



The topology and geometry of neural representations

Baihan Lin^{a,b,c,d} and Nikolaus Kriegeskorte^{d,e,f,1}

Affiliations are included on p. 11.

Edited by Terrence Sejnowski, Salk Institute for Biological Studies, La Jolla, CA; received October 13, 2023; accepted July 24, 2024

A central question for neuroscience is how to characterize brain representations of perceptual and cognitive content. An ideal characterization should distinguish different functional regions with robustness to noise and idiosyncrasies of individual brains that do not correspond to computational differences. Previous studies have characterized brain representations by their representational geometry, which is defined by the representational dissimilarity matrix (RDM), a summary statistic that abstracts from the roles of individual neurons (or responses channels) and characterizes the discriminability of stimuli. Here, we explore a further step of abstraction: from the geometry to the topology of brain representations. We propose topological representational similarity analysis, an extension of representational similarity analysis that uses a family of geotopological summary statistics that generalizes the RDM to characterize the topology while de-emphasizing the geometry. We evaluate this family of statistics in terms of the sensitivity and specificity for model selection using both simulations and functional MRI (fMRI) data. In the simulations, the ground truth is a data-generating layer representation in a neural network model and the models are the same and other layers in different model instances (trained from different random seeds). In fMRI, the ground truth is a visual area and the models are the same and other areas measured in different subjects. Results show that topology-sensitive characterizations of population codes are robust to noise and interindividual variability and maintain excellent sensitivity to the unique representational signatures of different neural network layers and brain regions.

cognitive neuroscience | computational neuroscience | representational similarity analysis

Geometrical and topological analyses can be applied profitably to the structure and connectivity of brains, on the one hand, and to neural population code representations on the other. Let us first consider the realm of brain structure and connectivity in computational neuroscience and cognitive neuroscience. Geometrical characterizations have been used to study the physical structure of brains [e.g., the geometry of the cortical surface (1–4)] and anatomical connectivity (5–7). Topological and graph-based characterizations have been extensively used in network neuroscience (8–14) to investigate the anatomical connections and functional correlations between brain regions and how structural and functional network topology is related to cognition. This paper is not about either the physical structure or the connectivity of the brain, but about the topology and geometry of neural representations.

In the realm of the neural representations, geometrical characterizations have been used to study the relationships between neural population activity patterns: the representational geometry (15–22) (Fig. 1, *Left* and *Middle*). The representational geometry provides a useful intermediate level of description capturing the information represented in a neuronal population code, while abstracting from the roles of individual measured responses (reflecting neurons or voxels) (23). Considering the representational dissimilarities, rather than the representational patterns, enables direct comparisons of population-code representations between different individuals and species, as well as between brains and computational models. The analysis of representational geometries, known as representational similarity analysis [RSA, (16)], has been successfully applied to understand diverse functions (17), including perception in the visual (15), auditory (24) and other modalities (25) and higher cognitive functions such as abstraction (20, 22), decision-making (26), working memory, social cognition (27–30), and planning (31). Representational geometries can be visualized by arranging stimuli in two or three dimensions, such that their distances approximately reflect the corresponding distances in the high-dimensional neural response space. Representational geometries, captured by the matrix of pairwise distances (the representational dissimilarity matrix, RDM), can also be used as a basis for model comparison (16, 32, 33), an approach

Significance

Characterizing the representations encoded in patterns of neural activity is key to understanding brain computation. Representational similarity analysis uses the geometry of neural population response patterns to compare representations between brains and models. We propose topological representational similarity analysis (tRSA) to also characterize the topology, providing a more robust signature of a region's computational function. tRSA compresses uninformative variation in small and large representational distances to focus on the critical intermediate scale. Analyzing functional MRI data and simulations, we show tRSA matches standard representational similarity analysis at identifying corresponding brain regions and model layers, while focusing on the core topological features distinguishing representations. Considering topology alongside geometry enables a comprehensive understanding of the information represented in population codes.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: nk2765@columbia.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2317881121/-/DCSupplemental>.

Published October 7, 2024.

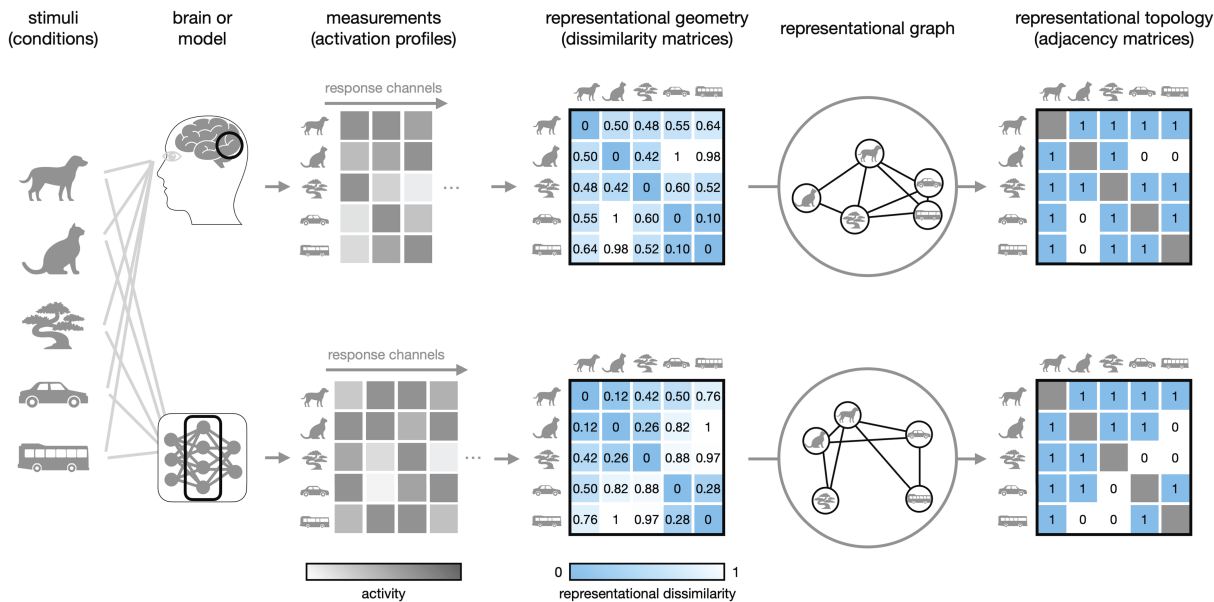


Fig. 1. Comparing representations between brains and models. To understand the degree to which a computational model can account for the cognitive process of a certain brain region, the same set of stimuli is presented to both the model and the biological system. The response patterns across measured response channels (e.g., neurons or voxels) are then characterized by a summary statistic, the representational dissimilarity matrix (RDM, center), which defines the metric configuration of the stimuli in the neural population response space. However, the metric configuration can be sensitive to measurement noise and idiosyncrasies of individual brains that do not reflect computational function. An alternative summary statistic that captures the topology would be the adjacency matrix (Right), which defines the unweighted graph of neighborhood relationships in the population response space. This summary statistic promises to be more robust to noise and idiosyncrasies, but may discard too much information. Considering the geometry (RDM) and topology (adjacency matrix) as extremes of a continuum suggests that it may be possible to get the best of both (Fig. 2).

that has enabled researchers to adjudicate among competing models of brain representations (34–37).

When we investigate the representational geometry by considering distances among neural activity patterns, we abstract from the roles of individual neurons. The representational topology provides a further step of abstraction. We may care less about the precise distances among the points in the high-dimensional response space that define the geometry than about the way the points hang together in what is sometimes called the neural manifold (18, 22, 38). We may hypothesize, for example, that the overall geometry of the representation in a given cortical area or layer of a neural network model may vary across individual people or instances of a neural network model trained from different random seeds (39). If the corresponding cortical areas in two people or the corresponding layers in two model instances served the same computational purpose, however, we may expect that stimuli that are neighbors in one individual's (or model instance's) representation remain neighbors in the other individual's (or instance's) representation. A graph of representational neighborhood relationships can be obtained by thresholding the distance matrix (Fig. 1, Right). The thresholding operation is well motivated when we care only about whether two points are in the same neighborhood or not. If they are in the same neighborhood, we consider them related and do not care whether they are close or very close. If they are not in the same neighborhood, we consider them unrelated and do not care whether they are very far apart or merely far enough not to count as neighbors. The neighborhood graph characterizes the representational topology.

An important question is to what extent the further step of abstraction involved in going from the distance matrix (geometry) to the neighbor graph (topology) is desirable or undesirable. It could be desirable for providing a more robust reduced signature of a region's computational function. However, it could be undesirable if it removes geometrical information important for discerning

regions that implement distinct computational functions. Here, we address this question empirically, using human functional MRI data and simulations based on neural network models.

Topological data analysis techniques (40) such as the persistent homology (41) and the Mapper algorithm (42) are popular in many fields of biology (43–48) and have also been used to directly analyze the representational space of the population activity (49). For instance, a study has discovered that the structure of spontaneous and evoked activity patterns in V1 can be mapped onto a manifold that has the topology of a sphere, whose two dimensions may reflect orientation and spatial frequency (50), with the population response selective to the extremes of spatial frequency mapped toward the two poles of the sphere. Similarly, a recent study applied topological analysis techniques to study the population activity of grid cells, which are thought to be involved in spatial navigation and orientation (51). This study found that the population activity of grid cells has a toroidal topology, meaning that the manifold wraps around like a donut, whose two surface dimensions correspond to the 2d space navigated, implementing a cyclic representation. These examples illustrate the power of topological data analysis techniques to reveal the structure of the neighborhood graph of neural population representations.

These inspiring studies notwithstanding, topological characterizations are more widely used in network neuroscience and only beginning to impact investigations of the relationships among neural population activity patterns. Here, we build on the early topological analyses of neural population activity patterns (38, 47, 50, 51) and introduce a family of summary statistics that can characterize the geometry as well as the topology of neural activity patterns. These geotopological summary statistics enable researchers to calibrate the geometric and topological sensitivity of the analysis, so as to define a good signature of the computational role of each brain region. The representational signatures can then be used not only for visualization of the

representational geometry and topology, but also as a basis for formal inferential model comparison in the framework of RSA (33), where our geotopological summary statistics can replace the RDM, which characterizes the geometry.

