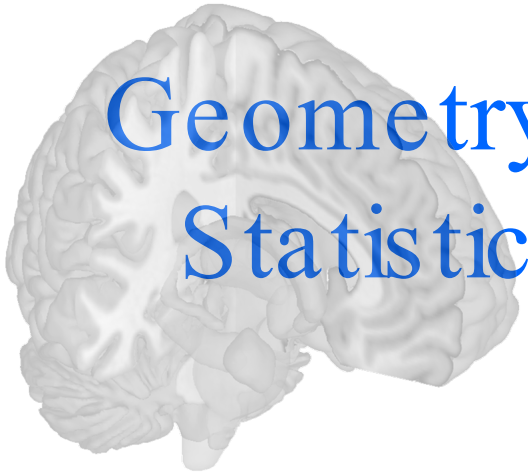
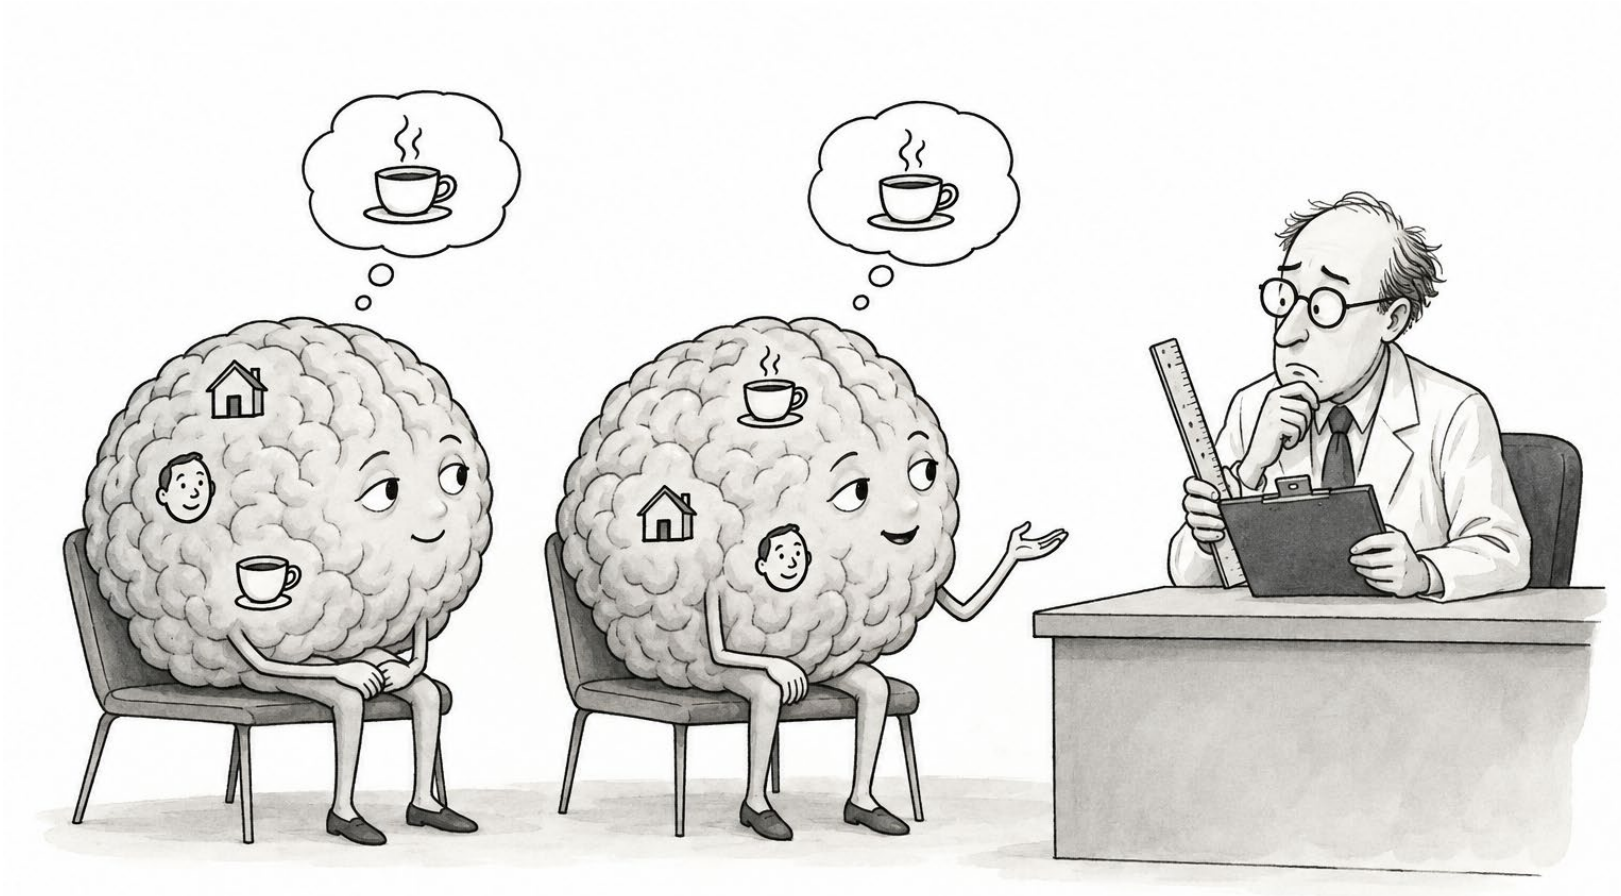


Geometry, Topography, and the Statistics of Functional Brain Alignment

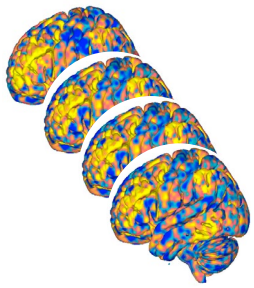
Martin Lindquist
Department of Biostatistics
Johns Hopkins University

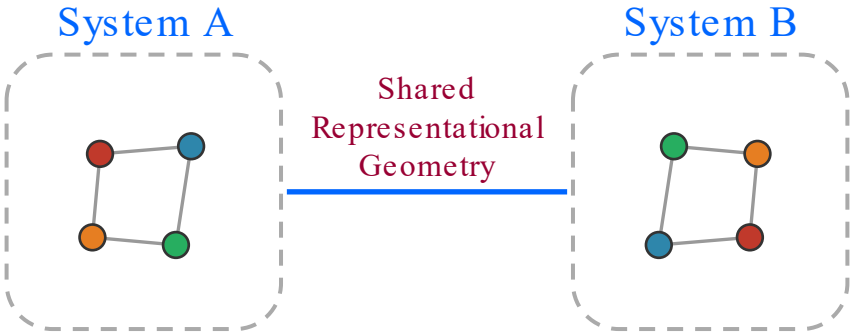


"We agree on everything. Must we agree on where?"

When Do Representations Match?

- System: the computational object being compared; for example, a brain or an artificial neural network.
- Representation: the pattern, or set of patterns, of internal activity through which a system encodes information about stimuli, task conditions, states, or variables.
- Question: What does it mean to say that two systems have the "same representation"?
  - Information is expressed in corresponding spatial locations.
  - Relationships among activity patterns are preserved.

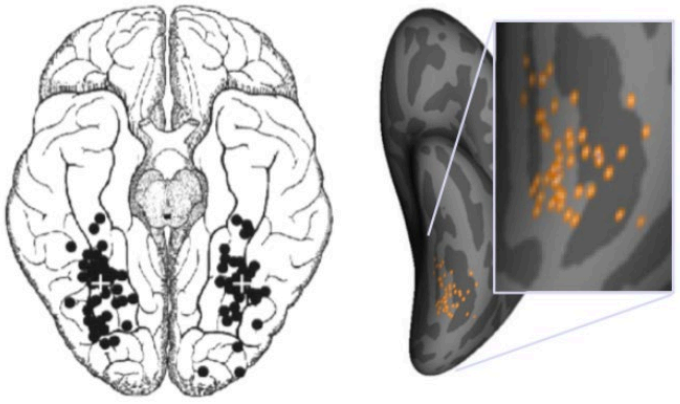




Different spatial arrangement → preserved pairwise relationships

Anatomical ≠ Functional Alignment

- Modern functional magnetic resonance imaging (fMRI) analyses increasingly rely on fine-grained spatial information.
- Anatomical normalization puts subjects into a common coordinate system.
- But corresponding anatomical locations do not necessarily correspond functionally.



Left: Variability in face-specific regions assessed with intracranial EEG in 93 patients. Right: inter-individual variability in peak word-selective areas in fMRI

Allison et al. (1999)

Misalignment Breaks Feature Correspondence



- Misaligned functional maps can dramatically lower the effective spatial resolution and cause problems for group-level inference.

  1. Brain mapping uses a mass-univariate approach, where the same voxel is compared across subjects.

  2. Connectivity investigates relationships between regions whose boundaries may be affected by misalignment.

  3. Predictive models often use voxels as features, and any variation in the location of functional regions across subjects would lead to feature misalignment.

- Additional functional alignment is needed, but what structure should it preserve?

Two Hypotheses About What Is Shared



Topographic alignment

What is shared?

- Spatial organization

What may vary?

- Location and smooth spatial deformation

→ Functional registration

Geometric alignment

What is shared?

- Representational geometry

What may vary?

