TOPOGRAPHIC SOFT MATCHING DISTANCE (TSMD): A FRAMEWORK FOR COMPARING THE SPATIAL ORGANIZATION OF NEURAL TUNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural populations are spatially organized according to their functional tuning, but how consistently this organization recurs across individuals remains unclear. Existing work typically begins with a predefined tuning variable (*e.g.,* orientation preference, category selectivity), and asks whether its spatial organization appears similar across individuals, often through qualitative comparison of the resulting maps. This becomes limiting for responses to complex natural stimuli, where tuning is high-dimensional and its relevant axes may not be known or easily interpretable. Here, we introduce Topographic Soft-Matching Distance (TSMD), a framework for quantifying shared spatial organization in high-dimensional neural representations without requiring predefined tuning features. TSMD extends soft-matching distance (Khosla & Williams, 2024) using fused Gromov–Wasserstein optimal transport, identifying correspondences between neurons that preserve both tuning similarity and relative spatial geometry. It tests for shared topographic organization by varying the relative weighting on matching neurons with similar tuning and preserving pairwise spatial distances, then comparing the resulting matches against a permutation null that shuffles tuning profiles across fixed spatial locations. After validating TSMD on synthetic orientation maps, we apply it to neural recordings across species and cortical areas. We find shared topographic organization across individuals in macaque V1 and V4, but little evidence in mouse V1 or macaque inferotemporal recordings. In human fMRI data collected during naturalistic viewing, TSMD reveals conserved large-scale semantic maps across individuals that cannot be explained by local smoothness alone. Finally, comparisons of developmental models of V1 topography reveal shared global organization across independent realizations of the same self-organizing model despite differences in their precise layouts, but no reliable correspondence beyond local smoothness between models with different map-forming rules. Together, these results establish TSMD as a general framework for testing the conservation of functional topography across brains and computational systems.

1 INTRODUCTION

For over half a century, neuroscientists have documented systematic relationships between neural tuning and cortical space. Classically, topographic organization refers to maps in which an interpretable tuning variable varies systematically across the cortex, as in the orientation pinwheels of primary visual cortex (Hubel & Wiesel, 1962; Bonhoeffer & Grinvald, 1991). More broadly, spatial organization can refer to the tendency for similarly tuned neurons to occur near one another, producing structures such as category-selective patches for faces (Kanwisher et al., 1997; Tsao et al., 2006), bodies (Downing et al., 2001), and scenes (Epstein et al., 1999), or more generally a decay in response similarity with cortical distance (Lee & DiCarlo, 2019; Margalit et al., 2024). Related forms of organization have been described across auditory (Reale & Imig, 1980; Formisano et al., 2003), entorhinal (Hafting et al., 2005), and olfactory cortex (Brann et al., 2026). Together, these examples illustrate that topography reflects two distinct properties of cortical organization: that nearby neurons tend to have similar tuning, and that specific *kinds* of tuning occupy specific, systematically related locations across cortex.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

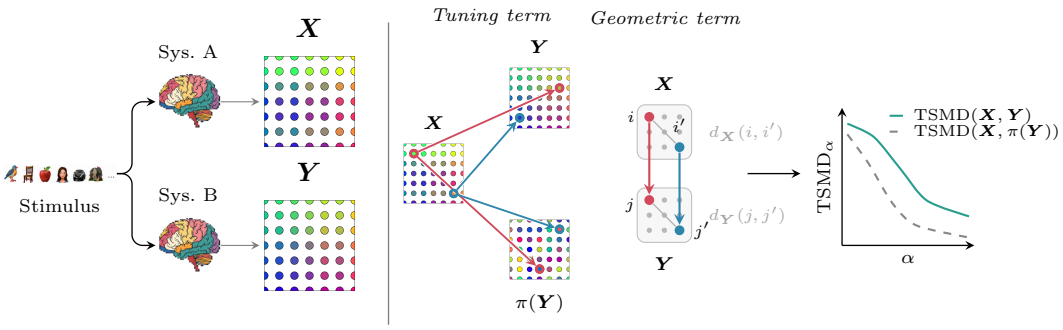

Figure 1: **Topographic Soft-Matching Distance.** The same stimuli are presented to two neural systems, A and B, each producing tuning to a continuous variable (*e.g.,* orientation), yielding tuning matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$. TSMD_α jointly optimizes a *tuning term*, which softly matches neurons across systems by tuning alone, and a *geometric term*, which favors matches preserving spatial distances, i.e., $d_{\boldsymbol{X}}(i,i') \approx d_{\boldsymbol{Y}}(j,j')$ for matched pairs $i \to j$, $i' \to j'$, with α balancing the two. We compare $\mathrm{TSMD}_\alpha(\boldsymbol{X},\boldsymbol{Y})$ against a shuffled control $\mathrm{TSMD}_\alpha(\boldsymbol{X},\pi(\boldsymbol{Y}))$, where $\pi(\boldsymbol{Y})$ preserves $\boldsymbol{Y}$'s tuning distribution but randomizes positions.

Much of the existing literature characterizes topography through the spatial layout of predefined, interpretable tuning properties, or through statistics summarizing that layout, such as the smoothness and size of category-selective patches (Margalit et al., 2024; Deb et al., 2025), spatial relationships among known feature maps (Nauhaus et al., 2012), or a decay in response similarity with cortical distance (Lee & DiCarlo, 2019; Margalit et al., 2024). These statistics, however, capture only the first property above: that nearby neurons share similar tuning. They say nothing about whether specific kinds of tuning occupy consistent locations relative to one another. A distance-decay statistic, for instance, is blind to this distinction — face-, body-, and scene-selective patches could each remain internally tight-knit while their relative arrangement across cortex was entirely rearranged, and the statistic would report no difference. Qualitative similarities across individual maps, together with such recurring statistical properties, suggest that some aspects of functional topography may be conserved across individuals, but similar summary statistics need not imply similar global spatial arrangements of tuning. This gap is especially important for responses to natural stimuli, where tuning is high-dimensional and the relevant features may not be readily interpretable. What remains missing is a general measure of how *closely* the global arrangement of tuning, not just its local smoothness, corresponds across populations, without reducing that organization to a predefined set of tuning features or map statistics.