Consider the example of visual perception (Fig. 1). We begin by measuring the brain-activity patterns elicited by each of a set of stimuli in a brain region or computational model. By estimating the distances among the stimulus representations (with full metric information), we can gain insights into distinctions the brain region or model layer emphasizes. Metric distance estimates promise detailed geometrical information, but are sensitive to noise and individual idiosyncrasies. Thresholding the distances provides a graph with binary edges, which captures how the neural manifold hangs together and also promises to be more resilient to nuisance variation. The methods we introduce here share a focus on neighborhood relationships with popular visualization techniques like Isomap (52), locally linear embedding (53), *t*-distributed stochastic neighbor embedding (54), and Uniform Manifold Approximation and Projection (55). However, while these techniques aim to visualize a single representation, our aims are to characterize multiple representations in models and brains, quantify their similarity, and statistically compare models in terms of their ability to account for the topology and geometry of brain representations. To combine the benefits of detailed metric information and a binary edge description, we seek to define a representational graph that captures aspects of both the representational geometry and topology using weighted edges.

To integrate geometric (distance-based) and topological (graph-based) characteristics, we define the edge weights of the graph by a nonlinear monotonic transformation of the distances that 1) emphasizes the distinction between small and large distances, 2) compresses very small distances, thus disregarding the distinctions among them, 3) represents a continuum of intermediate distances, and 4) compresses very large distances. An example of the effectiveness of this transformation is an adaptive generalization of distance correlation based on proposed transformed distances, which provides a dependence measure with robust sensitivity to geometric and topological characteristics (56, 57).

There are two motivations for defining the edge weights as a nonlinear monotonic transform of the distances, one theoretical and one data-analytical. From a theoretical perspective, differences among very large representational distances may not provide the most useful signature of the computational function of a brain region. It is the local geometry that determines which stimuli the representation renders indiscriminable, which it discriminates, but places together in a cluster, and which it places in different neighborhoods. The global geometry of the clusters (whether two stimuli are far or very far from each other in the representational space) may be less relevant to computation for two reasons. First, once two stimuli are perfectly discriminable, moving them even further apart does not improve discriminability. Second, in a high-dimensional space, a set of randomly placed clusters will tend to afford linear separability of arbitrary dichotomies among the clusters (58–60) independent of the exact global geometry.

Like in a storage room, related things may need to be placed together in a representational space and unrelated things in different locations. The requirement of colocalization strongly constrains the local geometry because there is only one direction toward a given location. The requirement that two things be far from each other, by contrast, only weakly constrains the global geometry, because there are many directions away from a given location, especially in a high-dimensional space. This argument suggests the hypothesis that variations among large distances are

idiosyncratic to an individual brain or model instance (39). If this hypothesis were true, then compressing this variation may help us focus on more functionally relevant features of the representation that are less variable across individuals and model instances.

From a data-analytical perspective, conversely, very small distances may be unreliable given the various noise sources that affect the measurements. Compressing small distances, thus, promises to reduce the influence of measurement noise on visualizations and inferential results. Compressing small and large distances is achieved by thresholding of the distances. However, thresholding may be too aggressive in that it completely removes all continuous information reflecting the geometry. As illustrated in Fig. 2, on one end, we have the distance matrix with the full metric information, and on the other, we simply have an adjacency matrix telling us whether two stimuli are neighbors in the representational space or not. To get the best of both worlds, it seems attractive to focus our sensitivity on a particular intermediate range of distances, so as to maintain reliable geometric information, while reducing the influence of noise (by compressing variation among small distances) and the influence of individual idiosyncrasies (by compressing variation among large distances). We show that this can be accomplished by a monotonic transform of the representational distances and that the resulting geotopological representational summary statistics robustly reveal the functional distinctions among human brain regions and Deep neural network (DNN) layers. We introduce a family of geotopological summary statistics that generalizes the RDM and provides a basis for topological RSA (tRSA), a generalization of RSA that balances sensitivity to the topology and geometry of neural representations.

Materials and Methods

Nonlinear Monotonic Transforms of Representational Dissimilarities Provide a Family of Geotopological Descriptors. Topological RSA builds on the literature on topological methods [e.g., persistent homology (61) and TDA mapping (62)]. In order to suppress noise, we would like to find a lower threshold *l* below which we consider stimuli as colocalized (i.e., the distance is 0). In order to abstract from idiosyncrasies of individual brains and highlight the representational properties that are key to their computational function, we would like to find an optimal upper distance threshold *u* above which we consider stimuli maximally distinct (i.e., we do not consider differences between larger distances meaningful). Two stimuli whose distance is larger than *u* are disconnected in the graph capturing the topology.

Between the two thresholds, we place a continuous transition so as to retain geometrical sensitivity in the range where it is meaningful (Fig. 3A). For simplicity, we propose a piecewise linear function as the monotonic distance transform. Note, however, that alternative monotonic transforms, such as the logistic function or a cosine transition, could be applied here. Given an original distance d_{ij} between the two neural signatures of two objects in the representational space, the piecewise linear geotopological (GT) transform is defined as

$$GT_{l,u}(d_{ij}) = \begin{cases} 0, & \text{if } d_{ij} \leq l \\ \frac{d_{ij}-l}{u-l}, & \text{if } l < d_{ij} < u \\ 1, & \text{if } u \leq d_{ij} \end{cases} \quad [1]$$

By varying the lower bound *l* and upper bound *u*, we obtain a family of GT transforms (Fig. 3B). Each member of this family transforms the original dissimilarity matrix (RDM) into a representational geotopological matrix (RGTM), which provides a multivariate summary statistic with particular degrees of topological and geometrical sensitivity (Fig. 3C). The RGTM replaces the RDM in topological RSA. Note that the RDM is itself a member of the family, where $l = 0$ and *u* are set to the maximum (*Upper Left* corner in Fig. 3B–D). The RGTM, thus, generalizes the RDM.

When applying topological RSA to neural data to understand a brain representation, we can benefit from considering not only the RDM, but also other members of the RGTM family to gain an understanding of the geometrical and topological features of a brain representation. For the particular purpose of model selection, we aim to choose thresholds l and u , so as to best recover the true (data-generating) model representation if it is among our models, or the best approximation among the models we are testing.

The Geotopological Descriptor Family Captures Both Geometric and Multiscale Topological Information. For interpretation of the family of GT transforms, we can consider the two diagonal dimensions of the triangular set spanned by l and u (Fig. 3D): From the *Upper Left* corner (RDM) toward the *Lower Right*, we go from geometrical to topological sensitivity, approaching the thresholding transforms ($l = u$) along the diagonal. From the *Lower Left* to the *Upper Right*, we go from local to global distance sensitivity.

For instance, the original RDM is in the *Top Left* corner ($l = 0, u = 1$), and around it, we have a region of RGTM that is *geometry-sensitive* where l is small and u is big. On the *Bottom Left* corner, where both l and u are small, we have *local extractors*, which are RGTM that are sensitive to whether or not two items are very close neighbors. On the *Upper Right* corner, where both l and u are big, we have *global extractors*, which are RGTM that are sensitive to whether or not two items are on opposite ends of the ensemble. If we go along the diagonal line from *Lower Left* to *Upper Right*, we have a belt zone of RGTM which are close to binary (i.e., l and u are close to each other).

Formally, we require $l < u$, so the diagonal, where the transform is a simple thresholding function, is excluded from the family. Allowing $l = u$, so as to formally include simple thresholding functions, is possible but would complicate Eq. 1. In either variant, the GT transform approaches a binary thresholding

function for points approaching the diagonal. The thresholding functions along the diagonal relate our approach here to the mathematical filtration process used to reveal persistent homology in topological data analysis.

By exploring the choices of l and u , we can identify the GT transforms that best enable us to match functionally corresponding cortical areas between different individuals. Similarly, we can generate data (with simulated measurement noise) from a layer in a deep neural network model, and determine which choices for l and u best enable us to identify the data-generating layer when using a range of layers from other instances of the DNN architecture (trained from different random seeds) as the models in analyses.

One possibility is that the ideal setting is $l = 0, u = \max$, i.e., the original RDM, which characterizes the geometry. Other settings of l and u remove information about the geometry. Whether removing information by choosing a larger l or a smaller u helps or hurts depends on the relative extent to which it reduces signal and noise (nuisance variation) in the context of a particular data-analytical objective. If a topology-sensitive summary statistic reduced the variation caused by measurement noise and individual idiosyncrasies (i.e., nuisance variation) more than variation reflecting computational roles of different representations, then inferential comparisons of deep neural network models would benefit from topological RSA.

Geodesic Distances Provide an Alternative Geotopological Descriptor.

In addition to RGTM, we propose the use of geodesic distances in the representation. Let us first consider the theoretical notion of a geodesic and then the practical analyses it motivates. Theoretically, the “representation” of a population of possible stimuli can be defined as the set of response patterns the stimuli elicit. If the stimulus population is a continuous set, we may hypothesize that so is the set of corresponding neural response patterns in our brain region