- Spatial implementation

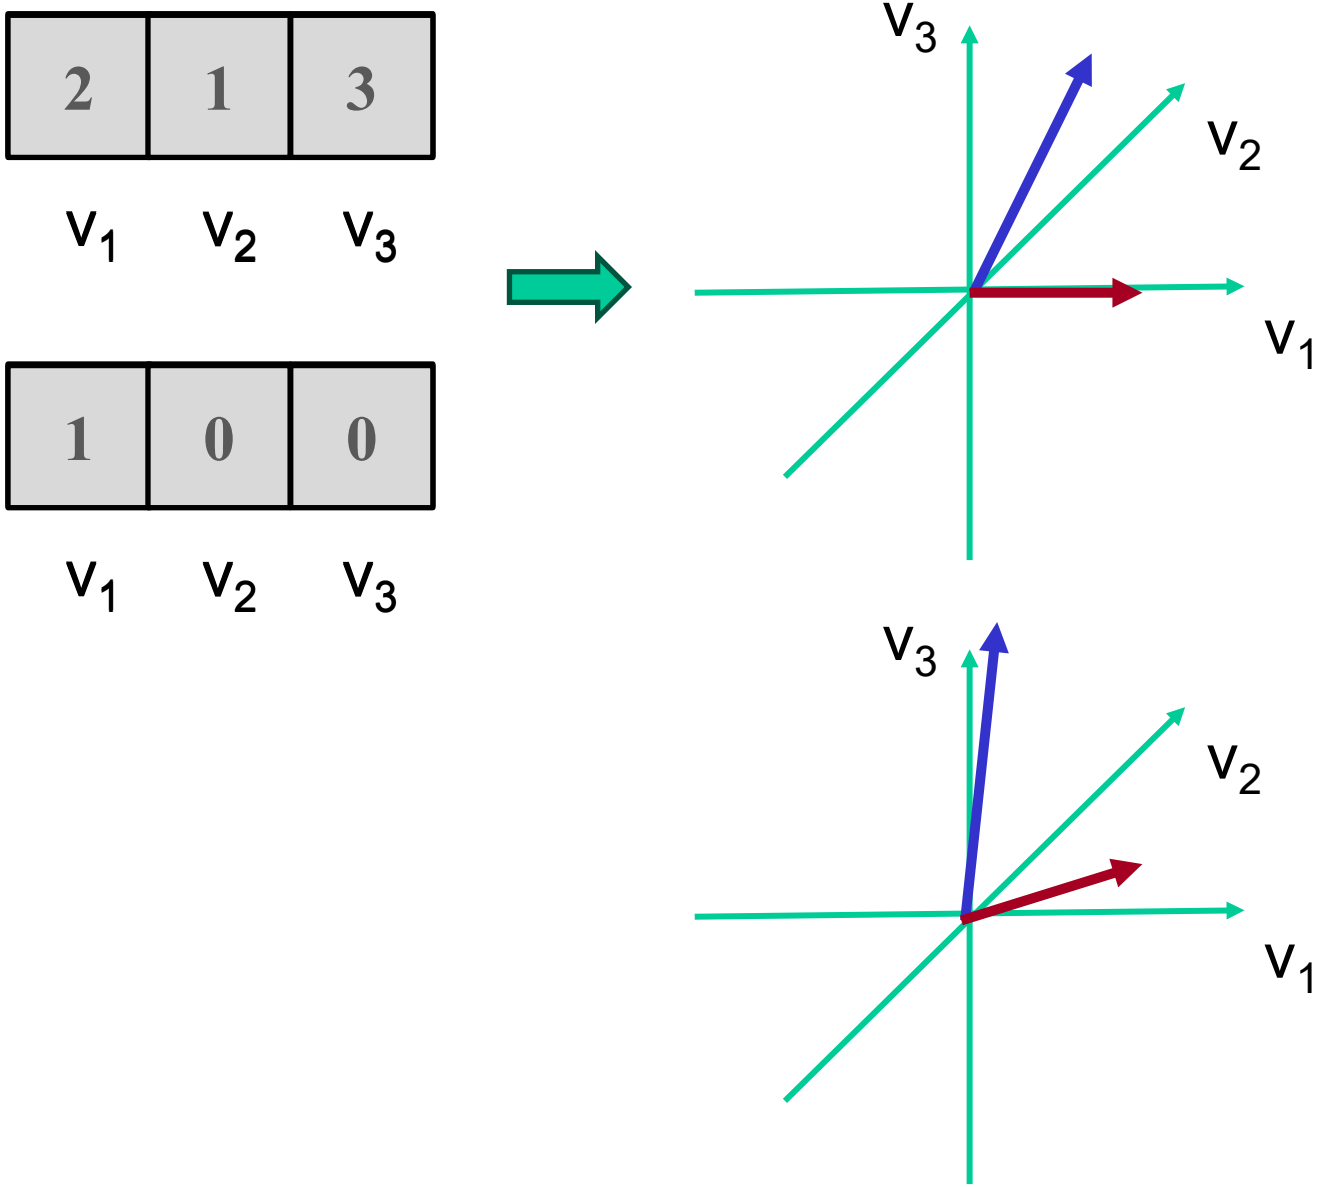
→ Hyperalignment

These approaches make different assumptions about what is shared across individuals: spatial organization or representational geometry.

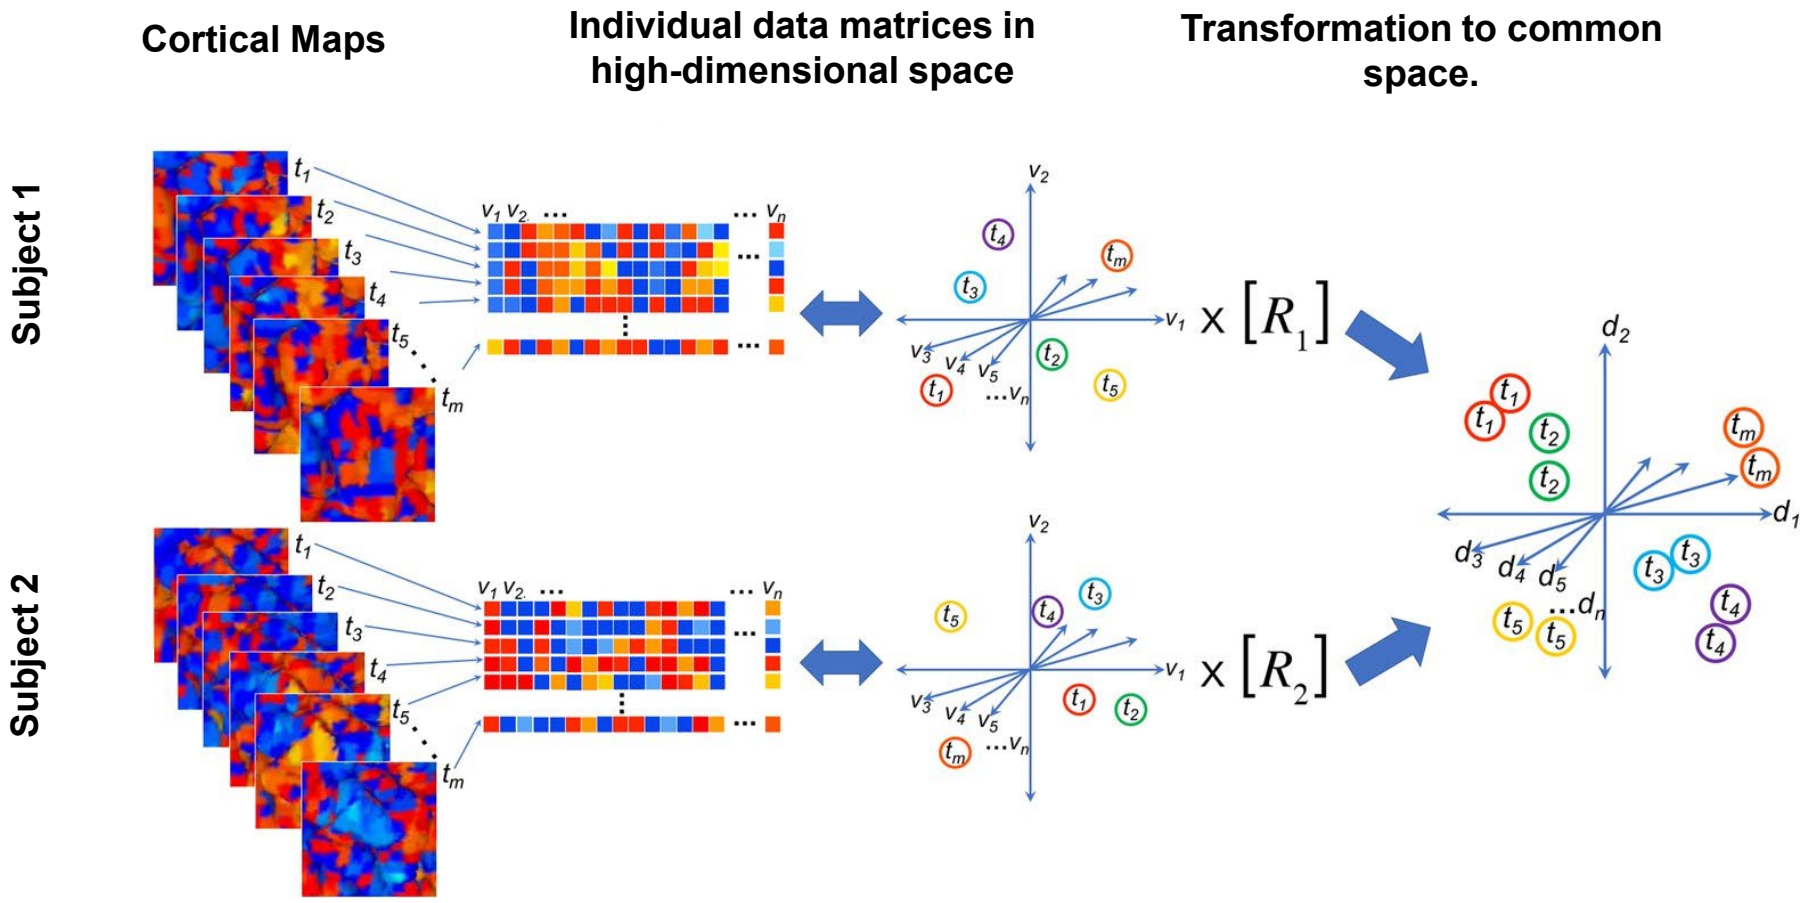
Hyperalignment

- Haxby et al. (2011) proposed hyperalignment (HA) as a method to align a set of functional maps.
  - Aligns functional brain representations across subjects based on shared representational geometry.
  - Information embedded in V-dimensional space, where V is the number of voxels being considered.
  - Vector representations of activation patterns are transformed into a common high-dimensional space shared across participants.
  - The transformation does not explicitly preserve spatial topography.