A separate line of literature provides quantitative tools for comparing neural representations across systems. Measures such as Centered Kernel Alignment (CKA) (Kornblith et al., 2019), Representational Similarity Analysis (RSA) (Kriegeskorte et al., 2008), and Procrustes shape distance (Williams et al., 2021) compare population response geometry but are invariant to rotations that mix the responses of individual neurons. Optimal transport (OT; (Peyré & Cuturi, 2019)) methods instead compare populations by matching neurons according to their response profiles. Soft-matching distance (Khosla & Williams, 2024) preserves sensitivity to individual tuning curves while remaining invariant to neuron ordering, and subsequent work allows partial correspondence between populations (Kapoor et al., 2026). These methods can compare tuning distributions across populations of different sizes, but they do not consider *where* neurons are located. Thus, neither population-level geometric similarity nor correspondence between individual tuning curves establishes whether the spatial arrangement of tuning is shared.

To address this gap, we introduce Topographic Soft-Matching Distance (TSMD), a framework that jointly compares functional tuning and pairwise spatial distances across neural populations (see Fig. 1). We define topographic correspondence as the extent to which neurons can be matched by their tuning while preserving their spatial relationships, independently of the populations' absolute locations. We validate TSMD on synthetic orientation maps, then use it to test for shared topographic organization in mouse and macaque visual cortex, human semantic maps, and developmental models of V1 topography. Together, these analyses move beyond describing topography within individual systems to testing whether the spatial organization of tuning is conserved across them.

2 METHODS

Consider two neural populations with tuning matrices (mean-centered, unit-normed) $\boldsymbol{X} \in \mathbb{R}^{M \times N_x}$ and $\boldsymbol{Y} \in \mathbb{R}^{M \times N_y}$, where M is the number of unique stimulus conditions and N_x and N_y neurons respectively. Each neuron also has a spatial location, $\boldsymbol{u}_i \in \mathbb{R}^2$ for neuron i in $\boldsymbol{X}$ and $\boldsymbol{v}_j \in \mathbb{R}^2$ for neuron j in $\boldsymbol{Y}$.

2.1 SOFT-MATCHING DISTANCE

We review the soft-matching distance (Khosla & Williams, 2024) before describing its topographic analogue. It treats each of the populations $\boldsymbol{X}$ and $\boldsymbol{Y}$ as a uniform empirical measure over tuning curves and computes the OT cost between them $d_{\boldsymbol{T}}(\boldsymbol{X}, \boldsymbol{Y}) = \min_{\boldsymbol{T} \in \mathcal{T}(N_x, N_y)} \langle \boldsymbol{C}, \boldsymbol{T} \rangle_F$, where $\boldsymbol{C}_{ij} = 1 - \boldsymbol{x}_i \boldsymbol{y}_j^\top$ is the cosine cost, $\langle \cdot, \cdot \rangle_F$ the Frobenius inner product, and $\mathcal{T}(N_x, N_y)$ the transportation polytope (De Loera & Kim, 2013):

$$\mathcal{T}(N_x, N_y) = \left\{ \boldsymbol{T} \in \mathbb{R}_{\geq 0}^{N_x \times N_y} \;:\; \boldsymbol{T}\mathbf{1}_{N_y} = \tfrac{1}{N_x}\mathbf{1}_{N_x},\;\; \boldsymbol{T}^\top \mathbf{1}_{N_x} = \tfrac{1}{N_y}\mathbf{1}_{N_y} \right\}, \tag{1}$$

whose uniform marginals assign each neuron equal mass and allow populations of unequal sizes to be compared. The minimizing coupling is invariant to permutations of the neuron indices, yet sensitive to rotations of the representational axes, which is ignored by rotation-invariant metrics like CKA and Procrustes shape distance. Moreover, $d_{\boldsymbol{T}}$ has a true metric structure, which has been shown to have appealing properties for downstream analyses of neural representations (Williams et al., 2021; Lange et al., 2023) like clustering. Its cost, however, is a function of tuning alone: two populations with identical tuning curves are assigned an identical coupling regardless of how their neurons are positioned in space. TSMD addresses this limtation.

2.2 TOPOGRAPHIC SOFT-MATCHING DISTANCE (TSMD)

TSMD compares $\boldsymbol{X}$ and $\boldsymbol{Y}$ through a soft assignment $\boldsymbol{T} \in \mathbb{R}^{N_x \times N_y}$ using two costs. The first is a cosine tuning cost, identical to soft-matching ($\boldsymbol{C}_{ij} = 1 - \boldsymbol{x}_i^\top \boldsymbol{y}_j$), which penalizes matches between functionally dissimilar neurons. Next, a geometric cost compares the intra-population distances $d_{\boldsymbol{X}}(i, i') = \|\boldsymbol{u}_i - \boldsymbol{u}_{i'}\|^2$ and $d_{\boldsymbol{Y}}(j, j') = \|\boldsymbol{v}_j - \boldsymbol{v}_{j'}\|^2$, and penalizes the coupling by the discrepancy $|d_{\boldsymbol{X}}(i, i') - d_{\boldsymbol{Y}}(j, j')|^2$ whenever it transports mass jointly through (i, j) and (i', j'). Intuitively, this cost is small when neurons adjacent in $\boldsymbol{X}$ are matched to neurons adjacent in $\boldsymbol{Y}$. In essence, this depends on the relative spatial arrangement of the two neural populations. The norm $\|\cdot\|$ may be Euclidean or any biologically appropriate metric, such as a geodesic distance on the cortical surface that accounts for gyri and sulci.