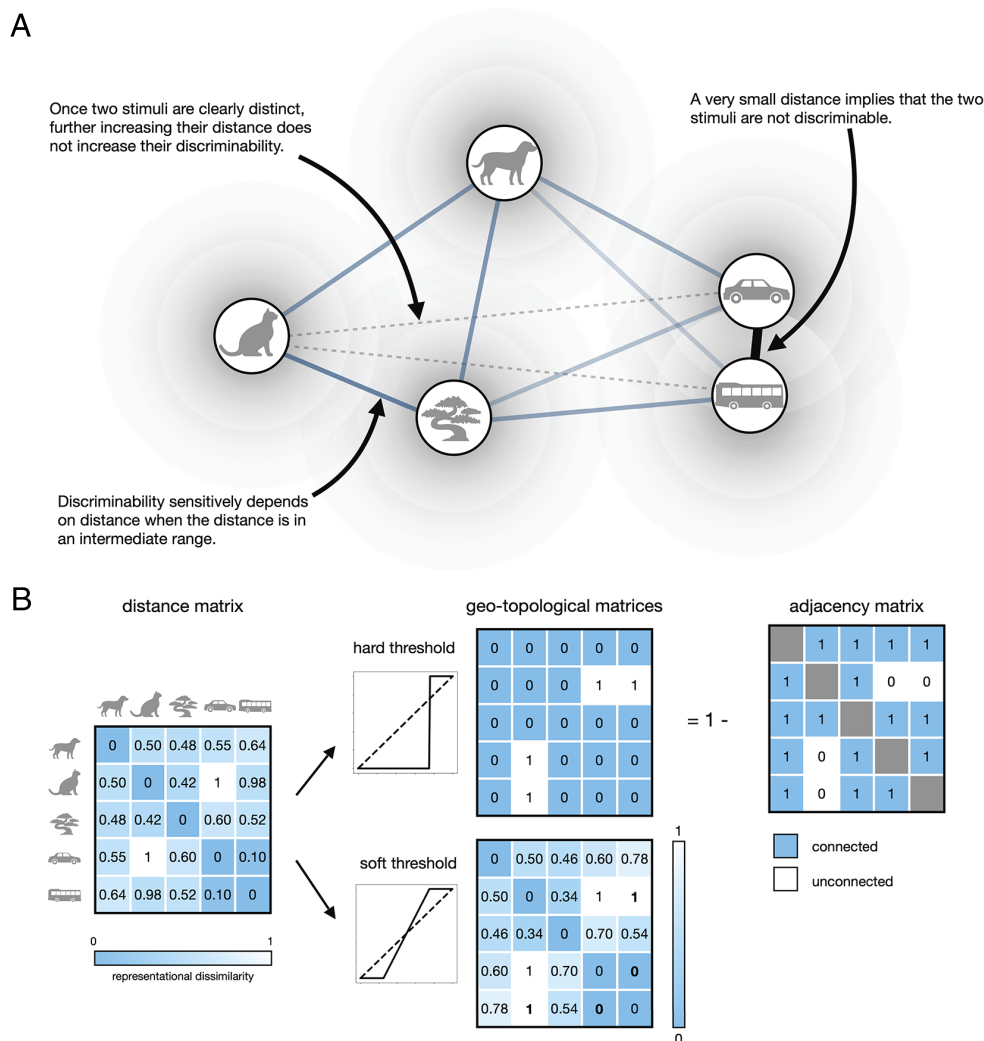


Fig. 2. Intuition of the geotopological transform of distances. (A) Consider the five visual stimuli from Fig. 1, whose representation in a visual cortical area can be characterized by its representational geometry. If the response patterns are all affected by isotropic noise (or the noise has been whitened by a transform), then the Euclidean distances monotonically reflect the discriminabilities. Variation among large distances, however, is not associated with great differences in discriminability, because all pairs of well-separated stimuli are nearly perfectly discriminable (dashed lines). Similarly, variation among very small distances is not associated with great differences in discriminability, because all pairs of neighboring stimuli are indiscriminable (thick black edge). This suggests that variation among small distances and variation among large distances can be suppressed in favor of emphasizing the transition from small to large distances. (B) To emphasize the transition from small to large distances while suppressing variation among small distances and variation among large distances, we can threshold the distances, such that small distances are pushed to zero and large distances are pushed to the maximum. We can either use a hard threshold (*Upper* row) or a soft threshold (*Lower* row). A hard threshold yields a binary matrix whose complement is the adjacency matrix of a graph that connects neighboring stimuli. A soft threshold creates a continuous transition, reflecting the graded increase in discriminability as the distance grows, and defines a weighted graph, where the weights reflect distances, but the pairs of stimuli that are furthest from each other are not directly connected (dashed lines in A).

of interest. This set of neural response patterns is often referred to as the neural manifold. (Note, however, that the set of response patterns would need to be locally homeomorphic to a Euclidean space to conform to the definition of manifold. Whether this is the case for a particular neural population is an empirical question.) A geodesic distance between two stimulus representations is the length of a shortest path traversing the representational set from one to the other. Unless the straight line between the two points is a subset of the representation, we will traverse a longer, curved shortest path through the representational set. The length of the geodesic path then will be larger than the Euclidean distance.

In practice, we will have data for a finite sample from the population of stimuli. To estimate the geodesic distance on the manifold, we can measure geodesics in the discrete graph characterizing our representation. In a graph, a geodesic distance is the length of the shortest path between two nodes. We define a representational graph for each member of the RGTM family. For each node i , edges exist only to other nodes j with dissimilarity $d_{ij} < u$. The edge weights are defined by the transform $GT(l, u)$ and are interpreted as distances between nodes. Note that edge weights, thus, can be zero. Zero-distance edges can be motivated as a correction of small positive distance estimates resulting from off-manifold noise displacements of patterns whose locations on the manifold are not significantly distinct. In order to maintain direct comparability of the edge-weight matrices of different brain and model representations, we do not collapse nodes with zero distance, but maintain one node for each stimulus.

The geodesic distance between two nodes is defined as the length of the shortest path that leads from one node to the other, where the length of a path is the sum of the internode distances (i.e., the edge weights) along it. The geodesic distance is infinite if there is no path connecting the two nodes through the edges. The shortest paths for all pairs of nodes (each corresponding to a stimulus) can be found using Dijkstra's algorithm (63). The result is a stimulus-by-stimulus matrix, which we refer to as the representational geodesic-distance matrix (RGDM). The RGDM can be used in place of the RGTM or the RDM for our visualization and model-comparative inference procedures. As an alternative to an RGTM-based graph, we could use a binary graph (edge weights $\in \{0, 1\}$)

to compute the RGDM. For example, we could use the binary graph in which each node is connected to its k nearest neighbors or the graph containing connections for the k smallest representational dissimilarities in the RDM. Our analyses here, however, use RGDMs computed from graphs with continuous internode distances, based on members of the RGTM family as described above.

Leave-One-Out Evaluation on Human fMRI Data and DNN Models Can Quantitatively Evaluate the Region Identification Power of the Summary Statistics.

We would like to quantitatively evaluate the power of topological RSA in the context of model selection, where each model predicts a representational geometry. We therefore consider cases, where the ground-truth model is known. This enables us to objectively evaluate the impact of different choices of l and u on the accuracy of model selection. If conventional RSA were optimal, then setting $l = 0$ and $u = \max$ would be the best settings. If other settings afforded equal or better accuracy for model selection, then topological signatures would deserve consideration in future studies applying RSA to adjudicate even among models that predict not just representational topologies, but full representational geometries.

Evaluating topological RSA's brain-region-identification accuracy (fMRI). The fMRI evaluation was performed on preexisting data from a human fMRI experiment (64, 65), in which 24 subjects were presented with 62 colored images depicting faces, objects, and places. We use 8 regions of interests here: the primary visual cortex (V1), the secondary visual cortex (V2), the extrastriate visual cortex (V3), the lateral occipital complex, the occipital face area, the fusiform face area, the parahippocampal place area, and the anterior temporal lobe.

We investigate the brain-region-identification accuracy (RIA), where each brain region is considered a model. The region labels provide the ground truth: For data from each brain region in a held-out subject, we would like to identify which region the data came from on the basis of the data for all the regions from the other subjects. We therefore perform leave-one-subject-out RIA evaluation. First, we randomly sample 10 sets of l 's and u 's in each of the five interpretable RGTM zones: topology-sensitive (TS), geometry-sensitive (GS), local extractor (LE), global extractor (GE), and intermediate (I). The 10 samples of l and u for

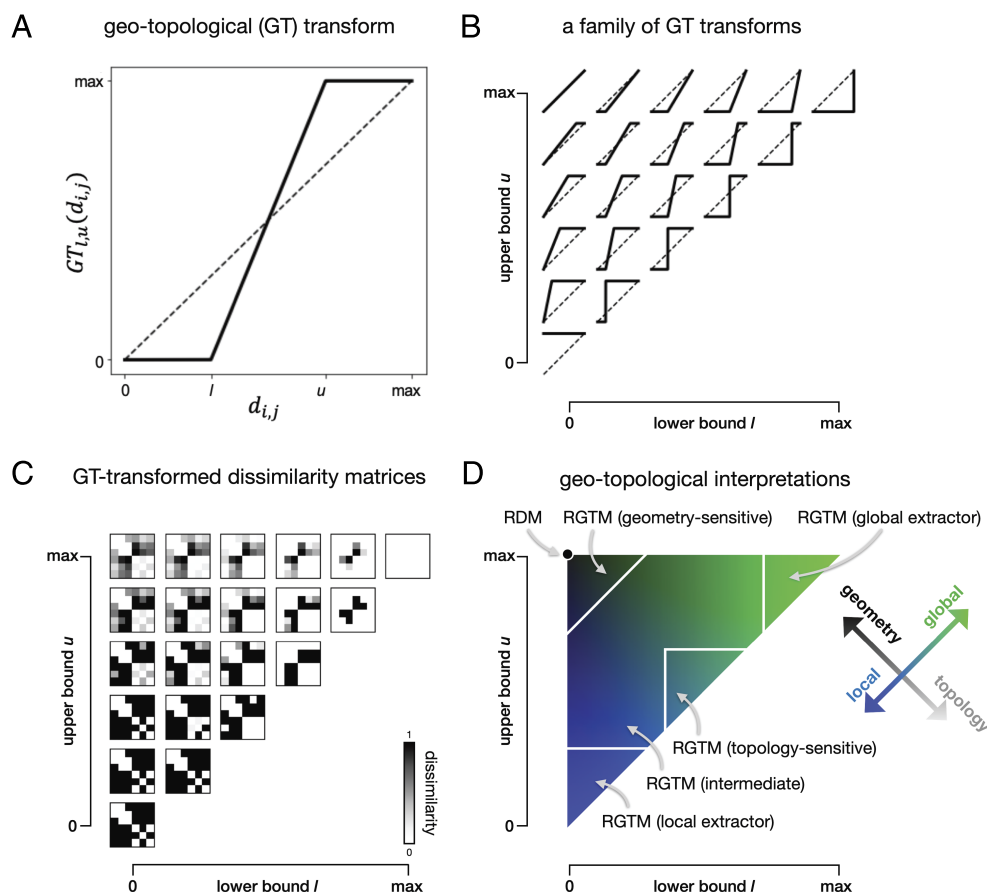


Fig. 3. A family of geotopological transforms of the RDM. (A) The geotopological (GT) transform is formulated as a linear piecewise function, such that any distances smaller than a *Lower* bound l will be mapped to zero, and any distances bigger than an *Upper* bound u will be mapped to 1. Between l and u , the transition is linear. (B) By varying the thresholds l and u , we select among a family of GT transforms. (C) By applying different GT transforms to the RDM, we obtain so-called representational geotopological matrices (RGTM). (D) To interpret the way the GT transforms reflect geometric and topological properties of the representation, we group the family members in different zones of the plane spanned by l and u . The closer a GT transform is to the *Upper Left* corner ($l = 0, u = \max$), the more similar the RGTM is to the RDM. As we approach the diagonal line ($l = u$), the GT transform approaches a hard threshold, emphasizing the topology rather than the geometry. As we move diagonally from the *Bottom Left* to the *Upper Right*, the RGTM goes from emphasizing the local neighbor relationships among the stimuli to emphasizing the global structure of the representation.

each region define 10 different GT transforms and provide an estimate of the RIA, averaged across regions and subjects, that reflects the performance at region identification of different zones of the family of RGTM.

For each region in a held-out subject, we assign the region label of the average RGTM from the other subjects that is closest (in terms of Euclidean distance) to the RGTM being identified. Each subject is held out once in a full cross-validation cycle and the RIA is the average identification accuracy.