Hyperalignment



Hyperalignment



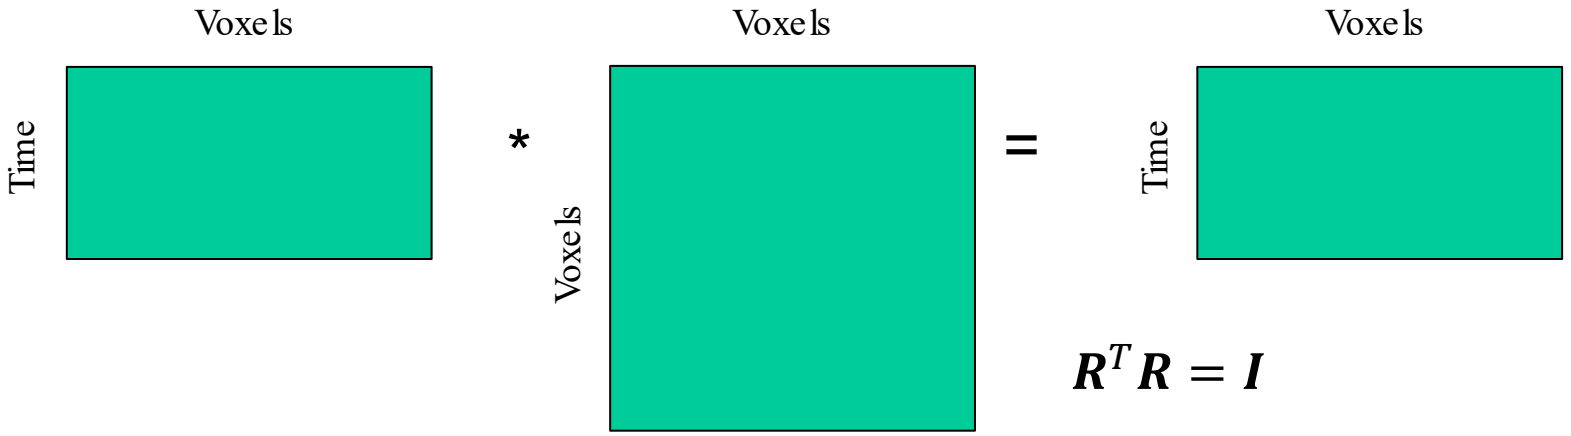
Haxby et al. 2020

Hyperalignment

$\mathbf{X}_i$ – individual data in individual brain space

$\mathbf{R}_i$ – individual transformation matrix

$\mathbf{M}_i$ – individual data in common model space.



$$\boldsymbol{R^T R = I}$$

$$\min_R \left\| \boldsymbol{X_i R - X_j} \right\|_{\boldsymbol{F}}^{\boldsymbol{2}} \longrightarrow \boldsymbol{R = UV^T} \qquad \boldsymbol{X_i^T X_j = U \Sigma V^T}$$

$$\min_{R_i} \sum_{\boldsymbol{i=1}}^{n} \left\| \boldsymbol{X_i R_i - M} \right\|_{\boldsymbol{F}}^{\boldsymbol{2}} \quad \text{where} \quad \boldsymbol{M} = \frac{1}{n} \sum_{\boldsymbol{i=1}}^{n} \boldsymbol{X_i R_i}$$

Hyperalignment

- The orthogonal Procrustes transformation can be used to align two sets of activation profiles.
  - Finds the rotation or reflection that best aligns two sets of activation patterns while preserving within-subject Euclidean distances.
    - Preserves the vector geometry of each data set

- The generalized Procrustes transformation can be used to align multiple activation profiles.
  - In practice this is performed using sequential application of the Procrustes transformation.

Problems?

- Several properties of the proposed implementation of HA should be considered:

  1. It remixes data across spatial loci.
     - Ideally, want to avoid remixing information from distant voxels.

  2. Solutions calculated using HA are not unique.
     - Sequential Procrustes depends directly on the order in which individuals are entered into the algorithm.

  3. Computational issues when working with the dimensions used in whole-brain fMRI studies.

- The first and third points can be partially mitigated by using a "searchlight" approach.

ProMises

- The ProMises (Procrustes von Mises–Fisher) model addresses some of these shortcomings.
  - Expresses functional alignment as a statistical model

$$\boldsymbol{X}_i = \boldsymbol{M}\boldsymbol{R}_i^T + \boldsymbol{E}_i \quad \text{each row of} \quad \boldsymbol{E}_i \sim \text{MVN}(\boldsymbol{0}, \boldsymbol{\Sigma}_v)$$

  - Prior distribution imposed on $\mathbf{R}_i$ (matrix von Mises–Fisher distribution), allowing for the embedding of anatomical information into the estimation

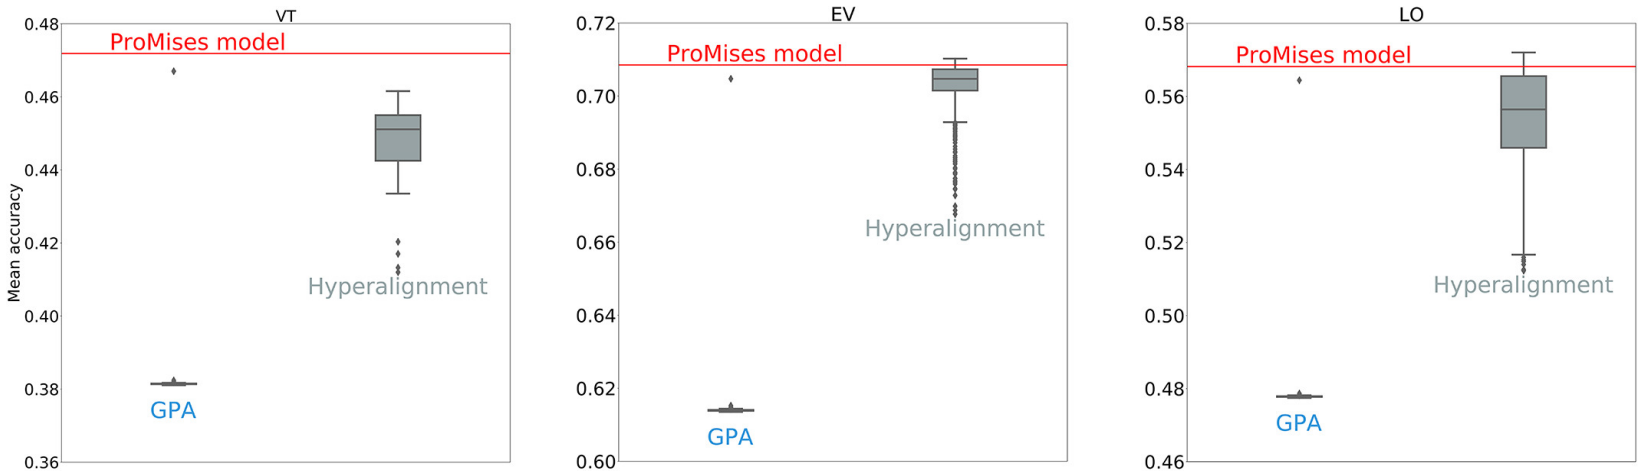
$$f(\boldsymbol{R}_i) = C(\boldsymbol{F}, k)exp\{tr(k\boldsymbol{F}^T\boldsymbol{R}_i)\} \qquad \boldsymbol{F} = exp\{-d_{uv}\}$$

  - The derived transformations, aligned images, and related results are all unique.

Andreella et al (2023)

ProMises

- Movie watching task with n = 31 subjects.
  - Used a one-nearest-neighbor classifier to identify movie segments from response vectors spanning six time points (18 seconds).
  - Classification was correct when the held-out subject's response correlates most strongly with the remaining subjects' mean response for the matching segment.
  - Repeated for each held-out subject and mean accuracy reported.



