EXPLORING GEOMETRIC REPRESENTATIONAL ALIGNMENT THROUGH OLLIVIER RICCI CURVATURE AND RICCI FLOW

Anonymous authors
Paper under double-blind review

ABSTRACT

Representational analysis explores the encoding of input data in high-dimensional spaces within distributed neural activations and facilitates the comparison of different systems, such as artificial neural networks and brains. Although existing methods offer relevant information, they typically do not account for local intrinsic geometrical properties within high-dimensional representation spaces. To overcome these limitations, we explore Ollivier Ricci curvature and Ricci flow as tools to study the similarity and alignment of representations between humans and artificial neural systems on a geometric basis. We used both simulations and a proof-of-principle study, in which we compared the representations of face stimuli between VGG-Face, a human-aligned version of VGG-Face, and the corresponding human similarity judgments from a large online study. Using this discrete geometric framework, we were able to identify global and local structural similarities and differences by examining distributions of node and edge curvature and higher-level properties by detecting and comparing community structure in the representational graphs.

1 INTRODUCTION

Artificial neural networks (ANN) can match human performance in image recognition and classification tasks, among others (Mohsenzadeh et al., 2020). This led to the investigation of how ANNs encode, transform, and generalize information, and if and how these processes can be related to the brain (representational alignment) (Richards et al., 2019). One key direction is to study how the geometry of internal representations reflects item similarity, categorical divisions, or latent variation within the input data (Chung & Abbott, 2021). The similarity between vectors in the representational space is commonly measured using a kernel or a Representational Dissimilarity Matrix (RDM); RDMs are then usually compared by applying Representational Similarity Analysis (RSA) by computing the correlations between RDMs (Kriegeskorte et al., 2008). Despite its popularity, RSA has notable limitations, e.g., Dujmović et al. (2022) showed that the correspondence between activation patterns in different systems can depend on the dataset, and that seemingly similar representational geometries actually encode different features. To improve this, it is crucial to capture the intrinsic geometric properties of high-dimensional data, assuming that these data lie on a lower-dimensional manifold within the ambient space, where the relationships between points reflect meaningful variations and similarities in the data (Lin & Zha, 2008). However, RSA imposes an ambient space geometry (Euclidean) that can distort these intrinsic geometrical relationships.

To respect the manifold's geometry and to overcome the limitations of classic RSA, we utilize a graph representation of the data (as a discrete counterpart of a manifold) and employ the Ollivier Ricci curvature (ORC) method, a discrete, graph-based analogue of Ricci curvature in Riemannian geometry (Ollivier, 2007). ORC reflects the geometry of the graph by considering local neighborhoods and computing the optimal transport between probability measures at each node of the graph. At the same time, ORC avoids imposing a metric on the data manifold from the ambient space. ORC measures the deviation of a neighborhood structure from being flat (in the discrete case, this is a grid-like topology). This notion of non-flatness relates to patterns of local connectivity in the graph. The positive curvature of an edge means that there are more connections, and the negative curvature indicates fewer connections compared to the grid. Eventually, the Ricci flow process allows us to reveal the communities of a graph, which resemble regions in Riemannian manifolds of large positive curvature.

We applied ORC to synthetic data and to data from a human online experiment (Anonymous), in which 1,397 participants were asked to judge similarities between 100 generated faces in a triplet-odd-one-out task (161,700 triplet combinations). Subsequently, human behavior was modeled using a custom adaptation of VGG-Face (Parkhi et al., 2015) that was trained to predict human choices given sets of three face images (Aligned VGG-Face), see Methods and Anonymous.
The tools from discrete geometry enabled us to compare similarity spaces from human responses and activation patterns from both the VGG-Face and the Aligned VGG-Face model. For the former, we used activations in the FC7 layer. For the latter, we took activations in the VGG-bridge, a dense layer connecting a frozen, pre-trained part of the original VGG-Face with a new decision block specific to the similarity judgment task (see Figure 2 and section 2.7). Then, we constructed graph representations based on human responses (Human Judgment graph), Aligned VGG-Face, and VGG-Face. We applied ORC to analyze local neighborhood structures in the graphs and to measure how the additional information introduced by the alignment process remodels the structural properties of representations in the original VGG-Face. We first compared the global structure between representations based on the adapted heat diffusion distance for our graph representations, computed between curvature-weighted graphs Hammond et al. (2013). To compare structural properties at an intermediate scale, we used edge properties derived from edge curvature values and applied Ricci flow to detect communities within graph representations. Finally, to identify which image features serve as indicators for the detected communities in each representation, we conducted an analysis comparing significant features within and between representations.