In order to inferentially compare the different zones of the RGTM family, we perform frequentist comparisons. We would like to consider a difference in performance as significant if we expect it to generalize to experiments performed with different samples of subjects and stimuli drawn at random from the same populations of subjects and stimuli. We therefore perform a 2-factor bootstrap procedure, resampling both subjects and stimuli simultaneously (32, 33). The SD of the RIA estimates across 1,000 bootstrap samples of subjects and stimuli serves as our estimate of the SE of the RIA estimate. Two-sided t tests are then applied to assess the significance of differences between RIA estimates for different choices of l and u . The degrees of freedom in this approach correspond to the smaller factor, which is the number of subjects (24) in this case.

Evaluating topological RSA's layer identification accuracy (DNNS). The DNN evaluation was performed on a convolutional neural network architecture (66) called the All Convolutional Neural Net (67). We investigate the layer identification accuracy (LIA), where each layer is considered a candidate brain-computational model. We perform leave-one-instance-out (LOIO) LIA evaluation. We trained 10 model instances, starting from 10 different random seeds, of the All-CNN-C network architecture (67), a 9-layer fully convolutional network that exhibits state-of-the-art performance on a well-known small object-classification benchmark task [Canadian Institute For Advanced Research (CIFAR-10) (68)].

We used the same numbers of feature maps (96, 96, 96, 192, 192, 192, 192, 10) and kernel dimensions (3, 3, 3, 3, 3, 3, 1, 1) as in the original paper. The training of All-CNN-C network instances involved 350 epochs using the ADAM optimizer with a momentum value of 0.9 and a batch size of 128. A preliminary learning rate of 0.01 was employed, along with an L2 regularization coefficient of 10^{-5} and gradient norm-clipping value of 500. Following ref. 39, we trained the DNNs on the complete CIFAR-10 image dataset (both training and test sets), which comprises 10 distinct object categories, each represented by 5,000 training and 1,000 test images, implemented with TensorFlow (version 1.3.0) and Python 3.5.4.

Like different individual subjects, these instances differ in their detailed connectivity, but perform the recognition task at similar levels of accuracy (39). In addition to evaluating the LIA across instances for different choices of l and u , we study the effect on the LIA of injecting Gaussian noise of a variety of variances σ^2 into the dissimilarity estimates.

The DNN-simulation-based evaluations follow the same procedures as the human-fMRI-based evaluations: The neural networks were presented with the same set of 62 object images to define the representational geometries. We perform LOIO LIA evaluation. Each layer in a held-out instance is identified as the layer whose average RGTM across the other instances is closest (in terms of Euclidean distance) to the RGTM being identified. Each instance is held out once in a full cross-validation cycle and the LIA is the average identification accuracy. The inference, likewise, employs the same 2-factor (instance and stimulus) bootstrap method.

Model-Comparative Statistical Inference for tRSA. tRSA can use the well-developed inferential techniques of RSA for comparison between representational models (16, 33, 69). The nonparametric inference methods of RSA3 (33) can simply use the geotopological statistics (RGTM and RGDM) in place of RDMs. As in conventional RSA, the representational similarity of brain regions and model layers can be quantified using various comparators, such as cosine similarity or correlation, but applied to geotopological statistics rather than RDMs.

Here we used Euclidean distance as the comparator for the geotopological summary statistics. The l and u are defined as quantiles (expressed as percentiles or ranks) relative to the set of dissimilarities in a given RDM. Defining l and u as quantiles enables matching choices for model and brain representations whose dissimilarities may have different magnitudes and may lack a common unit that would render them commensurable. In addition to defining l and u as quantiles, we use the ranks within each RDM to define the dissimilarities entering the GT transform. This has the benefit that the resulting RGTM have

identical distributions of values. In this scenario, the squared Euclidean distance as a comparator of representational summary statistics is proportional to the Pearson correlation distance and to the cosine distance. These comparators thus would have yielded identical model-selection results, rendering the analyses relevant to the most common choices of RDM comparator in RSA.

Another popular choice for the RDM comparator is a rank correlation coefficient, such as Kendall's τ_a (32) or the more computationally efficient Spearman-type coefficient ρ_a (33). In the present study, the RDMs are rank-transformed before computing the RGTM. Our comparators therefore benefit from the robustness afforded by the rank transform, obviating the need for another rank-transform at the level of the RGTM, as would happen if we chose a rank correlation coefficient as the comparator. Using a rank correlation coefficient is closely related (but not mathematically equivalent) to the present analyses.

Results

Proof of Concept: Geotopological Descriptors Can Discern Topological and Not Just Geometric Distinctions. As a proof of concept, consider the geometric and topological similarities among the four hypothetical representations in Fig. 4. We call the four representations the “flat 8,” the “bent 8,” the “untangled flat 8” and the “untangled bent 8.” Imagine an experiment in which the idealized continuous representational set is sampled using of 40 stimuli (balls in Fig. 4A). The flat 8 and the bent 8 share the self-intersection, which creates two holes, rendering the shapes topologically similar, although the bending of the latter greatly changes the geometry. Similarly, the untangled flat 8 and the untangled bent 8 are topologically similar (in that neither has a self-intersection and both have one hole) and geometrically dissimilar because of the bending. By contrast, the flat 8 and the untangled flat 8 are geometrically similar (both flat and similar in their RDM), and topologically distinct (one versus two holes). Likewise, the bent 8 and untangled bent 8 are geometrically similar (both bent) and topologically distinct (double arrows in Fig. 4A).

The different summary statistics reflect the topology and geometry of these hypothetical neural representations to different degrees (Fig. 4B). As expected, the RDM reflects the geometric distinctions but is not very sensitive to the topological distinctions, which are implemented here through minimal metric displacements that determine whether or not there is a self-intersection. The multidimensional scaling (MDS) arrangement of the four RDMs (Fig. 4C) shows the dominance of the flat versus bent distinction when characterizing the representations with RDMs.

The RGTM were obtained using small values for both l and u , yielding almost binary matrices with local topological sensitivity, which reveal whether or not there is a self-intersection. The MDS shows that the RGTM here balances sensitivity to the topology and the geometry of the representation. Note the prominent blue “eyes,” which reflect the self-intersection where the contour of the 8 crosses itself. The blue “eyes” are present for the two-hole representations with self-intersection (Left two representations in Fig. 4B), but absent for the untangled one-hole representations.

Finally, the RGDMs are even more sensitive to the topological distinctions. The RGDMs were defined on the basis of the graphs of the RGTM in the row above. The shortest-path lengths reflect the existence of “shortcut” routes through the self-intersection. For many pairs of stimuli, the shortcut (when available) provides a shorter path than going the straight route around the 8. The small geometrical distortion causing the self-intersection is therefore reflected across a large portion of the RGDM. The MDS plots show that the RGTM and RGDMs can effectively characterize both the geometric and topological similarities, whereas the RDMs mostly reflect the geometry of the representation.

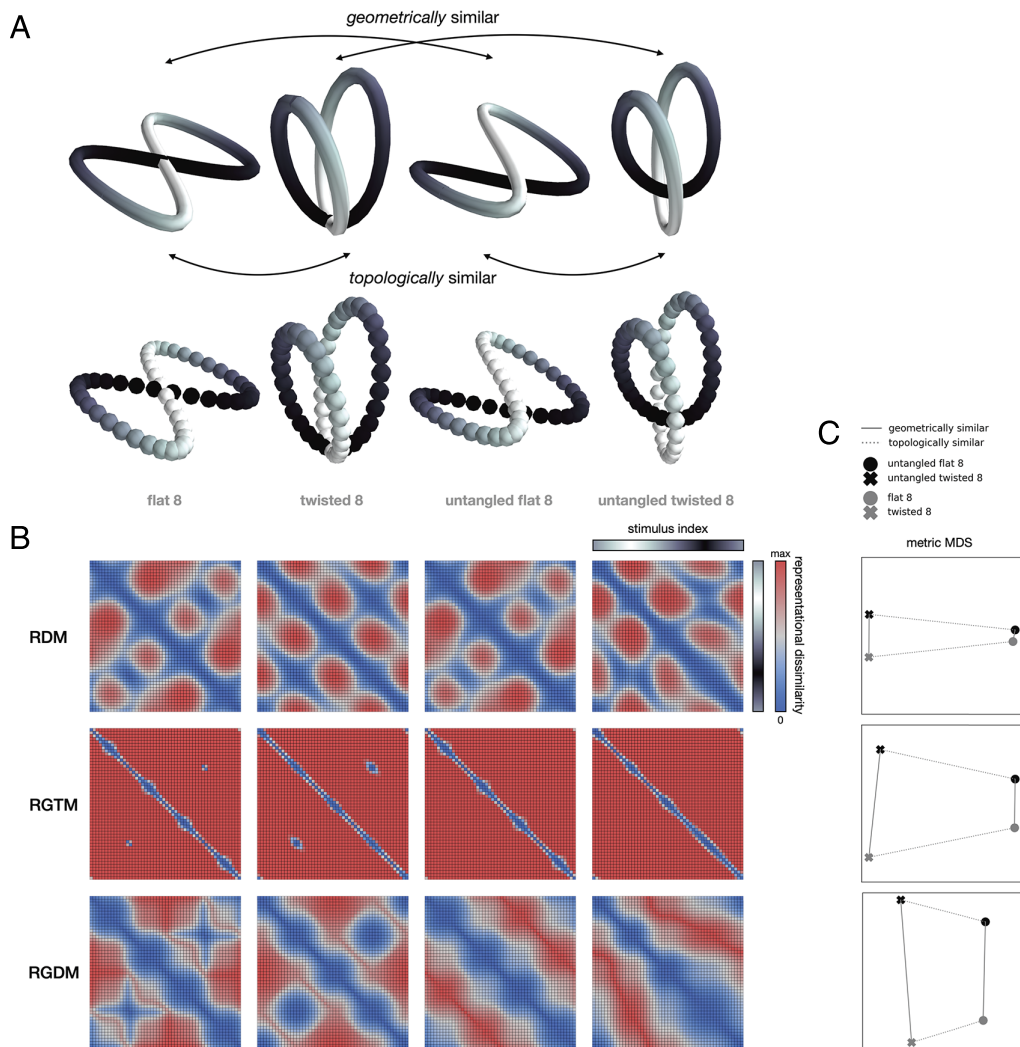


Fig. 4. The geometric and topological similarities between hypothetical neural representations (proof of concept). (A) We consider four hypothetical representations of 40 stimuli (balls in Lower row of A). The response patterns are sampled from idealized continuous sets of neural response patterns (Top row in A). Two of these continuous sets are manifolds (the untangled shapes) and the other two are not (the tangled shapes, where the set self-intersects, forming a neighborhood that is not homeomorphic to a Euclidean space). From Left to Right, we label the four representations the flat 8, the bent 8, the untangled flat 8, and the untangled bent 8. The flat 8 and the untangled flat 8 are geometrically similar, while being topologically dissimilar (with the former self-intersecting). Their bent versions, as well, are geometrically similar and topologically dissimilar. The flat 8 and the bent 8, on the other hand, are topologically similar (with the self-intersection creating two holes) and geometrically dissimilar (because the twisting substantially changes the metric geometry). Their untangled versions, as well, are topologically similar and geometrically dissimilar. (B) The RDMs (Euclidean distance, Top row) do a good job characterizing the geometric relationships but are insensitive to the topological relationships. The RGTM (Middle row) are more sensitive to topology. To achieve local topological sensitivity, we chose $l = 0$ and $u = 0.075$, revealing the self-intersection in the two leftmost representations. The RGDMs (Bottom row) are exquisitely sensitive to the topological relationships, while deemphasizing geometrical relationships between the representations. Note the “eyes” in the two leftmost

RGTM, representing the self-intersections. Each RGDM captures the lengths of the shortest paths in the graph of the RGTM shown Above it. The RGDMs more prominently reflect the topology as the “shortcut” paths enabled by the self-intersection affect shortest-path lengths for a broad swath of stimuli. (C) Multidimensional scaling (MDS) on the four representations shows the pairwise similarities among the four matrices of each row, confirming the increasing sensitivity to topological differences as we go from RDM to RGTM and on to RGDM.