	ROIs		
	VT	EV	LO
No functional alignment	0.289	0.534	0.238
ProMises model	0.472	0.709	0.568

Andreella et al (2023)

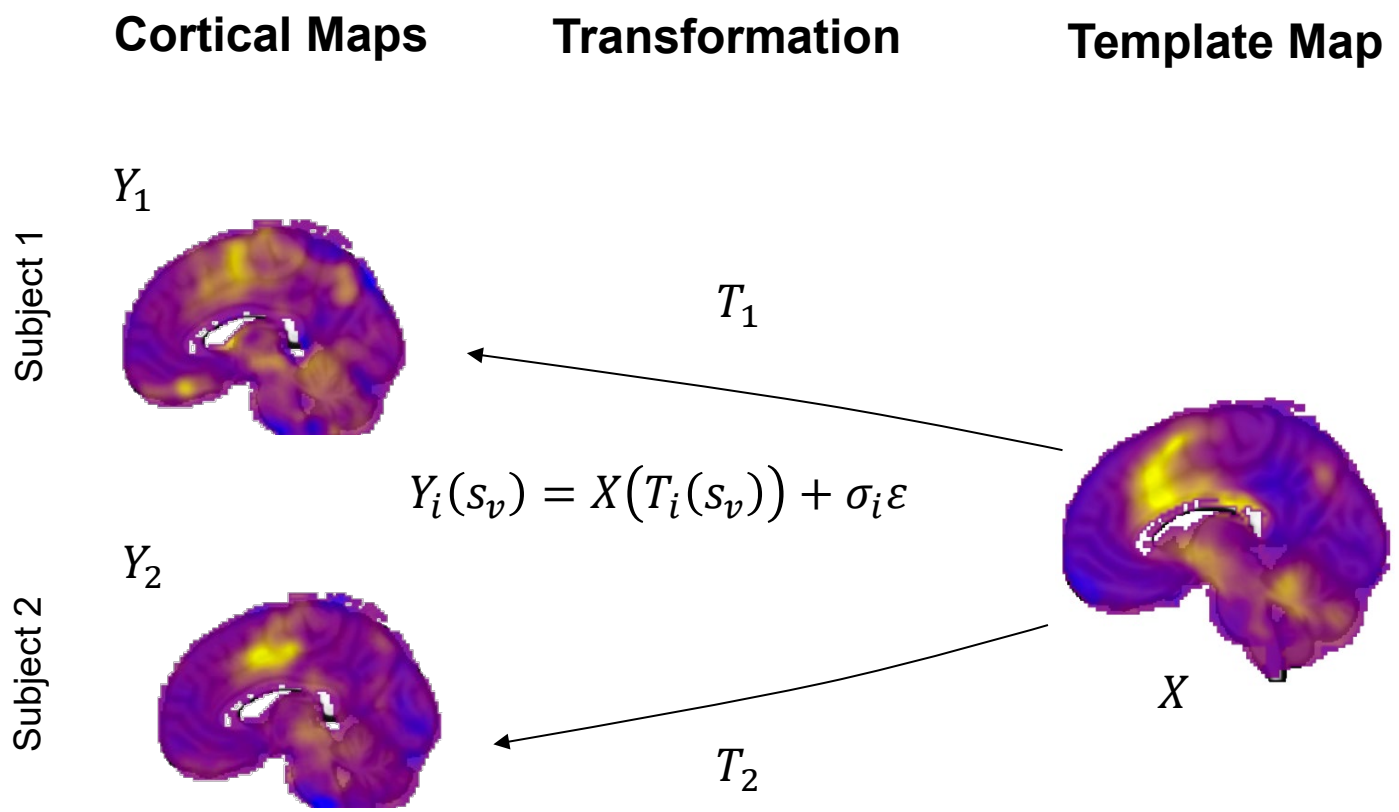
Topographic Alignment

- Topographic alignment models inter-individual differences using smooth, invertible spatial transformations that preserve neighborhood structure.

- Like HA, these approaches align individual functional maps to a common set of functional population-level reference spaces.

- In contrast, single-participant activation images are aligned based on topographical properties rather than representational reference space.

Topographic Alignment



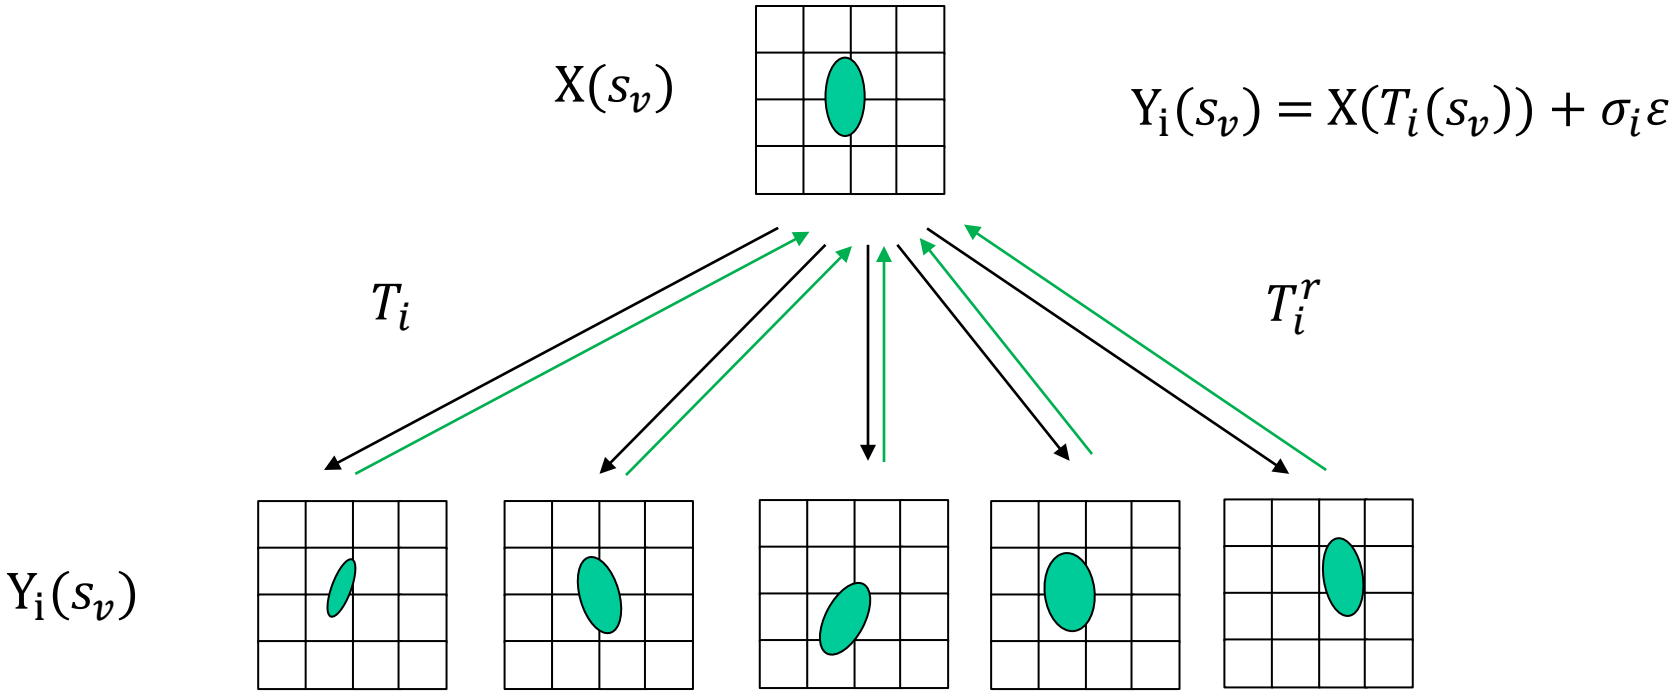
Functional Registration

- Wang et al. (2023) proposed a group-wise registration approach for spatially transforming each subject's functional data to a common latent template map.

- Used a generalized Bayes framework with a loss function for symmetric group-wise registration.

- The latent template is modeled using a Gaussian process, which helps capture spatial features in the template, producing more precise estimation.

Illustration - Functional Registration



$Y_i(s_v)$ is the intensity value in the spatial map of subject i at location s_v corresponding to voxel v, for $i = 1, \ldots, N$; $v = 1, \ldots, V$.

$X(s_v)$ is the latent activation map to which we seek to register the maps.

Model

Model:

$$Y_i(s_v) = X(T_i(s_v)) + \sigma_i \epsilon,$$

Loss function:

$$\ell_{sym}(X, \mathbf{T}, \mathbf{T}^r, \mathbf{b}, \sigma) = \sum_{i=1}^{N} \Big\{ \frac{1}{\sigma_i^2} \|Y_i(T_i^r) - \beta_i X\|_{L^2}^2 + \frac{1}{\sigma_i^2} \|Y_i - \beta_i X(T_i)\|_{L^2}^2 + \lambda^r \|T_i(T_i^r(\cdot)) - Id(\cdot)\|_F + \lambda^r \|T_i^r(T_i(\cdot)) - Id(\cdot)\|_F \Big\},$$

Generalized Bayes (Gibbs Posterior):

$$p(X, \mathbf{T}, \mathbf{T}^r, \beta, \sigma | \mathbf{Y}) \propto \exp\big\{ -\ell_{sym}(X, \mathbf{T}, \mathbf{T}^r, \beta, \sigma) \big\} p(X, \mathbf{T}, \mathbf{T}^r, \beta, \sigma),$$

Generalized Bayes



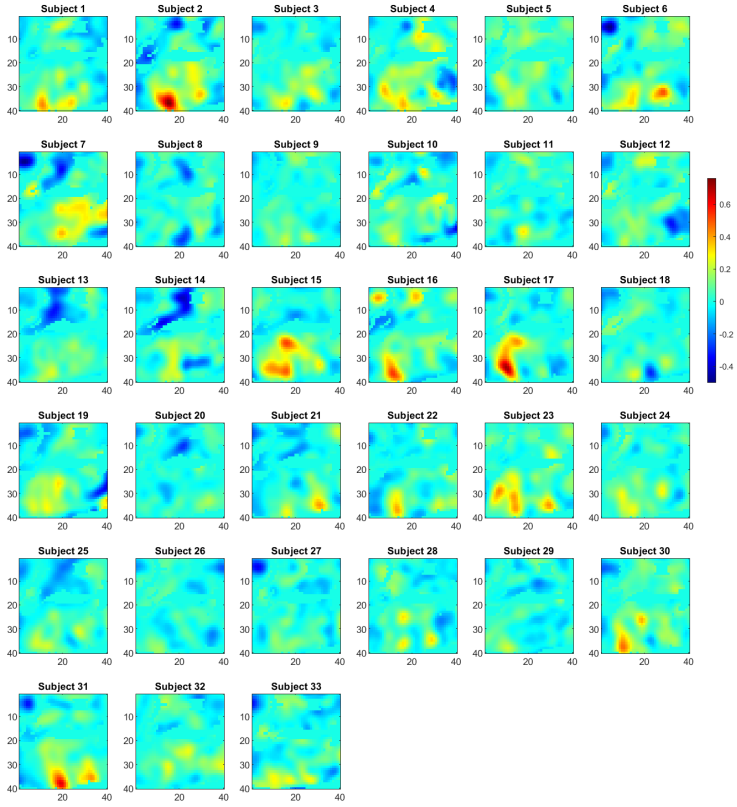
- A Gibbs posterior provides a coherent update of prior beliefs given the data and the choice of loss function.