A single hyperparameter $\alpha \in [0, 1]$ balances the two costs, giving the following TSMD objective:

$$\mathcal{L}_\alpha(\boldsymbol{T}) \;=\; \alpha \underbrace{\sum_{i,i',j,j'} \left| d_{\boldsymbol{X}}(i, i') - d_{\boldsymbol{Y}}(j, j') \right|^2 \boldsymbol{T}_{ij}\, \boldsymbol{T}_{i'j'}}_{\text{Geometric term}} + (1 - \alpha) \underbrace{\sum_{i,j} \boldsymbol{C}_{ij}\, \boldsymbol{T}_{ij}}_{\text{Tuning term}},$$

minimized over the transportation polytope $\mathcal{T}(N_x, N_y)$ defined in Eq. 1. We write $\mathrm{TSMD}_\alpha(\boldsymbol{X}, \boldsymbol{Y}) = \min_{\boldsymbol{T} \in \mathcal{T}(N_x, N_y)} \mathcal{L}_\alpha(\boldsymbol{T})$. Simply put, the marginals distribute each neuron's mass across its matches while conserving total mass, so the two populations may differ in size. Hence, under this objective, a pair of neurons is coupled strongly only when it is favored by both terms at once: when the neurons are functionally similar ***and*** lie in matching topographic neighborhoods.

This convex combination is exactly the fused Gromov–Wasserstein distance (Vayer et al., 2020), which also helps define the metric structure of TSMD. On the space of tuning-and-location populations, TSMD_α is symmetric and vanishes only for isomorphic populations (see A5.1), meaning each unit in one population has a matching unit in the other with identical tuning and location. It satisfies the triangle inequality only up to a constant factor, making it a **semi-metric** (Proposition 5.3 of Vayer et al., 2018). At $\alpha = 0$ the geometric term vanishes and $\mathcal{L}_\alpha$ reduces exactly to the soft-matching distance of Section 2.1, matching neurons by tuning alone; conversely, at $\alpha = 1$ it matches by spatial arrangement alone. Between these endpoints, TSMD balances tuning against topographic fidelity.

2.3 Null Test for Topography

Rather than fix α, we sweep $\alpha \in [0, 1)$, excluding $\alpha = 1$ where the geometric cost admits degenerate, physically implausible solutions. Denote $A(\alpha)$ as the statistic evaluated at the optimal coupling $\boldsymbol{T}(\alpha)$. This is the coupling-weighted mean tuning correlation between matched neuron pairs, where larger values indicate better functional correspondence under the spatial constraint. We refer to $A(\alpha)$ as ***TSMD similarity*** and summarize it as $\mathrm{AUC} = \int_0^1 A(\alpha)\, \mathrm{d}\alpha$.

We test for topography by comparing this *original* AUC against a permutation null, constructed by reassigning tuning curves to spatial locations. This conserves the tuning distribution while destroying spatial–tuning pairing. Repeating the sweep on the permuted population yields a null curve $A_\pi(\alpha)$ and area AUC_π. We assess null significance over $K = 15$ resamples. For within-population comparisons, each split divides electrodes into two disjoint halves, and for cross-population comparisons, each split subsamples half of each population's electrodes. In both cases, the null permutes tuning across electrode positions in one group. For each split k, we compute the paired difference $\Delta_k = \mathrm{AUC}_k - \mathrm{AUC}_{\pi,k}$, which should be near zero under the null hypothesis of no topographic structure, and test $\{\Delta_k\}_{k=1}^K$ against zero with a one-sided Wilcoxon signed-rank test.

3 Results

3.1 Synthetic Tuning Maps

To test whether TSMD jointly captures tuning and topographic correspondence, we validate it on synthetic orientation maps, similar to pinwheel maps observed in V1. Synthetic maps are useful here because their ground-truth structure is known. We describe their construction in Appendix A1.

Let $\boldsymbol{X}_{\text{base}} \in \mathbb{R}^{M \times N}$ denote the base cortical sheet, with M orientation directions and N neurons. We transform this base map in ways that selectively preserve or disrupt its tuning and spatial organization, and compute TSMD similarity (over the α-continuum) between $\boldsymbol{X}_{\text{base}}$ and two families of transformed copies. First, we consider two **rigid transformations**: translation, which shifts the sheet by (t_x, t_y) in coordinate space, and rotation, which turns it by γ° about the center. Neither changes the neuronal neighbors or how their tuning is organized relative to one another, so we expect alignment to remain nearly constant across α (see Appendix A5.1). Second, to isolate the effect of spatial organization from tuning, we apply a **permutation**: given a random permutation $\pi : [N] \to [N]$, a location (or neuron) i preserves its coordinates but is now assigned the tuning of neuron $\pi(i)$, preserving the overall set of tuning profiles while destroying their spatial organization. Alignment should therefore be high at $\alpha = 0$ (pure functional matching) and reduce monotonically as $\alpha \to 1$ (pure spatial matching). We demonstrate this in Fig. 2, where rotation and translation leave alignment near-constant, while permutation reduces monotonically as α increases.

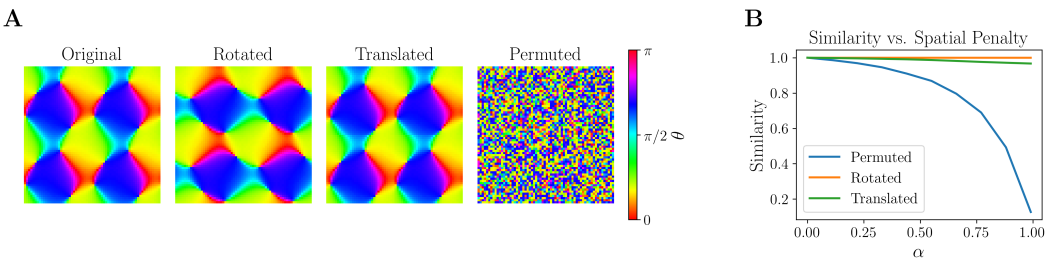

Figure 2: **Geometric Transforms for a Simulated Orientation Map.** (**A**) A synthetic pinwheel map is rotated, translated, and randomly permuted. Color indicates preferred orientation $\theta \in [0, \pi)$. (**B**) Rotation and translation preserve TSMD alignment (near-constant), while permutation reduces monotonically as the spatial penalty α increases.