Geotopological Descriptors Reveal What Aspects of Geometry and Topology Enable Accurate Identification of Brain Regions Across Subjects (fMRI).

Using human fMRI data, we quantitatively evaluated the performance of the RGTM summary statistics (including the RDM as a special case) at revealing the correspondences among ventral-stream cortical regions across subjects. A summary statistic will succeed in this evaluation, if it is robust to noise and interindividual variability while maintaining sensitivity to differences between different cortical areas.

We grouped the members of the RGTM family into interpretable zones in Fig. 3D. Results for each zone are shown in Fig. 5A). We plotted the RIA as a heatmap across choices for l and u (Fig. 5B). The highest RIA was achieved for $l = 0.40$ and $u = 0.65$ (red square in Fig. 5B). However, other settings yielded similar RIA. Results suggest that l and u should not be both low or both high. The RDM is an effective summary statistic. However, it contains information not needed for region identification. This result suggests that compressing less informative distance variation is possible without a reduction in model-selection accuracy. In this dataset, we did not find that the benefits of noise reduction significantly outweighed the loss of information.

We inferentially compared RIA between different zones within the RGTM family using simultaneous bootstrapping of both subjects and stimuli [2-factor bootstrap, (33)] for frequentist pairwise comparisons. The 2-factor approach serves to test for differences expected to generalize across random samples of subjects and stimuli drawn from the same populations. The summary statistics in the TS and GS zones did not perform significantly differently ($P = 0.434$). The TS, GS, and intermediate RGTM all performed better than those in the LE zone ($P = 4.776e-03$, $P = 4.972e-03$, and $P = 4.454e-02$, respectively). We found no other significant differences. Overall, geometry-sensitive and topology-sensitive summary statistics performed similarly, but the latter suffer when the upper bound is in the lower third of the distribution.

These findings suggest that a substantial portion of the smallest and largest dissimilarities can be compressed without reducing the RIA. Using the RGTM to focus on the intermediate range of dissimilarity variation (perhaps the middle third of the distribution of dissimilarities) while compressing the lower and upper third of the dissimilarities yields optimal model selection in this dataset.

Geotopological Descriptors Reveal What Aspects of Geometry and Topology Enable Accurate Identification of DNN Layers Across Instances. DNNs have emerged as important models of vision and brain-information processing in recent years (70–74). DNNs can learn nonlinear representational transformations through feedforward and recurrent processing, enabling them to capture the computations underlying task performance. When a neural network architecture is trained repeatedly, starting from different random connection weights, the resulting trained model instances have distinct parameters and distinct detailed unit tuning, despite performing the same task roughly equally well. Just like individual humans differ, thus, so do instances of DNNs trained from different random seeds (39).

Analogously to the brain-region identification analysis in the previous section, we investigated the usefulness of RGTM in identifying which layer of a neural network model a given RGT was computed from, given RGTs for all layers in other instances of the architecture. Like the human fMRI analysis, this analysis informs us about the best choice of representational summary statistic when performing DNN model comparisons using human data. We are able to objectively evaluate the summary statistics, because we know the ground-truth data-generating model. The DNN simulation analyses are complementary to the human fMRI analyses: On the one hand, they afford complete knowledge of the computational mechanisms (an advantage). On the other hand, they are less realistic with respect to the biology and structure of the noise in the data used as a basis for model comparison (a disadvantage).

Our layer-identification analysis reveals to what extent different RGTs (including geometry-sensitive and topology-sensitive members of the family of GT transforms) abstract

from nuisance variation across the same layer in different model instances, while capturing computationally meaningful variation between different layers (which are thought to play distinct computational roles).

As for the human brain regions, we quantitatively evaluated the performance of the summary statistics from the RGTM family (including the RDM). Results are shown in Fig. 6, grouped by noise level (Fig. 3D). We present results for two types of noise: Gaussian noise σ (added to the dissimilarity estimates, Fig. 6) and Bernoulli noise ϵ (applied during training as dropout rates, SI Appendix, Fig. S1).

Statistical comparisons of LIA for GT-transform zones showed that geometry-sensitive statistics outperformed topology-sensitive statistics at higher noise levels ($\sigma \geq 0.2$). The heatmaps showing LIA as a function of l and u in RGTs (Fig. 6B) likewise show that as the noise level increases, the optimal statistics (red frames) go from more topology-sensitive to more geometry-sensitive. The LIA heatmaps based on the RGDs show a prominent fall-off of accuracy with l (horizontal axis, as in Fig. 3) and little dependence on u : As l grows, more and more pairs of neighboring points collapse to 0 distance in the RGT, and thus more and more shortest-path lengths between points connected by paths of 0-distance edges collapse to 0 in the RGD. The weak dependence of LIA on u reflects the fact that more separated points can be connected either through a smaller number of larger steps (when u is large) or a larger number of smaller steps (when u is small), rendering the RGD less sensitive than the RGT to u . Overall, in RGTs, there is no significant LIA difference between the geometry-sensitive and the topology-sensitive zone when $\sigma < 0.2$. Selecting a summary statistic that emphasizes only local neighborhood (local extractor) can be significantly worse

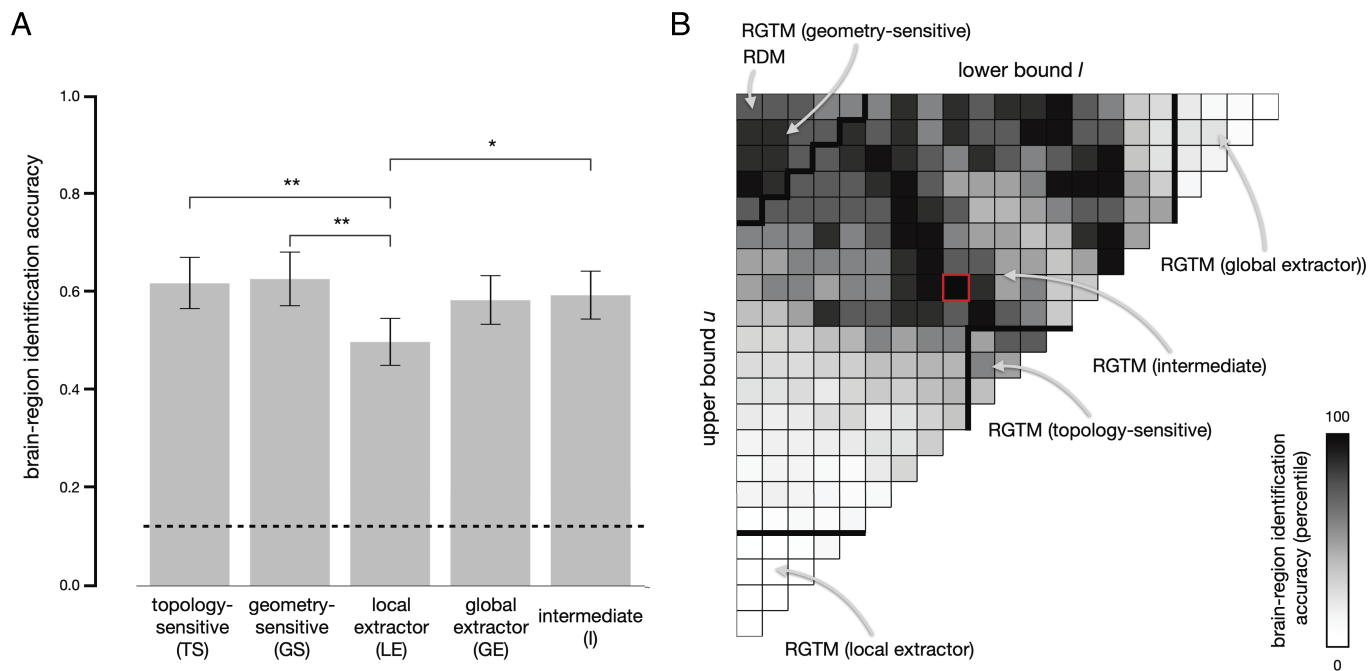


Fig. 5. Brain-region identification accuracy across the family of geotopological descriptors. Our analysis of human brain regions used fMRI data from 24 subjects, collected from 8 regions of interest, including the primary and secondary visual cortices, the lateral occipital complex, the occipital face area the fusiform face area, the parahippocampal place area, and the anterior temporal lobe. (A) Region-identification accuracy (RIA) is evaluated using leave-one-subject-out cross-validation, where a classifier (on which we compute RIA) is trained on all available data except for one subject and then tested on that *Left*-out subject. This process is repeated for each subject, with the final performance measure being the average across all iterations. We used bootstrapping to obtain an unbiased estimate of the SE as the error bound for region identification accuracy and used cross-validation to prevent overfitting. We randomly sampled 10 sets of *Lower* bounds l 's and *Upper* bounds u 's in each of the five interpretable RGTM zones as defined in Fig. 3 and applied a paired t test to compare the RIA between different RGTM zones as defined in panel (B) and Fig. 3 (**** $P < 1e-4$, *** $P < 1e-3$, ** $P < 1e-2$, * $P < 0.05$). (B) RIA percentile (gray colorscale) as a function of the combination of *Upper* and *Lower* bounds, l and u , defining the RGTM (layout as in Fig. 3) with the best-performing RGTM marked in red.

than selecting the geometry-sensitive or intermediate statistics, consistent with what we observed for the human fMRI data. This suggests that, in a low-noise regime, some of the information in the RDM is redundant and the extremes of the distance distribution can be compressed without compromising our power to distinguish representations of different computational modules (layers here). In the high-noise regime, having the full RDM information can better characterize the subtle distinctions among computational modules.