- It updates a prior using a specified loss, which need not arise from a sampling likelihood.

- Chernozhukov and Hong (2003) established desirable large sample properties of such generalized Bayes updates under suitable regularity conditions.
  - See Martin and Syring (2022) for overview.

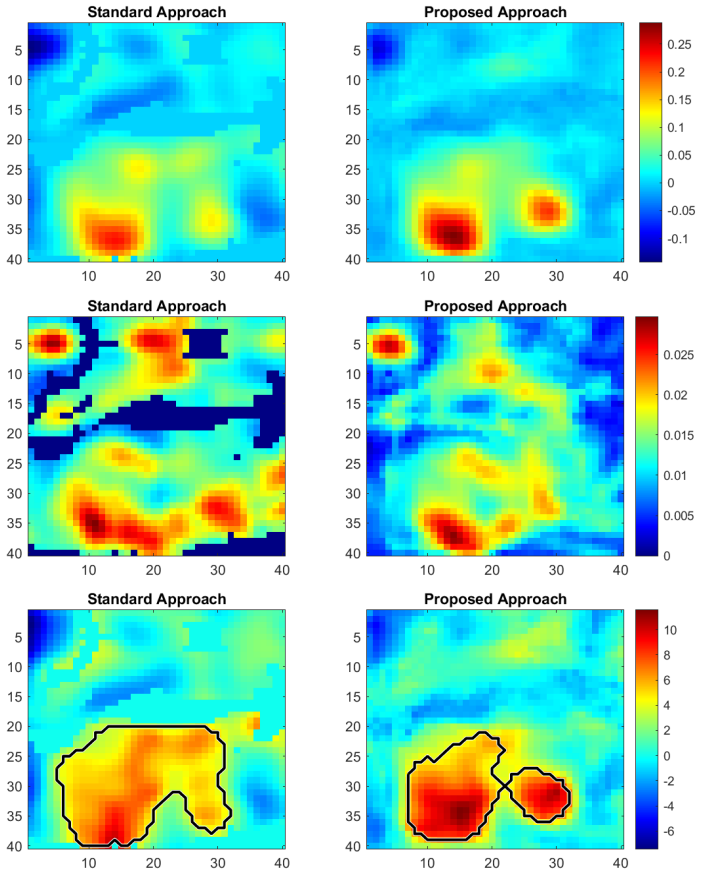
Thermal Pain Study

- Data are from an fMRI study of thermal pain
- 33 healthy subjects
- Stimuli administered at six different temperatures
- After each stimulus subjects rated their intensity of pain on a scale from 0 to 100.
- Performed functional registration on a somatosensory region of the brain.
- Derived activation maps.
- Performed group-level voxel-wise t-tests
- Fit a predictive model (Lasso-PCR) using activation maps to predict reported pain

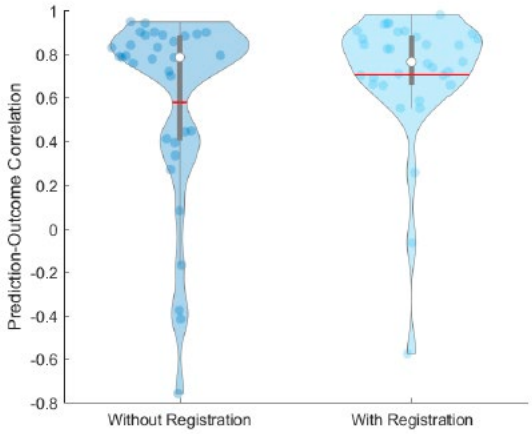


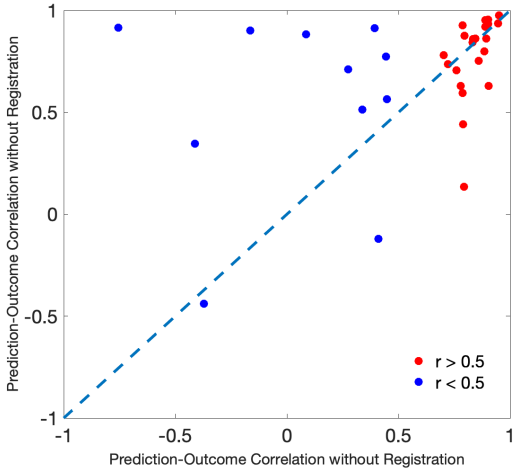
Results

Comparison of group-level inference



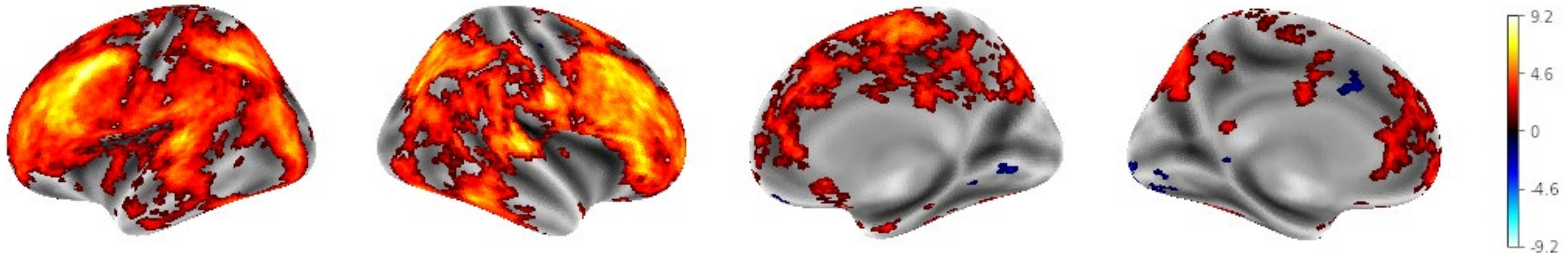
Comparison of predicted pain



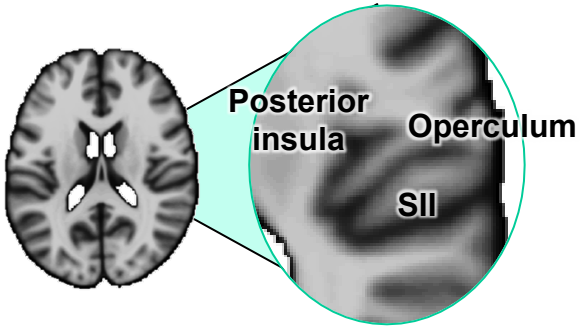


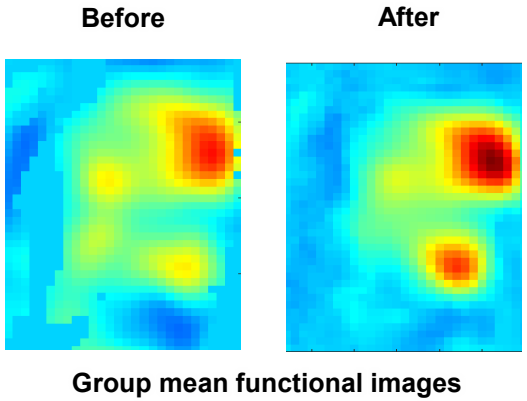
Regional Benefits of Functional Alignment

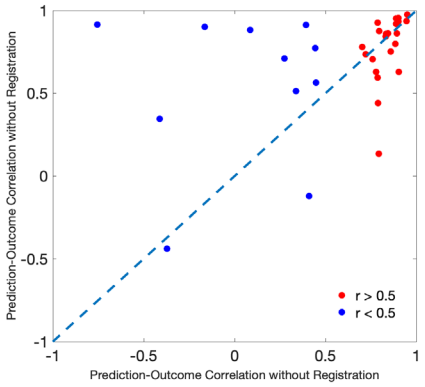
A. Improvement in pain classification after geometric alignment



B. Topography in somatosensory cortex before and after topographic alignment







What Is Shared Across Individuals?

- Could shared spatial organization and representational geometry help explain regional differences in alignment benefits?

- Petre et al. (2026) compared topographic and representational similarity across the brain.
  - 414 unrelated HCP–Young Adult participants, organized into 207 non-overlapping pairs.
  - Responses to 23 conditions spanning sensory, motor, and cognitive tasks.

- Do differences in spatial activity patterns reflect differences in the relationships among conditions, or different spatial implementations of similar representational geometry?