3.2 Topography in Neural Datasets

Topographic maps of tuning preference in the visual cortex are known to exhibit stereotyped spatial patterns. Yet, the degree of similarity both within and across individuals remains poorly quantified.

Prior work measures an autocorrelation of tuning and spatial frequency as a function of absolute cortical distance (Nauhaus et al., 2012). A larger and relatively recent body of work quantifies topography using absolute pairwise differences in preferred orientation as a function of cortical distance (Ringach et al., 2016; Chang et al., 2020; Ferreiro et al., 2021; Margalit et al., 2024). In the sections that follow, we instead apply TSMD to multi-unit electrophysiology recordings in mice and macaques.

3.2.1 Allen Brain Dataset

To test for topography in mouse V1, we use two-photon calcium imaging data from the Allen Brain Observatory (ABO) Visual Coding dataset (de Vries et al., 2020), which records somatic calcium fluorescence from genetically identified excitatory neurons in response to drifting gratings[1]. We selected $N = 10$ mice with imaged populations in primary visual cortex (`VISp`). For each neuron, we computed the mean $\Delta F/F$ response across trials for each of the 8 orientations, yielding a tuning matrix $\boldsymbol{X} \in \mathbb{R}^{8\times N}$ per mouse, with spatial coordinates taken from the two-photon field of view (a 2D tangential plane through the cortical surface). Full stimulus parameters, cell-selection criteria, and preprocessing steps are provided in Appendix A2.

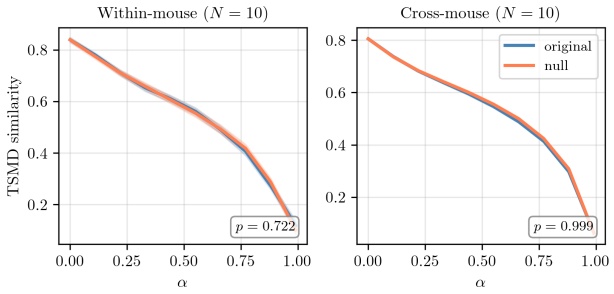

Figure 3: **Similarity vs. Spatial Penalty Within and Across Mice Pairs.** We randomly permute the orientation tunings associated with neurons at different spatial locations (null model), and observe that the topographic similarity increases with respect to the original model over the entire α continuum.

We observe that similarities both within- and across-mice remain indistinguishable from the permutation null across α (Fig. 3). We use our permutation null defined in Section 2.3 to obtain significance values $p = 0.722$ within and $p > 0.999$ across $N = 10$ mice. In other words, TSMD finds no significant topographic structure in mouse V1, which is consistent with the widely documented *"salt-and-pepper"* organization (Ohki et al., 2005; Ohki & Reid, 2007; Bonin et al., 2011). This motivates extending our analysis to a species with known columnar organization, where TSMD should detect topographic structure if the metric is sufficiently sensitive.

3.2.2 THINGS Ventral Stream Spiking Data

Unlike mice, macaque V1 exhibits well-established columnar organization of orientation preference (*e.g.,* pinwheel-like structures) that has been extensively documented by means of two-photon imaging (Ohki et al., 2006; Nauhaus et al., 2012). This serves as our positive control, in that TSMD should detect topography where it is known to exist. We analyze multi-unit spiking activity (MUA) from the THINGS Ventral Stream Spiking Dataset (Papale et al., 2025), which records from two monkeys (N and F). Each monkey was implanted with Utah electrode arrays across three ventral stream regions: V1, V4, and the inferotemporal cortex (IT). We take the normalized MUA responses to 22,248 natural images from the THINGS database (Hebart et al., 2019) as the tuning representation[2]. Electrode coordinates on the cortical surface are obtained from the array placement

[1] We used full-field drifting gratings because response correlations to spatially structured natural scenes can reflect overlapping receptive fields rather than columnar feature organization; Appendix A3 demonstrates this confound empirically.

[2] This needs to be interpreted in a rather nuanced manner. Natural-image responses can conflate feature tuning with receptive-field position. We address this using factorized readouts that separate spatial and feature

images (see Appendix A4 for details). We run two analyses per region: a within-monkey test and a

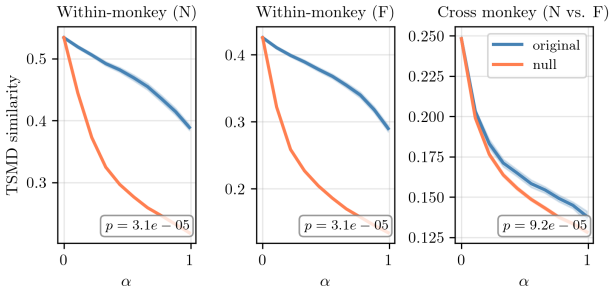

Figure 4: **Topographic Structure in Macaque V1.** (**Left**, **Center**) Within-monkey split-half comparisons for monkeys N and F. (**Right**) Cross-monkey comparison between N and F. In all panels, the original (blue) scores lie consistently above the permutation null (orange), indicating shared topographic organization within and across monkeys (Wilcoxon signed-rank test on AUC differences). Error bands denote standard error over splits.

cross-monkey comparison. For the within-monkey test, we split a region's electrodes into two halves and compute TSMD between them and compare this score to a null obtained by permuting tuning profiles across electrode positions. If the region has topographic structure, the true split should score above this null. For the cross-monkey comparison, we compute TSMD directly between the populations of monkeys N and F, constructing the null by permuting the tuning assignments within one monkey. A score above this null indicates shared topography between the two monkeys.

In V1, the within-monkey score lie above its null across all α for both monkeys N ($p = 3.1 \times 10^{-5}$) and F ($p = 3.1 \times 10^{-5}$), indicating that nearby electrodes share tuning to natural images (Fig. 4). We further demonstrate that this signal is not just a retinotopic artifact by performing checkerboard and subsampling controls that break receptive-field overlap (Appendix A5). Moreover, we factorize the signal into its "where" and "what" components and note identical qualitative and quantitative results (see Sec. 3.3). A similar pattern holds when we go up the hierarchy of the ventral stream (Fig. 5). Both V4 and IT exhibit within-monkey topography (V4: $p = 3.1 \times 10^{-5}$; IT: $p = 3.1 \times 10^{-5}$). We note that V4 also exhibits shared cross-monkey structure ($p = 3.1 \times 10^{-5}$), while IT does not. We caution readers to interpret the latter result, given IT's loosely defined anatomical boundaries, a caveat we discuss further in Section 3.3.