1.1 Related work

Representational alignment Sucholutsky et al. (2023) examine representational alignment—the correspondence between internal representations in biological and artificial systems. The authors propose a framework to measure and enhance this alignment, drawing from cognitive science, neuroscience, and machine learning and they address three key challenges: (1) measuring alignment, (2) mapping representations into a shared space, and (3) improving alignment across systems. Using RSA Kriegeskorte et al. (2008), various studies aim to compare the representations of artificial and biological systems, e.g., Khaligh-Razavi & Kriegeskorte (2014) compare supervised and unsupervised learning models to determine which model explains inferior temporal (IT) neural representations better and explore how well computational models replicate the representational geometry of the IT cortex. In addition to the RSA method, other techniques have been developed to explore neural network representations. Kornblith et al. (2019) compares neural network representations using Canonical Correlation Analysis (CCA), showing that CCA and similar statistics fail when representations exceed the number of data points. They introduce Centered Kernel Alignment, which replaces invariance to invertible linear transformations with orthogonal invariance.

Curvature of graphs and networks The concept of curvature, a geometric measure of how space deviates from flatness, has been utilized in various studies to modify underlying systems or to compare different representations derived from these systems, uncovering structural (dis)similarities. The study by Gosztolai & Arnaudon (2021) introduces an extension of ORC by incorporating the similarity of dynamical processes, like diffusion, at neighboring nodes, captures process evolution, revealing network geometry. They show that curvature distribution evolves with gaps at key timescales, marking bottleneck edges that restrict information flow and it effectively detects community structures in synthetic and real-world networks. In another work, ORC has been employed to assess the robustness of connections in brain structural networks Farooq et al. (2019). By applying curvature-based measures, the authors identify robust and fragile brain regions in healthy individuals and demonstrate that curvature effectively tracks age-related changes and alterations in brain connectivity associated with autism spectrum disorder. Sandhu et al. (2015) apply a discrete Ricci curvature, adapted for weighted graphs, to assess the ”shape” or ”bending” of the network at each edge. This approach provides insights into local connectivity and structural robustness, distinguishing cancerous from normal biological networks.

2 METHODS

2.1 GRAPH CONSTRUCTION

To construct the graph from vector embeddings, we employ an adaptive nearest-neighbor method. First, we calculate the density as the inverse of the distance using a k-nearest neighbors (KNN) density kernel, incorporating the parameters k_{min} and k_{max}, which specify the minimum and maximum number of neighbors, respectively. Next, we normalize the local density by scaling it between the minimum and maximum density values, thereby defining the density at each data point. Based on this normalized density, we determine the number of neighbors for each data point and proceed to construct the graph.

The same approach is applied to the similarity matrix derived from human similarity judgments. First, a distance matrix is constructed from the similarity matrix provided. Using this distance matrix, the method estimates an appropriate k-value for each data point (image) based on the distance distribution from that point. This results in an adaptive KNN structure for the dataset, which in turn forms the foundation for an adaptive graph construction.

2.2 OLLIVIER RICCI CURVATURE

In Riemannian geometry, curvature describes how a manifold deviates from being locally similar to Euclidean space, with Ricci curvature specifically measuring this deviation in various tangent directions Samal et al. (2018). Geometrically, Ricci curvature influences the rate at which the volume of a ball expands as its radius increases, as well as the volume of the overlap between two balls, depending on their radii and the distance between their centers. In addition, the overlap volume between two balls is directly connected to the transport cost needed to move one ball to the other: a greater overlap volume implies a lower transport cost. This relationship highlights a connection between Ricci curvature and optimal transport. Using this concept, Ollivier introduced a generalized form of the Ricci curvature in metric measure spaces based on optimal transport Ollivier (2007). For a metric space (X, d) equipped with a probability measure m_x for each $x \in X$, the Ollivier Ricci curvature (ORC) κ_{xy} along a path xy is defined as follows:

$$\kappa_{xy} = 1 - \frac{W_1(m_x, m_y)}{d(x, y)} \tag{1}$$

where $W_1(m_x, m_y)$ is the Wasserstein distance.