Two Ways to Compare Brain Representations

Topographic similarity

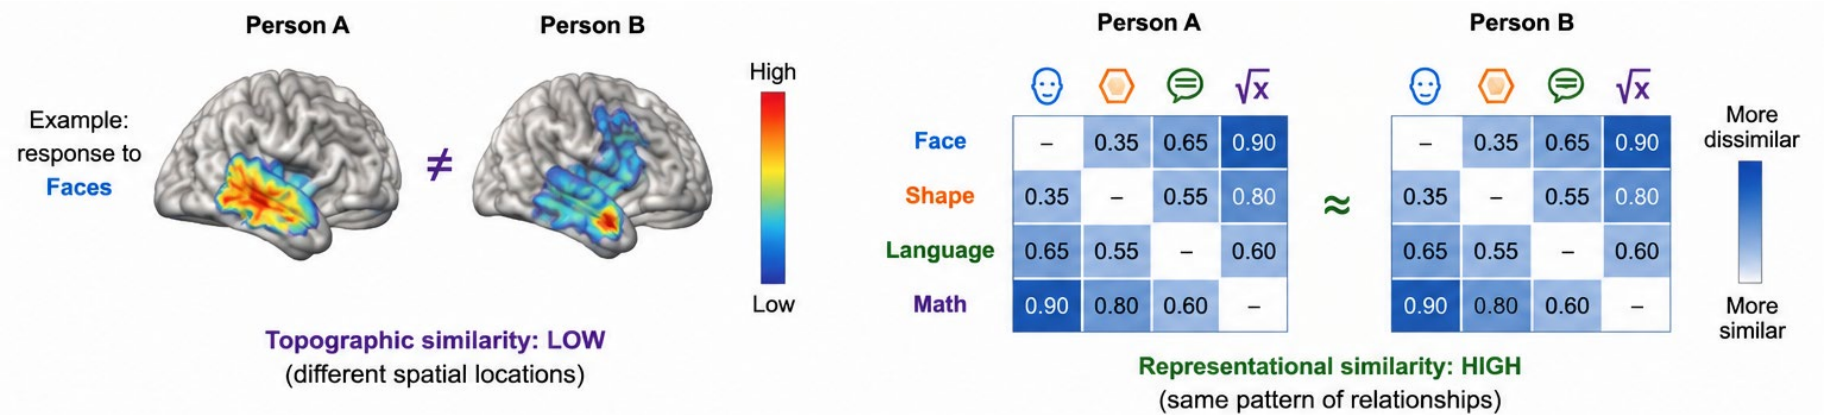
Are activity patterns expressed in corresponding spatial locations?

- Compare responses to the same conditions across people, voxel by voxel.
- High similarity means that corresponding locations show similar responses.

Representational similarity

Are relationships among activity patterns preserved across people?

- Within each person, calculate dissimilarities between conditions to form a representational dissimilarity matrix (RDM).
- Compare the RDMs across people.
- High similarity means that conditions have similar relationships, even when their spatial activity patterns differ.



- Different topographies can preserve the same representational geometry.

Measuring Similarity

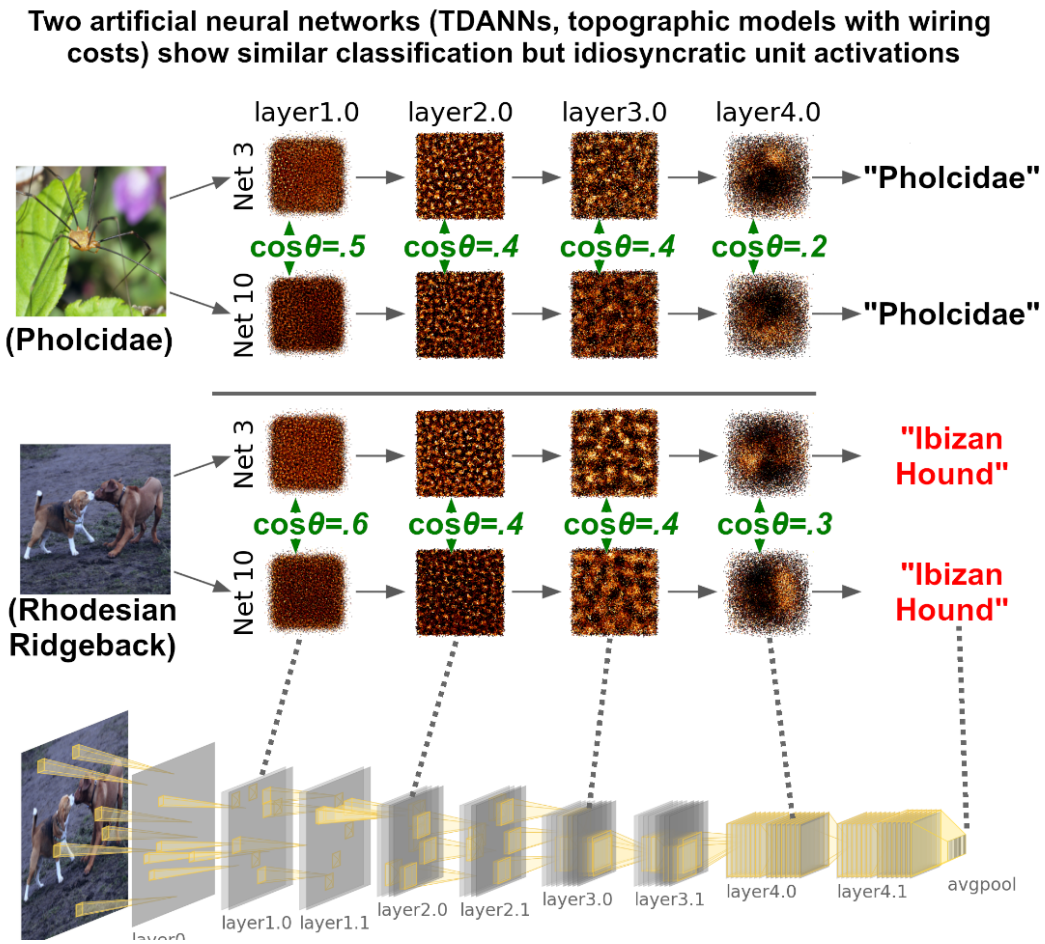
- Topographic similarity
  - Cosine similarity compares responses at corresponding voxels or units.
  - Brain analyses use spatially noise-whitened maps.

- Representational similarity
  - Crossnobis distances form an RDM describing differences between conditions.
  - WUC compares brain RDMs while accounting for dependencies among distance estimates.
  - Linear CKA compares response structure across networks without matching individual units.

- Similarity scores from different measures are not directly interchangeable.

Lesson from Artificial Neural Networks

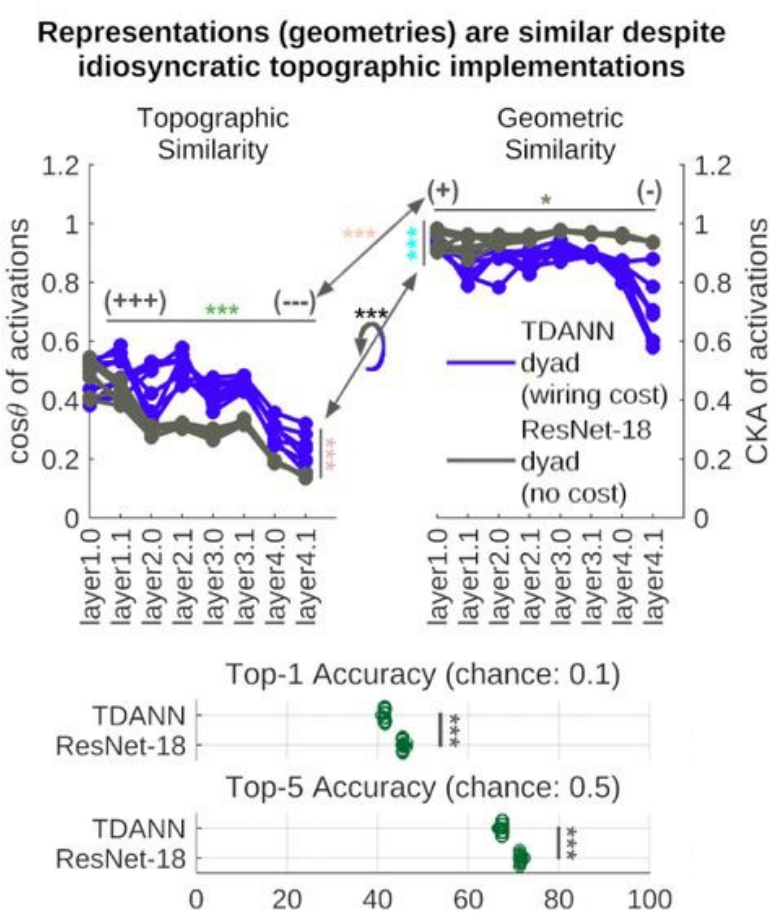
- Different internal activity patterns can produce the same classifications.



Petre et al. (2026)

Lesson from Artificial Neural Networks

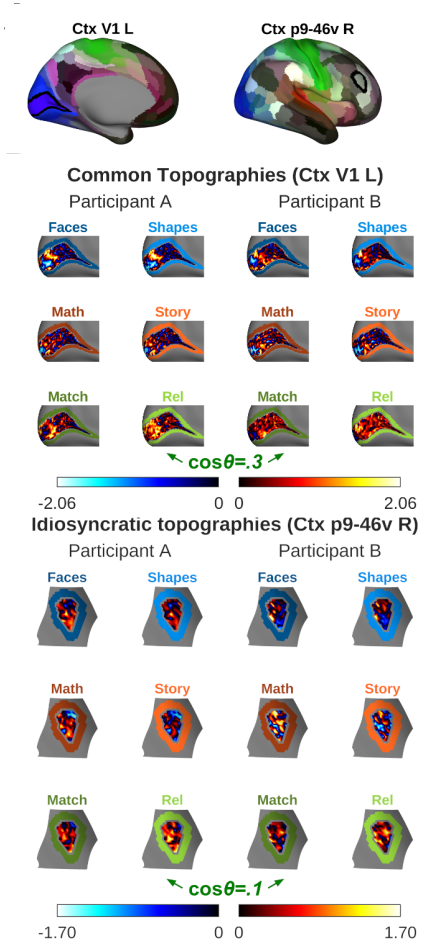
- Representations are similar even when topography isn't.
- Networks without the wiring-cost constraint show lower topographic similarity and higher representational similarity.