3.3 Separating Feature Topography from Retinotopy

Our macaque results in Sec. 3.2.2 rely on an assumption we swept under the rug and now tackle directly. A response to a natural image encodes both *what* features a neuron prefers and *where* in the visual field it looks. Therefore, two adjacent electrodes can naïvely have correlated responses because they share a receptive-field location with no genuine similarity in feature tuning. We refer to the spatial organization of feature selectivity as **feature topography** and of receptive-field locations as **retinotopy**. Since TSMD does not distinguish between them, we factor the "where" component out of the tuning to cleanly attribute our macaque signal to feature topography.

We do this with a factorized encoding model in the spirit of Klindt et al. (2017). Each image is passed through a frozen, ImageNet-pretrained VGG16 backbone (Simonyan & Zisserman, 2015). This produces a feature tensor $\boldsymbol{F} \in \mathbb{R}^{C \times H \times W}$ shared across all electrodes. For each electrode n, we then fit a factorized readout that factorizes into a spatial mask $\boldsymbol{m}_n \in \mathbb{R}^{H \times W}$ and a feature-weight vector $\boldsymbol{w}_n \in \mathbb{R}^C$, which are the "where" and "what" components, respectively. The predicted response is $r_n = f\left(\boldsymbol{w}_n^\top \sum_{h,w} \boldsymbol{m}_n[h,w]\, \boldsymbol{F}[:,h,w]\right)$, where $f(\cdot)$ is the pointwise output non-linearity to ensure non-negative predicted responses. The mask pools $\boldsymbol{F}$ over the electrode's receptive-field and $\boldsymbol{w}_n$ reads out its tuning. We hold the backbone fixed, fit only the readout, sweep across VGG16 layers, and select for each electrode the layer with the highest noise-corrected held-out correlation.

components (Sec. 3.3) together with checkerboard and subsampling controls that reduce sensitivity to short-range electrode correlations (Appendix A5).

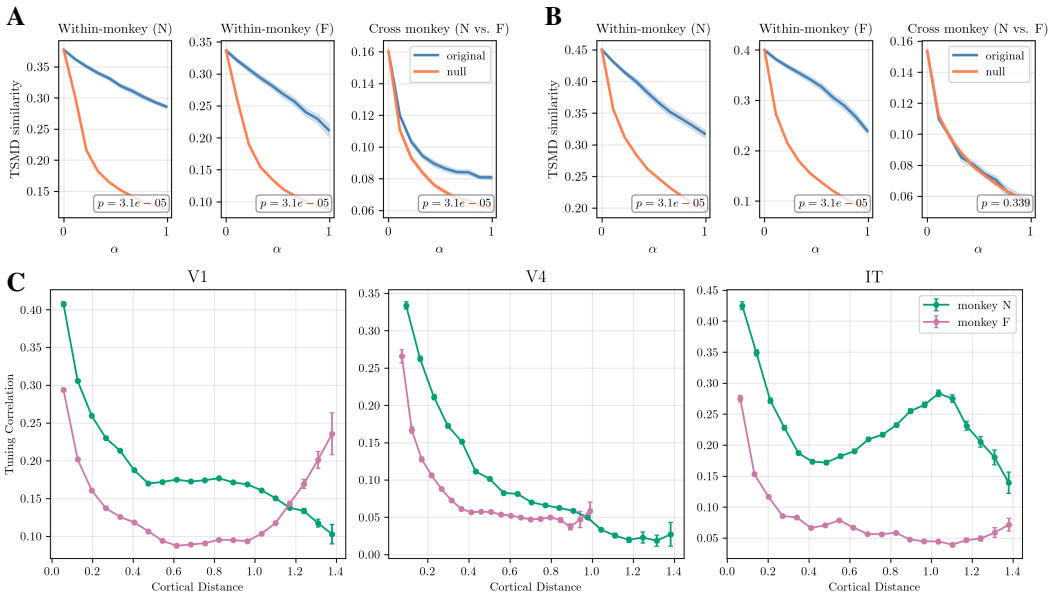

Figure 5: **Topographic Structure in V4 and IT.** Macaque V4 (**A**) shows both local and shared topography, whereas IT (**B**) shows local but no shared topography. (**C**) Using the canonical definition of topographic alignment, we observe that tuning correlation decays with cortical distance in all regions. Error bars denote SEM over electrode pairs in each of $N = 20$ distance bins.

We run this analysis on TVSD only, since its image stimuli provide sufficient coverage to estimate receptive-field locations. However, the Allen natural-scene set is too small for a reliable "where" estimate (118 natural images). Since retinotopic position is absorbed into $\boldsymbol{m}_n$, the feature vector $\boldsymbol{w}_n$ carries only the tuning profile with location factored out. Using the feature weights, we construct a new tuning matrix $\boldsymbol{W} \in \mathbb{R}^{C \times N}$ (N electrodes), and recompute TSMD. If the macaque signal is an artifact of retinotopy, this control should yield no statistically significant result. Instead, within- and cross-monkey topography in V1 and V4 persists under the factorized readout, matching the raw-MUA analysis both qualitatively and quantitatively (Fig. 6).

IT Cross-Monkey Result. The factorized readout also qualifies our IT result. Within each monkey, IT is topographically organized, but this organization is not shared across the pair. We do not read this as evidence against canonical IT organization, since, unlike V1 and V4, IT lacks anatomical landmarks for registering electrode coordinates across animals. Therefore, a cross-monkey test requires homologous sampling that cannot be guaranteed here. To our knowledge, no dataset combines functional IT localization with two-dimensional probe coverage. The closest is the Triple-N (Li et al., 2026) dataset, which records IT along a single Neuropixels transect per animal, precluding such a comparison. We therefore treat the absent cross-monkey IT signal as a limitation of available recordings rather than evidence against shared IT topography.