To further interpret the distributions of representations associated with different layers and instances, we performed MDS on the RDMs, RGTMs, and RGDMs (using the optimal threshold pairs (l, u) marked in red in the corresponding heatmaps in Fig. 6B). We used the Procrustes alignment (75) to obtain a consistent orientation and reflection of the MDS plots (Fig. 6C). We observe that the RDMs cluster the representations of several layers together, whereas the RGTMs and RGDMs manage to separate out layers in a smoother progression from the early layers to the later layers. This suggests that the geotopological and geodesic transforms can help emphasize topological invariants that consistently characterize layer representations across variation among model instances.

Discussion

Theories about neural codes and computations can predict how the neural manifold hangs together (i.e., the topology) without predicting a specific representational geometry. Such theories are consistent with an infinite number of distinct geometries, and so do not predict a specific RDM. Topological RSA enables researchers to evaluate and statistically compare such theories as expressed by a representational graph and associated RGTMs. Our results do not support the conclusion that topological descriptors should replace geometrical descriptors in RSA in general. Studies that aim to adjudicate among brain-computational models that predict representational geometries can continue to use geometrical descriptors. Topology may not be needed as a tool to suppress nuisance variation in RSAs of representational geometries. However, the topology of neural representations is interesting in its own right.

For theories and brain-computational models that do predict geometries, evaluation on the basis of their topological predictions provides an alternative and complementary perspective. If our models' RDM predictions do not reach the noise ceiling, the model that best predicts the geometry may not be the model that best predicts the topology. We can simply apply a geotopological transform to the brain RDM before performing model-comparative RSA inference with a rank-based comparator. The choice of parameters l and u (which can be defined as percentiles within the dissimilarities of each RDM) enables us to calibrate the relative sensitivity to the topological and geometrical properties of the representation. This approach enables comparisons of both topology-predicting and geometry-predicting models within the same tRSA framework. The fact that region identification and layer identification do not suffer when we drastically reduce the information in the RDM suggests that the bulk of the information distinguishing the representations in brain regions and neural network layers is captured by topological descriptors.

Characterizing Topology and Geometry Provides a Comprehensive View of Neural Representations. Considering topology-sensitive representational summary statistics alongside geometrical ones can provide a more comprehensive and nuanced

understanding of the data. We introduced geotopological descriptors as a class of summary statistics for RSA. These descriptors emphasize the topology by compressing variation among the smallest representational distances and among the largest representational distances. Variation among small representational distances is sensitive to noise, and variation among very large distances may reflect individual idiosyncrasies more than computationally meaningful information. While tRSA did not outperform conventional RSA at region or layer identification, it matched the model-selection performance of RSA while substantially reducing the information, thus revealing what range of variation among representational distances captures the discriminable computational signatures of different representations. Topological RSA provides a robust way to identify neural representations in spite of noise and interindividual variation. We emphasize the synergistic potential of combining the topological and geometrical perspectives.

Testing Topological Representational Hypotheses Requires

tRSA. From a theoretical perspective, the choice of using topological or geometrical summary statistics in RSA depends on the hypotheses about the neural representations that are to be tested. If the hypothesis is that a neural representation conforms to a particular topology, then topological RSA provides a straightforward method to test the hypothesis. Note that a hypothesis about the representational topology cannot be straightforwardly tested with conventional RSA because the hypothesis corresponds to a complex set of RDMs any of which conforms to the hypothesized topology. This is a major motivation for tRSA independent of the question of whether it can help reduce nuisance variation in adjudicating among models that predict geometries.

Do theories about neural mechanisms imply specific predictions about the representational topology? Investigating the relationship between brain-computational theories and the topology of neural population representations is an important direction for future computational work. Here we focused on visual representations in both the model-based simulations and the fMRI dataset, where the representational space has many dimensions and we do not have simple topological hypotheses about the structure of the representations. In the absence of simple topological hypotheses, a data-driven approach can reveal to what extent the representational topologies are consistent between individuals and distinct across stages of processing or across stages of learning or development. The relationships among brain-representational topologies can be visualized using MDS of the RGTMs or RGDMs as we show for DNN layers in Fig. 6C. Such analyses have been performed previously for RDMs (e.g., refs. 16 and 76) and require independent data for each region to prevent correlated noise fluctuations from confounding the RDM estimates as shown in ref. 77. Future studies could use tRSA to test strong predictions about the topology of lower-dimensional stimulus sets sampling, e.g., orientations and spatial frequencies of gratings, or directions and velocities of visual motion, which are known to be represented at different levels of the primate visual hierarchy. Another example where tRSA could support a strongly theory-driven approach seeking to adjudicate among alternative topological hypotheses is the head and traveling direction system in mice and fruit flies (49, 78–80).

Testing Geometrical Representational Hypotheses May Benefit from tRSA. If the models to be tested predict specific representational geometries, conventional RSA can be used to adjudicate among them. Topological descriptors are functions