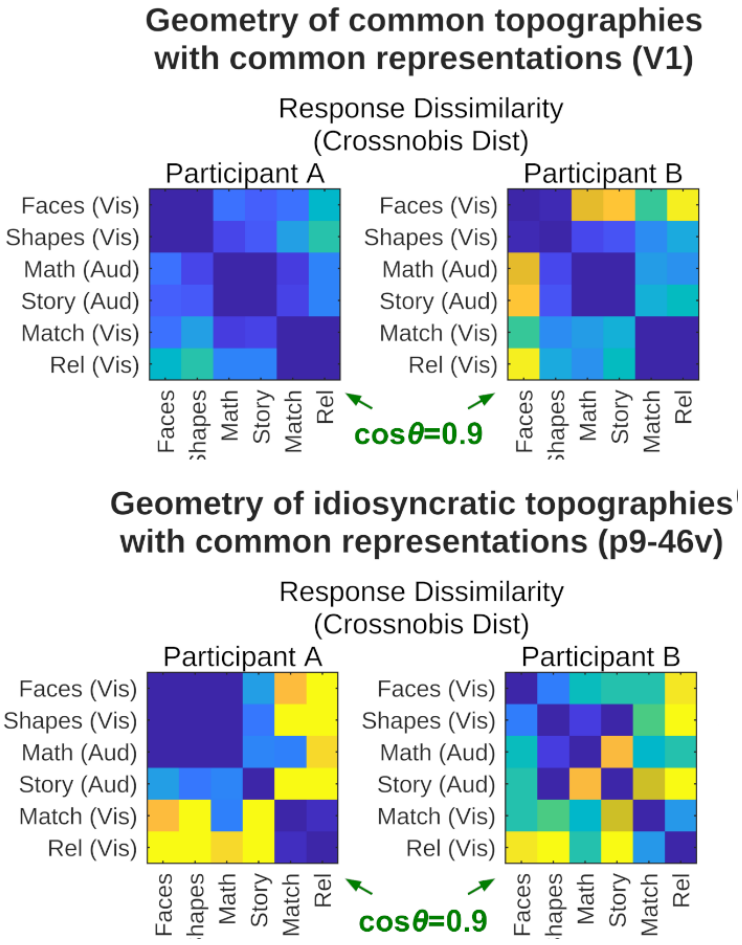


Petre et al. (2026)

Same Comparison Applied to the Brain

- Topographic similarity is higher in V1 than in dlPFC.
- Both regions show similar representational geometry across the sampled task conditions despite differing topographic similarity.

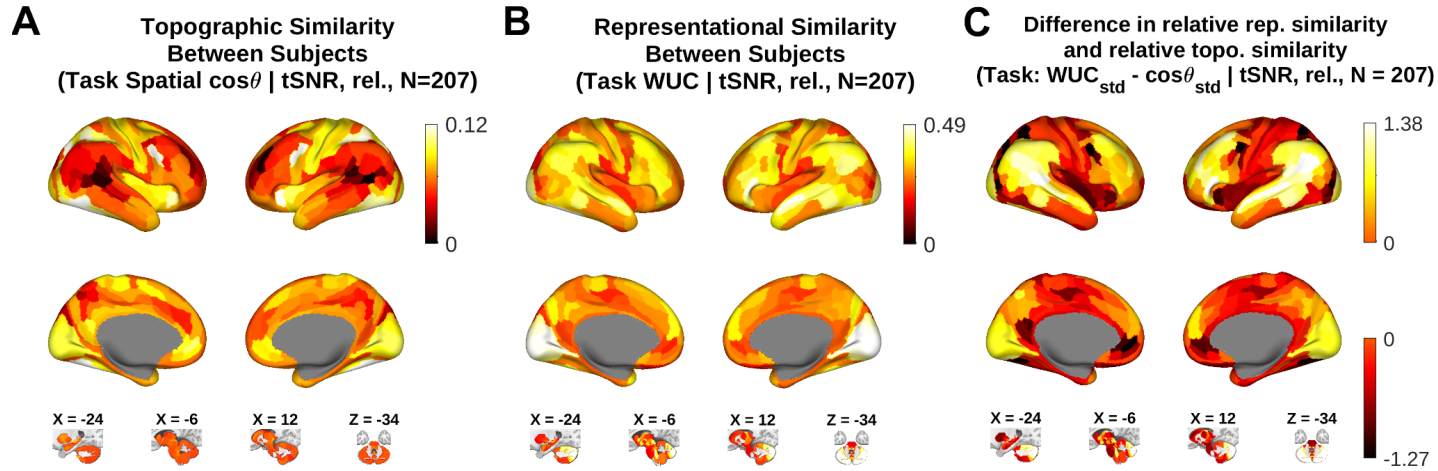




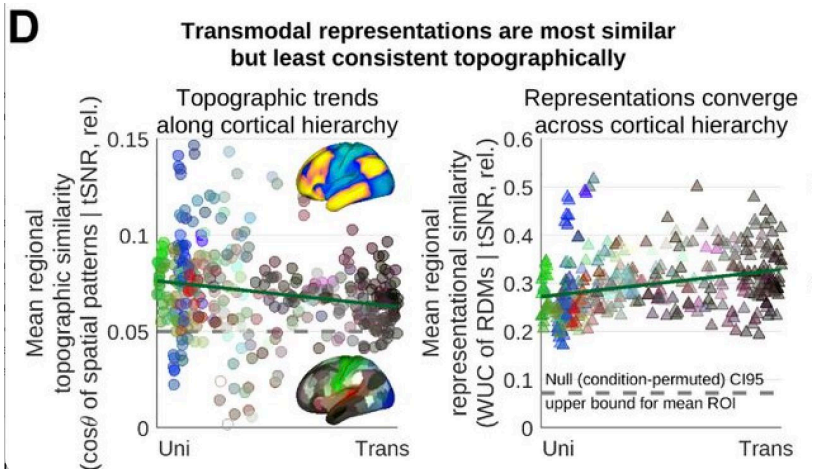
Petre et al. (2026)

Topography Diverges While Representation Converges

- The same pattern extends across the cortical hierarchy.
- Topographic similarity decreases while representational similarity remains high or increases.



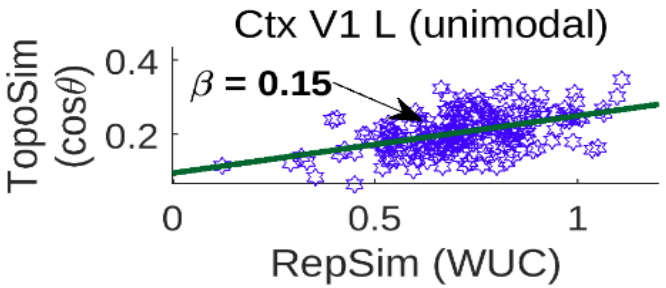
Cortical hierarchy is the progression from primary sensory and motor regions toward higher-order association cortex



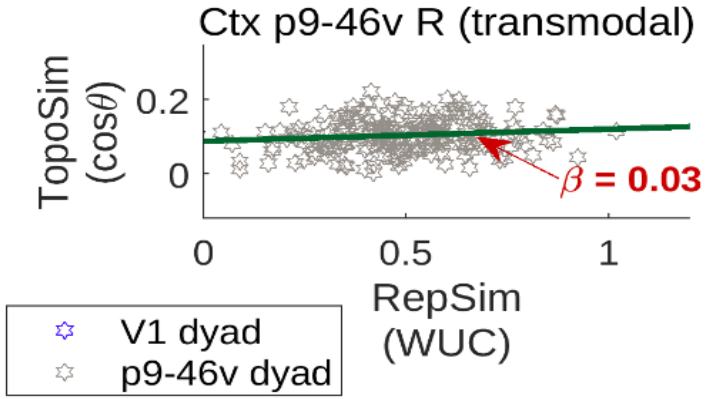
Petre et al. (2026)

Influence of Architectural Constraints

- The association between topographic and representational similarity is stronger in V1 (β=.15) than in dlPFC (β=.03).
- Stronger architectural constraints may favor more similar spatial implementations; weaker constraints may permit similar representations across different topographies.



Unimodal → stronger architectural constraints → fewer viable implementations → topography tracks representation

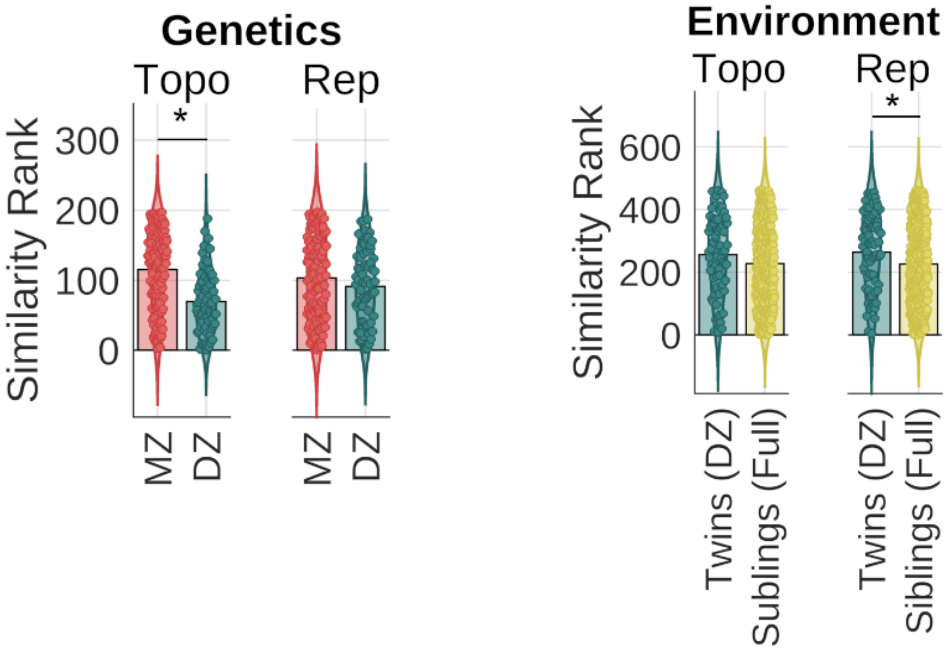


Transmodal → weaker spatial constraints → more viable implementations → representation can converge while topography diverges

Petre et al. (2026)

Genetic and Environmental Contributions

- Twin comparisons suggest stronger genetic influences on topographic organization than on representational geometry.
- Representational geometry also shows evidence of shared environmental influences.
- Environmental effects were not significantly different between the two measures.