3.4 Semantic topography is conserved across individuals

Prior voxel-wise encoding studies of fMRI acquired during narrative listening produced cortical semantic maps that appeared qualitatively similar across individuals (Huth et al., 2016). We used our TSMD framework to test this apparent conservation quantitatively in four participants from the CNeuroMod *Friends* dataset (Gifford et al., 2024). Using semantic embeddings of the viewed narrative, we fitted ridge encoding models predicting each Schaefer-1000 parcel's response (Schaefer et al., 2018). The ridge coefficient vector defined each parcel's semantic tuning, while its MNI centroid defined its cortical location (see Appendix A6 for further details). Across all six participant pairs, the observed TSMD AUC exceeded both a full tuning-to-location permutation null and a conservative patch null (Fig. 7, Appendix A5). All comparisons remained significant after correction across participant pairs, patch scales, and null families ($q_{\mathrm{FDR}} \leq 6.1 \times 10^{-5}$). The patch null further suggests that this result cannot be attributed to local fMRI smoothness alone. These findings

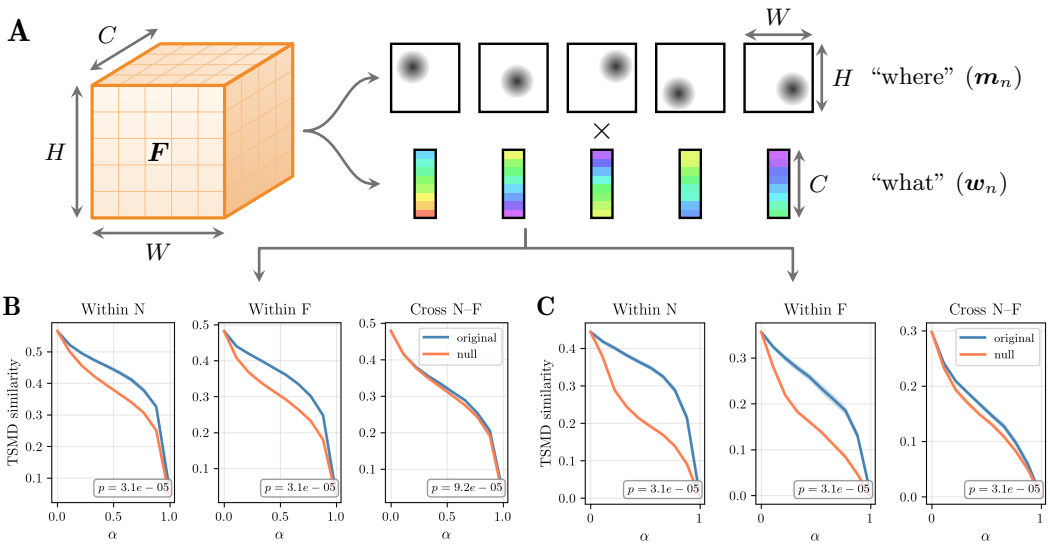

Figure 6: **Feature Topography after Factoring Out Retinotopy.** (**A**) A frozen visual backbone produces a feature tensor ($\boldsymbol{F}$) that is mapped to each electrode by a factorized readout comprising a spatial mask ("where") and feature weights ("what"). (**B–C**) TSMD computed from the feature-weight vectors in V1 and V4, respectively. Within- and cross-monkey observed curves remain above their tuning-to-location permutation nulls, indicating that the correspondence is not explained solely by shared receptive-field position. Lines and shading denote mean ± SEM across resamples.

provide the first quantitative evidence that the large-scale spatial organization of semantic tuning is conserved across participants.

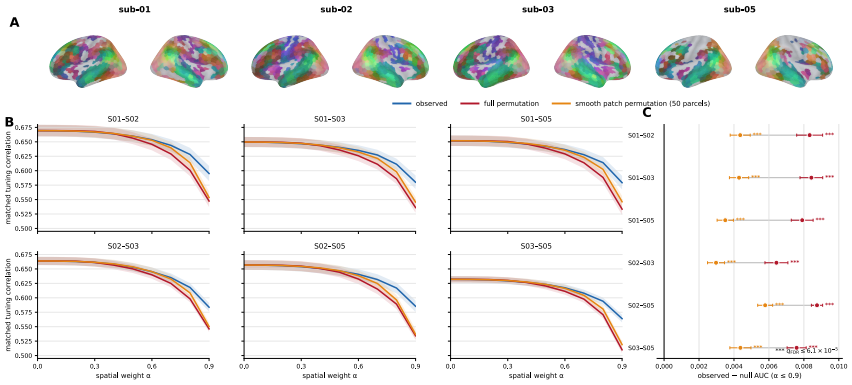

Figure 7: **Conserved Semantic Topography in CNeuroMod.** (**A**) Semantic tuning projected onto a common RGB basis; unreliable parcels are gray. (**B**) TSMD curves for all participant pairs under the observed, full-permutation, and 50-parcel smooth-patch conditions. Lines and shading show mean ± SEM over 15 resamples. (**C**) Observed-minus-null AUC. Both nulls are rejected for every pair ($q_{\mathrm{FDR}} \leq 6.1 \times 10^{-5}$).

3.5 Applying TSMD to Developmental Models of Cortical Topography

We next apply TSMD to topographies produced by computational models. This helps answer whether independent developmental realizations converge on a shared spatial organization and whether distinct map-forming mechanisms produce related topographies. We answer both questions using generative models of V1-like organization, long studied for how retinotopy and columnar feature maps emerge from pressures for feature coverage and cortical continuity (Durbin & Mitchison, 1990; Obermayer et al., 1992; Erwin et al., 1995). In a self-organizing map (SOM), units on a

two-dimensional sheet compete to represent sampled visual features, and updating the winning unit and its neighbors yields a smooth map balancing feature coverage and spatial continuity (Kohonen, 1982). Different random seeds can therefore place individual domains and pinwheels differently while preserving the same developmental rule.