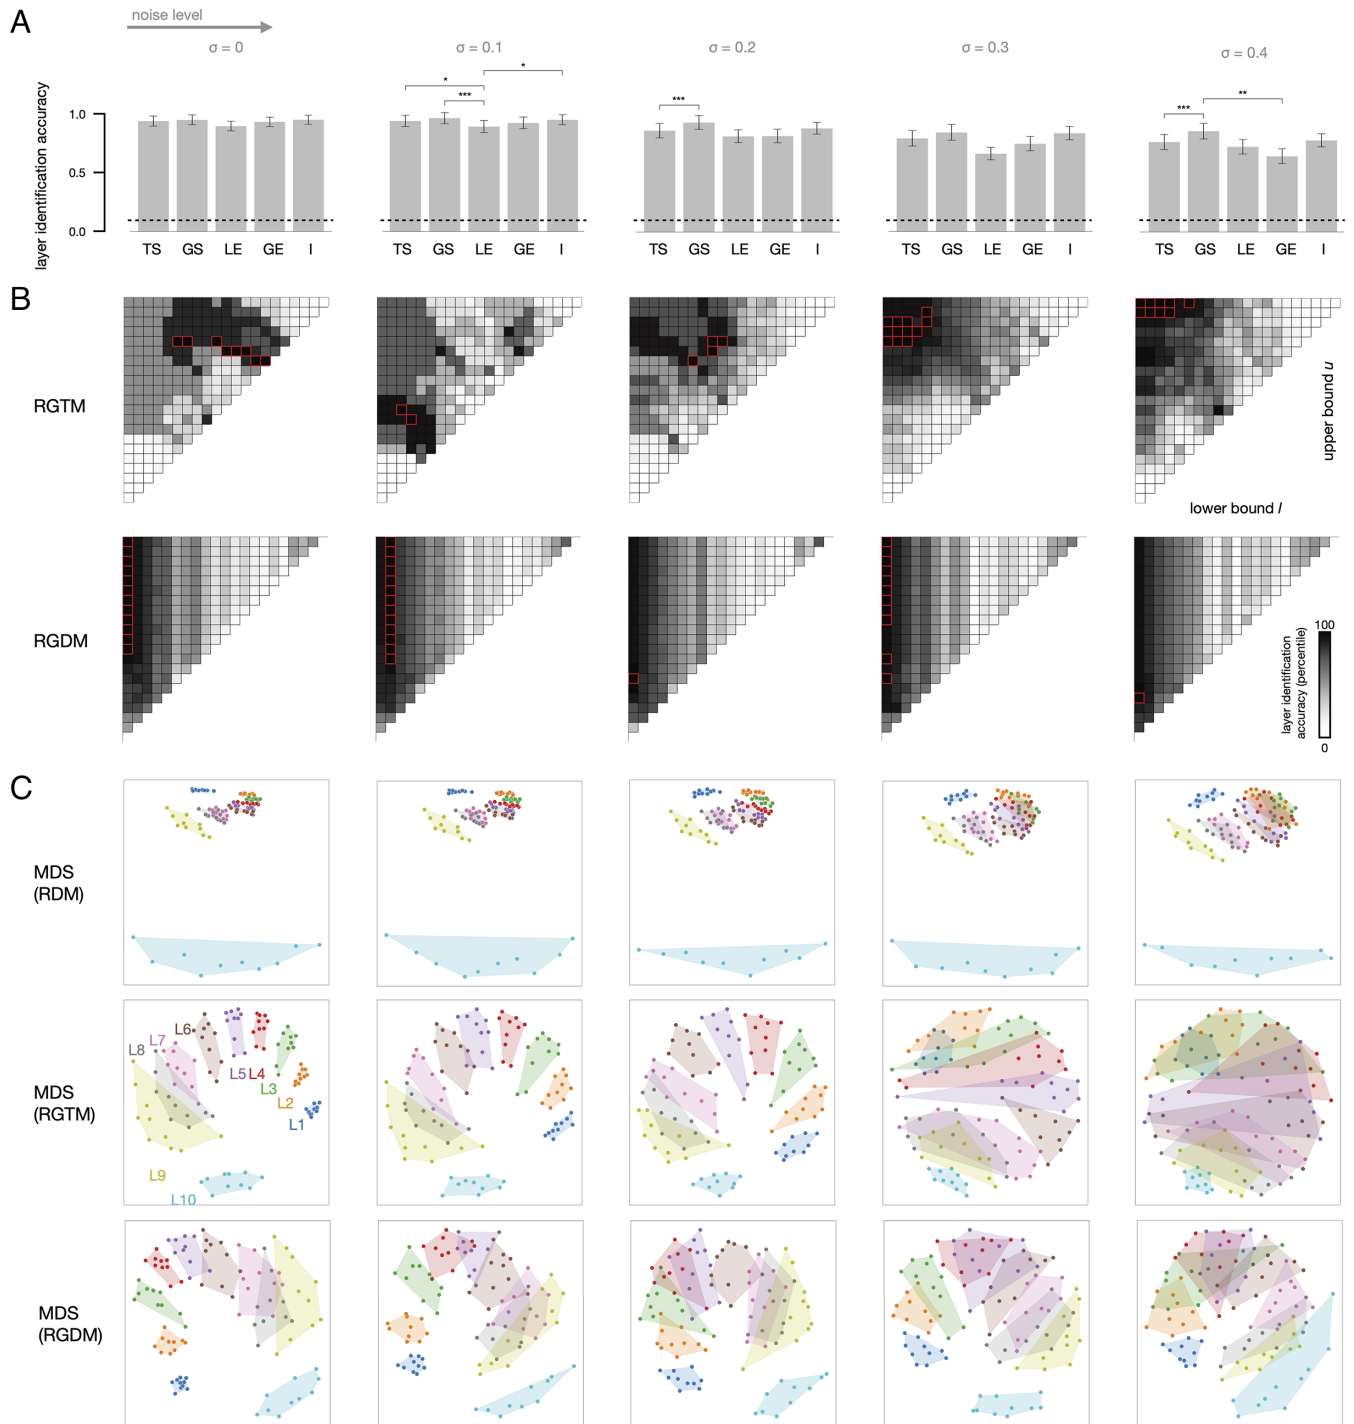


Fig. 6. Neural-network-layer identification accuracy across the family of geotopological descriptors. The performance of different RGTM at neural-network layer identification with gradually increasing Gaussian noise σ on the dissimilarity estimates (columns). Results are for All Convolutional Neural Nets, a simplified model architecture with many convolutional layers (with 10 separate layers of interest in total). Layer identification accuracy (LIA) was estimated in leave-one-instance-out cross-validation, where a classifier is trained on all available model instances except for one and then tested on that *Left-out* model instance. This process is repeated for each instance, with the final performance measure being the average across all iterations. To enable cross-validated evaluation similar to the brain-activity data in the previous figure, for each model, we trained 10 instances (analog of individual subjects) from different random seeds. Cross-validation was performed within a bootstrapping process to estimate LIA and its estimation error. We randomly sampled 10 sets of l 's and u 's in each of the five interpretable RGTM zones as defined in Fig. 3 and applied a paired t test to compare the LIAs afforded by different RGTM zones. (A) Bar graphs showing the LIA for each RGTM zone and noise level with statistical comparisons. LIA is similar for topology-sensitive RGTM and geometry-sensitive RGTM. (B) The LIA of RGTM and corresponding RGDM across the (l, u) threshold pairs are represented in a percentile heatmap, with the maximum performance marked in red. (C) The MDS plots of the RDMs (*Top* row) show geometrical similarities among layer representations. Each dot is the RDM for one layer in one model instance. For each layer, a convex hull is drawn to group its representations across all instances. MDS plots reflecting geotopological similarities are shown for the best-performing RGTM and RGDM (red in B), revealing the good discrimination of layers these geotopological summary statistics provide (*Middle* and *Lower* rows, respectively).

of geometrical ones and reduce the information, so may appear undesirable if geometrical hypotheses are to be tested. Consider the example of a set of visual gratings covering the entire cycle of orientations with equal spacing. For a wide variety of hypothetical representations, we expect the topology to be that of a circle, whereas the geometry may not be that of a circle, for example, if the neural code is anisotropic, more precisely representing cardinal orientations (which are more frequent in natural images) (18, 81, 82). The topology, in that case, fails to capture an important feature of the neural code. In other scenarios, however, a topological descriptor may capture the essential features of the neural code. If our topological descriptor reduces nuisance variation in the data more than the representational signatures that distinguish different models, it can support improved model-adjudication accuracy. For the visual representations investigated here in models (deep neural networks) and humans (fMRI), we did not find evidence for such an advantage. Future studies will reveal whether a focus on the topology reduces the noise more than the signal in other circumstances.

However, even if our models make predictions about representational geometries (as do DNNs) and tRSA does not improve model-adjudication power, it can still provide complementary information by revealing to what extent the topology implicit to the geometrical predictions enables model adjudication and which model best predicts the representational topology. If the best model fully explains the geometry, it will also correctly predict the topology. However, the relative performance of models that do not fully explain the geometry can change when we evaluate models by their topological predictions. Our analyses here showed that tRSA, using reduced topological summaries, can match the performance of geometrical RSA at model adjudication. This suggests that topological descriptors that compress the variation among small and among large representational distances capture the information essential to the computational signature of different representations.

Practical Considerations for Implementing tRSA. Topological RSA model comparison can rely on the RSA3 statistical inference framework (33), which has been implemented in Python in the open-source [RSA Toolbox](#)). The topological descriptors introduced here and the RSA3 inference methods enable analyses based on a wide range of brain-activity data, including invasive neural recordings, fMRI, electroencephalogram, and magnetoencephalography. A topological representational hypothesis can be expressed directly in a representational graph and associated RGTMs. Alternatively, topological hypotheses can be derived from representational models that predict geometries, i.e., any descriptors of the experimental conditions. If the conditions correspond to visual stimuli, for example, representational models can be derived from properties of the images or of the objects depicted and their categories (15). In particular, topological hypotheses can be derived from brain-computational models, such as neural network models that implement candidate hypotheses about the computations performed by the brain and predict a geometry for each stage of representation (70, 73, 83).

If the theories to be compared predict representational topologies (RGTMs), then we can avoid having to choose l and u . We can compare the RGTMs predicted by the theories to the brain RDM directly using the ρ_a rank correlation coefficient (33).

If the theories to be compared predict geometries (RDMs), but we intend to compare models in terms of their predictions of the topology of the brain representation, then we can convert the brain RDM to a brain RGTMs by selecting l and u to choose the desired calibration of the summary statistic to the geometry and the topology of the brain representation. The brain RGTMs can then be compared to the model RDMs using the ρ_a rank correlation coefficient and model-comparative inference can proceed on this basis. Finally, if the goal is to assess whether the geometrical predictions of a model (the model RDM) significantly exceed its topological predictions (model RGTMs, based on a choice of l and u), we can use the ρ_a rank correlations with the brain RDM to statistically compare the model's topological and geometrical predictions.

The values for l and u can be set by a priori considerations, e.g., on the basis of the level of graph connectivity predicted by the theories to be evaluated, or on the basis of results of earlier studies that use a different dataset (e.g., this study). If different settings of l and u are to be explored in the analyses, it is important to avoid biases to the analyses that can result from selecting these two parameters (84, 85). One approach is to choose a small number of settings for l and u and add each resulting RGTMs as a separate model in the model-comparative inference, where standard adjustments for multiple testing can then be used (controlling the family-wise error rate or the false-discovery rate as implemented in the [RSA Toolbox](#)). Alternatively, an independent analysis, for layer- or brain-region identification accuracy, could inform the choice of l and u . Importantly, whatever exploration or search procedures were employed in selecting l and u must be fully reported.

It is important to note that both topological and geometrical descriptors depend on the experimental conditions (e.g., stimuli) for which activity patterns are included in the analysis. However, topological descriptors can be even more sensitive to the set of conditions included. For example, the geodesic distance between the representations of two conditions depends on the other conditions included (which determine the shortest path), whereas the Mahalanobis distance does not.

In conclusion, tRSA is essential for testing topological representational hypotheses. Even when testing models that predict representational geometries, tRSA can reveal to what extent the models capture, in particular, the representational topology observed in a neural population. The combination of topological and geometrical descriptors offers a promising comprehensive approach for analyzing neural representations.

Data, Materials, and Software Availability. Previously published data were used for this work (39, 64, 65). The codes necessary to reproduce the analytical results are available at <https://github.com/doi1bh/TopologicalRSA> (86).

Author affiliations: ^aDepartment of Artificial Intelligence and Human Health, Hasso Plattner Institute for Digital Health, Icahn School of Medicine at Mount Sinai, New York, NY 10029; ^bDepartment of Psychiatry, Center for Computational Psychiatry, Icahn School of Medicine at Mount Sinai, New York, NY 10029; ^cDepartment of Neuroscience, Friedman Brain Institute, Icahn School of Medicine at Mount Sinai, New York, NY 10029; ^dZuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027; ^eDepartment of Psychology, Columbia University, New York, NY 10027; and ^fDepartment of Neuroscience, Columbia University, New York, NY 10027

Author contributions: B.L. and N.K. designed research; B.L. performed research; B.L. contributed new reagents/analytic tools; B.L. analyzed data; and B.L. and N.K. wrote the paper.

1. A. M. Dale, B. Fischl, M. I. Sereno, Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage* **9**, 179–194 (1999).
2. R. S. Desikan *et al.*, An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**, 968–980 (2006).

3. B. Fischl *et al.*, Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron* **33**, 341–355 (2002).
4. N. Kriegeskorte, R. Goebel, An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. *Neuroimage* **14**, 329–346 (2001).