Petre et al. (2026)

Shared Structure May Guide Alignment Choice



- Somatosensory cortex
  - Topographic alignment improved performance.
  - Preserved spatial organization may make topographic alignment particularly useful.

- dlPFC / posterior parietal cortex
  - Hyperalignment improved performance.
  - Shared representational geometry despite variable topography may favor geometric alignment.

- The most effective alignment strategy may depend on whether spatial organization or representational geometry is more consistently shared across individuals.
  - This hypothesis requires direct comparisons of the methods on held-out data.

Take-Home Messages

- Functional correspondence is not the same as spatial correspondence.
  - Different individuals or networks can show similar representational relationships despite different topographies.
- Structure shared across individuals provides a starting point for choosing an alignment target.
  - Strongly constrained systems may support meaningful topographic alignment; more flexible systems may benefit from representational alignment.
  - Aligning the appropriate structure can improve downstream inference and prediction.
- The broader AI lesson is to identify which relationships are shared across systems.
  - Do not force correspondence of units, locations, or coordinates when those may be arbitrary; align the structure that is genuinely preserved across systems.

Collaborators

- Johns Hopkins University
  - Guoqing Wang
  - Farzad Farahani
  - Mary Beth Nebel
  - Abhirup Datta
- University of Venice
  - Angela Andreella
- University of Padova
  - Livio Finos
- Dartmouth College
  - Tor Wager
  - Bogdan Petre

Thank you

- References:

Wang, Datta, and Lindquist (2022). "Bayesian Functional Registration of fMRI Data", The Annals of Applied Statistic, 16 (3), 1676-1699.

Wang, Datta, and Lindquist (2023). "Improved fMRI-based Pain Prediction using Bayesian Group-wise Functional Registration", Biostatistics, 25 (3), 885-903.

Andreella, Finos, and Lindquist (2023). "Enhanced hyperalignment via spatial prior information", Human Brain Mapping, 44 (4), 1725-1740.

Farahani, Nebel, Wager, and Lindquist (2023). "Effects of connectivity hyperalignment on brain network properties: from coarse-scale to fine-scale", arXiv.

Jung H, et al. (2025). "Spacetop: A multimodal fMRI dataset unifying naturalistic processes with a rich array of experimental tasks". Sci Data 12(1):1465.

Petre, Lindquist, and Wager (2026). "Representations converge as brain maps diverge along the cortical hierarchy". bioRxiv.

- Funding:

NIH - R01 EB016061
NIH - R01 EB026549

New Book

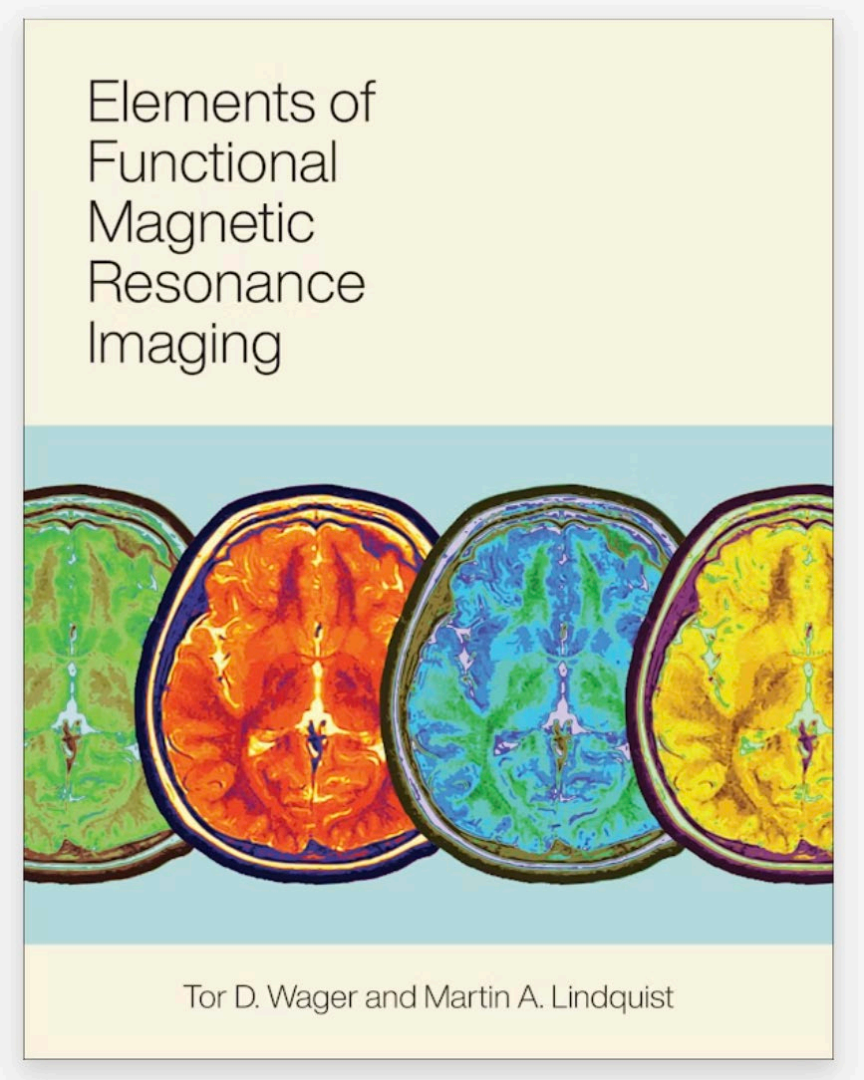
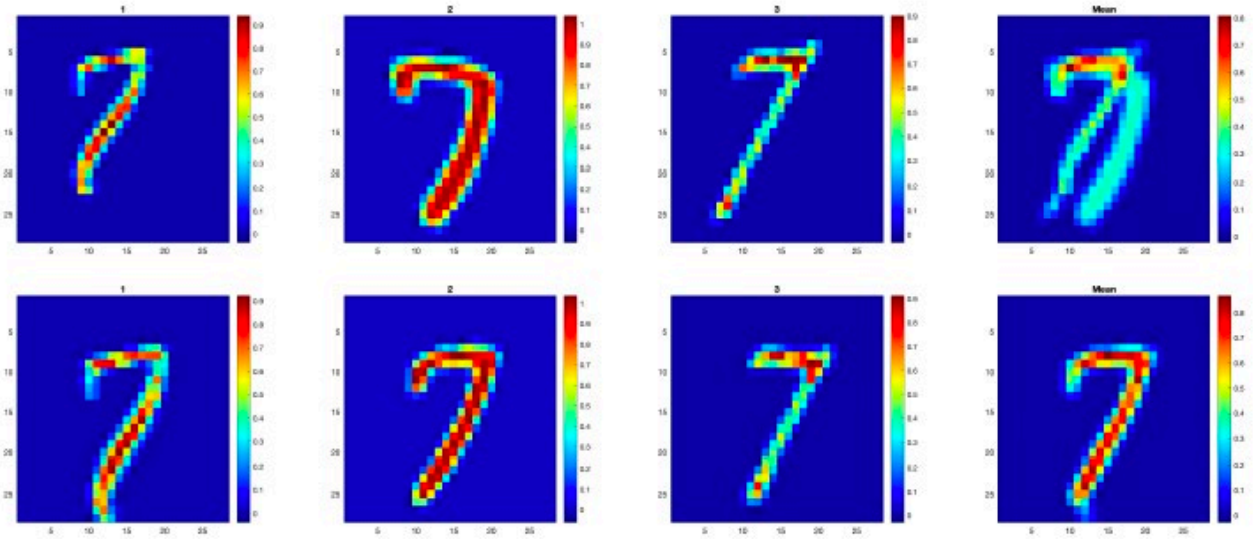

Elements of
Functional
Magnetic
Resonance
Imaging
Tor D. Wager and Martin A. Lindquist

Illustration - Functional Registration



Model

- The latent template is modeled using a Gaussian process with exponential covariance kernel.
  - A major limitation is the computational difficulty
  - We incorporate a nearest-neighbor Gaussian Process (NNGP) to enhance computational efficiency

- T is assumed to be an affine transformation.
  - The transformed latent map $X(T_i)$ given the latent map X is derived using Kriging.
  - The priors for T_i and T_i^r act as regularization terms as the registration problem is inherently ill-posed
    - They follow multivariate t-distributions.

MCMC



- Samples drawn from the posterior using MCMC.
  - Gibbs sampling applied to the latent map, and other parameters updated through Metropolis-Hastings.
  - Updates on the latent template map performed using NNGP with dynamic neighborhood updates.
  - The space of transformation T is generally not linear.
    - Addition of two transformations does not necessarily result in a transformation in the same class.
  - Hence, the proposal distribution is updated on the corresponding tangent space where it is feasible to apply linear operations.