We generated 96×96 maps from a common feature space comprising retinal position, periodic orientation, eye identity, and spatial frequency. We first compare independent SOM realizations to each other and then to an elastic-net model that balances the same pressures through a different optimization procedure (Durbin & Mitchison, 1990). Both mechanisms produced maps with smooth retinotopy, orientation pinwheels, ocular-dominance domains, and spatial-frequency domains (representative maps and simulation details in Appendix A7).

For the TSMD analysis, maps were pooled to 24×24 and tuning was restricted to columnar features (orientation, ocular dominance, and spatial frequency) so that shared coarse retinotopy could not determine the result. We compared against a full-permutation null that destroyed tuning-location dependence and a patch-permutation null that preserved local smoothness while disrupting the global arrangement of tuning domains. We found that independent SOM realizations exhibited greater correspondence than both full-permutation ($q = 2.4 \times 10^{-4}$) and patch-permutation nulls ($q = 0.010$; Fig. 8). By contrast, SOM-elastic-net correspondence exceeded the full null ($q = 0.0084$) but not the patch null ($q = 0.319$). Thus, TSMD revealed reproducible global columnar organization across realizations of the same developmental rule despite differences in layouts, but found no evidence that the two mechanisms share global organization beyond generic local smoothness.

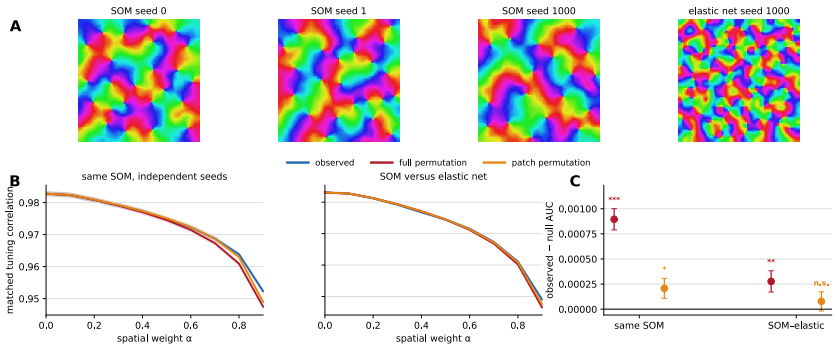

Figure 8: **TSMD on Self-Organizing V1-like Maps.** (**A**) Orientation maps from independent SOMs and an elastic net. (**B**) Columnar-only TSMD curves under the observed, full, and patch-permutation conditions. (**C**) Observed-minus-null AUC over 15 independent map pairs.

4 Discussion

We introduced the Topographic Soft-Matching Distance (TSMD), a framework for comparing neural populations jointly by their functional tuning and relative spatial organization. In synthetic orientation maps, TSMD alignment remains stable under rigid transformation, and reduces monotonically under permutations. In real neural data, TSMD recovers well-documented absence of columnar organization in mouse V1, detects significant local and shared organization in macaque V1 and V4 which persists even after factoring out retinotopy. In IT however, topography was present within each monkey but not shared reliably across monkeys, which we attribute to the lack of dataset which combines functional IT localization with two-dimensional probe coverage rather than due to absence of underlying organization. Beyond visual cortex, TSMD also demonstrates conservation of large-scale organization of semantic tuning during natural language listening in human participants. Last, in developmental models of cortical organization, we show that independent seeds of the same self-organizing rule converge on a shared global columnar arrangement. On the other hand, maps produced from two distinct rules (SOM and elastic net) produce maps that agree only at the level of local smoothness, but not global organization. Taken together, these results indicate that convergent topographic is complex signature that depends on species, cortical region, and, in models, the specific developmental rule. TSMD can thus adjudicate between competing hypotheses about when, where, and why such organization emerges, across biological and artificial systems alike.

REFERENCES

Tobias Bonhoeffer and Amiram Grinvald. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, 353(6343):429–431, 1991.

Vincent Bonin, Mark H Histed, Sergey Yurgenson, and R Clay Reid. Local diversity and fine-scale organization of receptive fields in mouse visual cortex. *Journal of Neuroscience*, 31(50):18506–18521, 2011.

David H Brann, Tatsuya Tsukahara, Cyrus Tau, Dennis Kalloor, Rylin Lubash, Lakshanyaa Thamarai Kannan, Nell Klimpert, Mihaly Kollo, Martín Escamilla-Del-Arenal, Bogdan Bintu, et al. A spatial code governs olfactory receptor choice and aligns sensory maps in the nose and brain. *Cell*, 189(11):3358–3379, 2026.

Jeremy T Chang, David Whitney, and David Fitzpatrick. Experience-dependent reorganization drives development of a binocularly unified cortical representation of orientation. *Neuron*, 107 (2):338–350, 2020.

Jesús A De Loera and Edward D Kim. Combinatorics and geometry of transportation polytopes: An update. *Discrete geometry and algebraic combinatorics*, 625:37–76, 2013.

Saskia EJ de Vries, Jerome A Lecoq, Michael A Buice, Peter A Groblewski, Gabriel K Ocker, Michael Oliver, David Feng, Nicholas Cain, Peter Ledochowitsch, Daniel Millman, et al. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature neuroscience*, 23(1):138–151, 2020.

Mayukh Deb, Mainak Deb, and Apurva Murty. Toponets: High performing vision and language models with brain-like topography. In *International Conference on Learning Representations*, volume 2025, pp. 39099–39119, 2025.

Paul E Downing, Yuhong Jiang, Miles Shuman, and Nancy Kanwisher. A cortical area selective for visual processing of the human body. *Science*, 293(5539):2470–2473, 2001.

Richard Durbin and Graeme Mitchison. A dimension reduction framework for understanding cortical maps. *Nature*, 343(6259):644–647, 1990.

Russell Epstein, Alison Harris, Damian Stanley, and Nancy Kanwisher. The parahippocampal place area: recognition, navigation, or encoding? *Neuron*, 23(1):115–125, 1999.