5. A. M. Bastos, J. M. Schoffelen, A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Front. Syst. Neurosci.* **9**, 175 (2016).
6. C. J. Honey *et al.*, Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 2035-2040 (2009).
7. M. P. Van Den Heuvel, H. E. H. Pol, Exploring the brain network: A review on resting-state fMRI functional connectivity. *Euro. Neuropsychopharmacol.* **20**, 519-534 (2010).
8. D. S. Bassett, O. Sporns, Network neuroscience. *Nat. Neurosci.* **20**, 353-364 (2017).
9. J. D. Medaglia, F. Pasqualetti, R. H. Hamilton, S. L. Thompson-Schill, D. S. Bassett, Brain and cognitive reserve: Translation via network control theory. *Neurosci. Biobehav. Rev.* **75**, 53-64 (2017).
10. S. Gu *et al.*, Controllability of structural brain networks. *Nat. Commun.* **6**, 8414 (2015).
11. P. Srivastava *et al.*, Models of communication and control for brain networks: Distinctions, convergence, and future outlook. *Netw. Neurosci.* **4**, 1122-1159 (2020).
12. E. Bullmore, O. Sporns, The economy of brain network organization. *Nat. Rev. Neurosci.* **13**, 336-349 (2012).
13. F. V. Farahani, W. Karwowski, N. R. Lighthall, Application of graph theory for identifying connectivity patterns in human brain networks: A systematic review. *Front. Neurosci.* **13**, 585 (2019).
14. O. Sporns, Graph theory methods: Applications in brain networks. *Dialogues Clin. Neurosci.* **20**, 111-121 (2022).
15. N. Kriegeskorte *et al.*, Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* **60**, 1126-1141 (2008).
16. N. Kriegeskorte, M. Mur, P. A. Bandettini, Representational similarity analysis: connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
17. N. Kriegeskorte, R. A. Kievit, Representational geometry: Integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17**, 401-412 (2013).
18. N. Kriegeskorte, X. X. Wei, Neural tuning and representational geometry. *Nat. Rev. Neurosci.* **22**, 703-718 (2021).
19. B. Sorscher, S. Ganguli, H. Sompolinsky, Neural representational geometry underlies few-shot concept learning. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2200800119 (2022).
20. E. H. Nieh *et al.*, Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80-84 (2021).
21. J. B. Freeman, R. M. Stolier, J. A. Brooks, B. S. Stillerman, The neural representational geometry of social perception. *Curr. Opin. Psychol.* **24**, 83-91 (2018).
22. S. Chung, L. Abbott, Neural population geometry: An approach for understanding biological and artificial neural networks. *Curr. Opin. Neurobiol.* **70**, 137-144 (2021).
23. N. Kriegeskorte, J. Diedrichsen, Peeling the onion of brain representations. *Annu. Rev. Neurosci.* **42**, 407-432 (2019).
24. B. Sievers *et al.*, Visual and auditory brain areas share a representational structure that supports emotion perception. *Curr. Biol.* **31**, 5192-5203 (2021).
25. A. Fournel, C. Ferdenzi, C. Sezille, C. Rouby, M. Bensafi, Multidimensional representation of odors in the human olfactory cortex. *Hum. Brain Mapp.* **37**, 2161-2172 (2016).
26. J. M. van Baar, L. J. Chang, A. G. Sanfey, The computational and neural substrates of moral strategies in social decision-making. *Nat. Commun.* **10**, 1483 (2019).
27. M. A. Thornton, J. P. Mitchell, Theories of person perception predict patterns of neural activity during mentalizing. *Cereb. Cortex* **28**, 3505-3520 (2018).
28. M. A. Thornton, D. Tamir, The brain represents situations and mental states as sums of their action affordances. *PsyArXiv [Preprint]* (2023). <https://doi.org/10.31234/osf.io/7eskf> (Accessed 1 August 2023).
29. D. I. Tamir, M. A. Thornton, J. M. Contreras, J. P. Mitchell, Neural evidence that three dimensions organize mental state representation: Rationality, social impact, and valence. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 194-199 (2016).
30. D. I. Tamir, M. A. Thornton, Modeling the predictive social mind. *Trends Cogn. Sci.* **22**, 201-212 (2018).
31. D. B. Ehrlich, J. D. Murray, Geometry of neural computation unifies working memory and planning. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2115610119 (2022).
32. H. Nili *et al.*, A toolbox for representational similarity analysis. *PLoS Comput. Biol.* **10**, e1003553 (2014).
33. H. H. Schütt, A. D. Kipnis, J. Diedrichsen, N. Kriegeskorte, Statistical inference on representational geometries. *eLife* **12**, e82566 (2023).
34. S. M. Khaligh-Razavi, N. Kriegeskorte, Deep supervised, but not unsupervised, models may explain cortical representation. *PLoS Comput. Biol.* **10**, e1003915 (2014).
35. T. C. Kietzmann *et al.*, Recurrence is required to capture the representational dynamics of the human visual system. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 21854-21863 (2019).
36. R. M. Cichy, D. Pantazis, Multivariate pattern analysis of MEG and EEG: A comparison of representational structure in time and space. *NeuroImage* **158**, 441-454 (2017).
37. T. Konkle, C. Conwell, J. S. Prince, G. A. Alvarez, What can 5.17 billion regression fits tell us about the representational format of the high-level human visual system? *J. Vis.* **22**, 4422-4422 (2022).
38. R. J. Low, S. Lewallen, D. Aronov, R. Nevers, D. W. Tank, Probing variability in a cognitive map using manifold inference from neural dynamics. *bioRxiv [Preprint]* (2018). <https://doi.org/10.1101/418939> (Accessed 1 August 2023).
39. J. Mehrer, C. J. Spoeer, N. Kriegeskorte, T. C. Kietzmann, Individual differences among deep neural network models. *Nat. Commun.* **11**, 1-12 (2020).
40. L. Wasserman, Topological data analysis. *Annu. Rev. Stat. Appl.* **5**, 501-532 (2018).
41. H. Edelsbrunner *et al.*, Persistent homology-a survey. *Contemp. Math.* **453**, 257-282 (2008).
42. G. Singh, F. Mémoli, G. E. Carlsson, Topological methods for the analysis of high dimensional data sets and 3d object recognition. *PBG Eurogr.* **2**, 091-100 (2007).
43. A. H. Rizvi *et al.*, Single-cell topological RNA-seq analysis reveals insights into cellular differentiation and development. *Nat. Biotechnol.* **35**, 551-560 (2017).
44. B. Lin, Topological data analysis in time series: Temporal filtration and application to single-cell genomics. *Algorithms* **15**, 371 (2022).
45. C. Geniesse, O. Sporns, G. Petri, M. Saggari, Generating dynamical neuroimaging spatiotemporal representations (DyNeuSR) using topological data analysis. *Netw. Neurosci.* **3**, 763-778 (2019).
46. B. Lin, "Single-cell topological simplicial analysis reveals higher-order cellular complexity" in 2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), D. Adjeroh *et al.*, Eds. (IEEE, 2022), pp. 1933-1940.
47. C. T. Ellis, M. Lesnick, G. Henselman-Petrusek, B. Keller, J. D. Cohen, Feasibility of topological data analysis for event-related fMRI. *Netw. Neurosci.* **3**, 695-706 (2019).
48. J. A. Pike *et al.*, Topological data analysis quantifies biological nano-structure from single molecule localization microscopy. *Bioinformatics* **36**, 1614-1621 (2020).
49. R. Chaudhuri, B. Geřek, B. Pandey, A. Peyrache, I. Fiete, The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **22**, 1512-1520 (2019).
50. G. Singh *et al.*, Topological analysis of population activity in visual cortex. *J. Vis.* **8**, 11-11 (2008).
51. R. J. Gardner *et al.*, Toroidal topology of population activity in grid cells. *Nature* **602**, 123-128 (2022).
52. J. B. Tenenbaum, V. Silva, J. C. Langford, A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319-2323 (2000).
53. S. T. Roweis, L. K. Saul, Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**, 2323-2326 (2000).
54. L. Van der Maaten, G. Hinton, Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579-2605 (2008).
55. L. McInnes, J. Healy, J. Melville, Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv [Preprint]* (2018). <https://arxiv.org/abs/1802.03426> (Accessed 1 August 2023).
56. B. Lin, N. Kriegeskorte, Adaptive geo-topological independence criterion. *arXiv [Preprint]* (2018). <https://arxiv.org/abs/1810.02923> (Accessed 1 August 2023).
57. B. Lin, "Geometric and topological inference for deep representations of complex networks" in *Companion Proceedings of the Web Conference 2022, WWW '22*, F. Laforest, R. Troncy, L. Médini, I. Herman, Eds. (Association for Computing Machinery, New York, NY, USA, 2022), pp. 334-338.
58. M. Rigotti *et al.*, The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585-590 (2013).
59. L. Kushnir, S. Fusi, Neural classifiers with limited connectivity and recurrent readouts. *J. Neurosci.* **38**, 9900-9924 (2018).
60. T. M. Cover, Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Trans. Elect. Comput.* **EC-14**, 326-334 (1965).
61. A. J. Zomorodion, *Topology for Computing* (Cambridge University Press, 2005), vol. 16.
62. G. Carlsson, Topology and data. *Bull. Am. Math. Soc.* **46**, 255-308 (2009).
63. E. W. Dijkstra, "A note on two problems in connexion with graphs" in *Edsger Wybe Dijkstra: His Life, Work, and Legacy*, K. R. Apt, T. Hoare, Eds. (Association for Computing Machinery, 2022), pp. 287-290.
64. A. Walthner, "Beyond brain decoding: Representational distances and geometries," PhD thesis, University of Cambridge (2015).
65. A. Walthner, J. Diedrichsen, M. Mur, S. M. Khaligh-Razavi, N. Kriegeskorte, Sudden emergence of categoricity at the lateral-occipital stage of ventral visual processing. *J. Vis.* **16**, 407-407 (2016).
66. Y. LeCun *et al.*, Convolutional networks for images, speech, and time series. *Handb. Brain Theory Neural Netw.* **3361**, 1995 (1995).
67. J. T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, Striving for simplicity: The all convolutional net. *arXiv [Preprint]* (2014). <https://arxiv.org/abs/1412.6806> (Accessed 1 August 2023).
68. A. Krizhevsky, *Learning Multiple Layers of Features from Tiny Images* (University of Tront, 2009).
69. J. Diedrichsen *et al.*, Comparing representational geometries using whitened unbiased-distance-matrix similarity. *arXiv [Preprint]* (2020). <https://arxiv.org/abs/2007.02789> (Accessed 1 August 2023).
70. D. L. Yamins, J. J. DiCarlo, Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356-365 (2016).
71. C. F. Cadieu *et al.*, Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS Comput. Biol.* **10**, e1003963 (2014).
72. B. Long, C. P. Yu, T. Konkle, Mid-level visual features underlie the high-level categorical organization of the ventral stream. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E9015-E9024 (2018).
73. N. Kriegeskorte, Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annu. Rev. Vis. Sci.* **1**, 417-446 (2015).
74. K. Dwivedi, M. F. Bonner, R. M. Cichy, G. Roig, Unveiling functions of the visual cortex using task-specific deep neural networks. *PLoS Comput. Biol.* **17**, e1009267 (2021).
75. B. Lin, M. Mur, T. Kietzmann, N. Kriegeskorte, "Visualizing representational dynamics with multidimensional scaling alignment" in *Conference on Cognitive Computational Neuroscience*, H. Nienborg, R. Poldrack, T. Naselaris, Eds. (Berlin, Germany, 2019).
76. M. Visconti di Oleggio Castello, Y. O. Halchenko, J. S. Guntupalli, J. D. Gors, M. I. Gobbini, The neural representation of personally familiar and unfamiliar faces in the distributed system for face perception. *Sci. Rep.* **7**, 12237 (2017).
77. L. Henriksson, S. M. Khaligh-Razavi, K. Kay, N. Kriegeskorte, Visual representations are dominated by intrinsic fluctuations correlated between areas. *NeuroImage* **114**, 275-286 (2015).
78. C. Langdon, M. Genkin, T. A. Engel, A unifying perspective on neural manifolds and circuits for cognition. *Nat. Rev. Neurosci.* **24**, 363-377 (2023).
79. S. S. Kim, A. M. Hermundstad, S. Romani, L. Abbott, V. Jayaraman, Generation of stable heading representations in diverse visual scenes. *Nature* **576**, 126-131 (2019).
80. C. Lyu, L. Abbott, G. Maimon, Building an allocentric travelling direction signal via vector computation. *Nature* **601**, 92-97 (2022).
81. S. Appelle, Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychol. Bull.* **78**, 266 (1972).
82. A. R. Girshick, M. S. Landy, E. P. Simoncelli, Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926-932 (2011).
83. N. Kriegeskorte, P. K. Douglas, Cognitive computational neuroscience. *Nat. Neurosci.* **21**, 1148-1160 (2018).
84. N. Kriegeskorte, W. K. Simmons, P. S. Bellgowan, C. I. Baker, Circular analysis in systems neuroscience: The dangers of double dipping. *Nat. Neurosci.* **12**, 535-540 (2009).
85. M. Hosseini *et al.*, I tried a bunch of things: The dangers of unexpected overfitting in classification of brain data. *Neurosci. Biobehav. Rev.* **119**, 456-467 (2020).
86. B. Lin, TopologicalRSA. *GitHub*. <https://github.com/doerlbh/TopologicalRSA>. Deposited 5 September 2023.