Ed Erwin, Klaus Obermayer, and Klaus Schulten. Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural computation*, 7(3):425–468, 1995.

Dardo N Ferreiro, Sergio A Conde-Ocazionez, Joao HN Patriota, Lua C Souza, Moacir F Oliveira, Fred Wolf, and Kerstin E Schmidt. Spatial clustering of orientation preference in primary visual cortex of the large rodent agouti. *Iscience*, 24(1), 2021.

Elia Formisano, Dae-Shik Kim, Francesco Di Salle, Pierre-Francois Van de Moortele, Kamil Ugurbil, and Rainer Goebel. Mirror-symmetric tonotopic maps in human primary auditory cortex. *Neuron*, 40(4):859–869, 2003.

Alessandro T Gifford, Domenic Bersch, Marie St-Laurent, Basile Pinsard, Julie Boyle, Lune Bellec, Aude Oliva, Gemma Roig, and Radoslaw M Cichy. The algonauts project 2025 challenge: How the human brain makes sense of multimodal movies. *arXiv preprint arXiv:2501.00504*, 2024.

Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.

Martin N Hebart, Adam H Dickter, Alexis Kidder, Wan Y Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I Baker. Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PloS one*, 14(10):e0223792, 2019.

DH Hubel and TN Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1):106–154, 1962.

Alexander G Huth, Wendy A De Heer, Thomas L Griffiths, Frédéric E Theunissen, and Jack L Gallant. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532 (7600):453–458, 2016.

Nancy Kanwisher, Josh McDermott, and Marvin M Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of neuroscience*, 17(11): 4302–4311, 1997.

Chaitanya Kapoor, Alex H Williams, and Meenakshi Khosla. Partial soft-matching distance for neural representational comparison with partial unit correspondence. *arXiv preprint arXiv:2602.19331*, 2026.

Meenakshi Khosla and Alex H Williams. Soft matching distance: A metric on neural representations that captures single-neuron tuning. In *Proceedings of unireps: the first workshop on unifying representations in neural models*, pp. 326–341. PMLR, 2024.

David Klindt, Alexander S Ecker, Thomas Euler, and Matthias Bethge. Neural system identification for large populations separating "what" and "where". *Advances in neural information processing systems*, 30, 2017.

Teuvo Kohonen. Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69, 1982.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pp. 3519–3529. PMlR, 2019.

Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:249, 2008.

Richard D Lange, Devin Kwok, Jordan Kyle Matelsky, Xinyue Wang, David Rolnick, and Konrad Kording. Deep networks as paths on the manifold of neural representations. In *Topological, Algebraic and Geometric Learning Workshops 2023*, pp. 102–133. PMLR, 2023.

Hyodong Lee and James J DiCarlo. Topographic deep artificial neural networks (tdanns) predict face selectivity topography in primate inferior temporal (it) cortex. *arXiv preprint arXiv:1909.09847*, 2019.

Yipeng Li, Xieyi Liu, Wanru Li, Jia Yang, Baoqi Gong, Wei Jin, Zhengxin Gong, Kesheng Wang, Jingqiu Luo, Zishuo Zhao, et al. Triple-n dataset: large-scale fmri-guided dense recordings of nonhuman primate neural responses to natural scenes. *Nature Neuroscience*, pp. 1–13, 2026.

Eshed Margalit, Hyodong Lee, Dawn Finzi, James J DiCarlo, Kalanit Grill-Spector, and Daniel LK Yamins. A unifying framework for functional organization in early and higher ventral visual cortex. *Neuron*, 112(14):2435–2451, 2024.

Ian Nauhaus, Kristina J Nielsen, Anita A Disney, and Edward M Callaway. Orthogonal micro-organization of orientation and spatial frequency in primate primary visual cortex. *Nature neuroscience*, 15(12):1683–1690, 2012.

Klaus Obermayer, Gary G Blasdel, and Klaus Schulten. Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A*, 45(10):7568, 1992.

Kenichi Ohki and R Clay Reid. Specificity and randomness in the visual cortex. *Current opinion in neurobiology*, 17(4):401–407, 2007.

Kenichi Ohki, Sooyoung Chung, Yeang H Ch'ng, Prakash Kara, and R Clay Reid. Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, 433 (7026):597–603, 2005.

Kenichi Ohki, Sooyoung Chung, Prakash Kara, Mark Hübener, Tobias Bonhoeffer, and R Clay Reid. Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*, 442(7105):925–928, 2006.

Paolo Papale, Feng Wang, Matthew W Self, and Pieter R Roelfsema. An extensive dataset of spiking activity to reveal the syntax of the ventral stream. *Neuron*, 113(4):539–553, 2025.

Gabriel Peyré and Marco Cuturi. *Computational optimal transport: With applications to data science*. Now Foundations and Trends, 2019.

Richard A Reale and Thomas J Imig. Tonotopic organization in auditory cortex of the cat. *Journal of Comparative Neurology*, 192(2):265–291, 1980.

Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pp. 3982–3992, 2019.

Dario L Ringach, Patrick J Mineault, Elaine Tring, Nicholas D Olivas, Pablo Garcia-Junco-Clemente, and Joshua T Trachtenberg. Spatial clustering of tuning in mouse primary visual cortex. *Nature communications*, 7(1):12270, 2016.

Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral cortex*, 28(9):3095–3114, 2018.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd international conference on learning representations (ICLR 2015)*. Computational and Biological Learning Society, 2015.

Doris Y Tsao, Winrich A Freiwald, Roger BH Tootell, and Margaret S Livingstone. A cortical region consisting entirely of face-selective cells. *Science*, 311(5761):670–674, 2006.

Titouan Vayer, Laetita Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused gromov-wasserstein distance for structured objects: theoretical foundations and mathematical properties. *arXiv preprint arXiv:1811.02834*, 2018.

Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused gromov-wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020.

Alex H Williams, Erin Kunz, Simon Kornblith, and Scott Linderman. Generalized shape metrics on neural representations. *Advances in neural information processing systems*, 34:4738–4750, 2021.