2.3 RICCI FLOW

The Ricci flow method, based on the geometric concept of curvature introduced by F. Gauss and B. Riemann, describes how the space bends at each point Perelman (2002); Gauss (1828); Jost (2016). Areas with high positive curvature are denser, while regions with negative curvature are less so. Hamilton developed the Ricci flow, a curvature-driven diffusion process, which deforms space similarly to heat diffusion; regions with large positive curvature contract, while those with strong negative curvature expand Hamilton (1982). Ni et al. (2019) adapted Ricci flow from Riemannian geometry to discrete networks, using it to detect community structures within graphs. The discrete Ricci flow algorithm on a network is an evolving process. In each iteration, all edge weights update simultaneously by the following flow process:

$$w_{xy}^{i+1} = d^i(x, y) - \kappa^i(x, y) \cdot d^i(x, y) \tag{2}$$

where w_{xy}^i is the weight of the edge xy in the i-th iteration, and κ_{xy}^i is the Ricci curvature at the edge xy in the i-th iteration, and $d_{(x,y)}^i$ is the shortest path distance in the graph induced by the weights w_{xy}^i. Initially, we set $w_{xy}^0 = w_{xy}$ and $d_{xy}^0 = d_{xy}$

2.4 HEAT DIFFUSION DISTANCE

To compare graphs based on the curvature of their edges, we employ the graph diffusion distance (Hammond et al., 2013). This metric quantifies the average similarity of heat diffusion across each graph and is rooted in the framework of diffusion maps, enabling comparisons between weighted graphs. The edge weights represent the conductivity between the vertices, capturing how changes in the structure of the graph influence the transmission of heat, information, or other quantities across the graph. The process involves generating a diffusion pattern centered around a vertex i by initializing it with a localized delta impulse at that vertex and allowing the diffusion to evolve over a specified time t. Different adjacency matrices yield distinct diffusion patterns. The graph diffusion distance is calculated as the average norm of the differences between such patterns for any two adjacency matrices. The distance between two graphs G_1 and G_2 and for time t could be calculated with:

$$d(G_1, G_2; t) = \|exp(-tL_1) - exp(-tL_2)\|_F^2 \tag{3}$$

where $\|.\|_F$ is the Frobenius norm and L_1 and L_2 are Laplacian matrices of G_1 and G_2, respectively.

In this study, we model our graphs as weighted graphs, with curvature values assigned as edge weights. This approach allows us to analyze and compare variations between different graph representations. Our study assessed representations based on the heat diffusion distance between these weighted graphs. The edge weights in the graphs were determined using curvature values.

2.5 GRAPH COMMUNITY MEASURES

We computed a number of quantities to measure the community structure in our representational graphs. *Modularity* measures the strength of clustering by comparing the density of edges within communities to the density of edges on a random graph with the same degree distribution. *Average-Embeddedness:* measures the number of shared neighbors for pairs of nodes within a community. It captures the cohesiveness (i.e., interconnectedness) of a community. *Internal Edge Density:* measures the density of edges within the community compared to the maximum possible number of edges in that community.

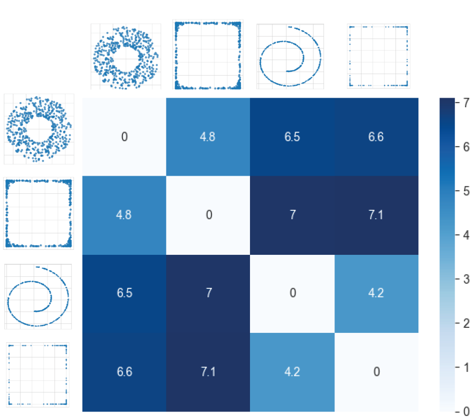

Figure 1: Pairwise heat diffusion distances between the different synthetic datasets. From top to bottom (left to right): 2D torus, transformed 2D torus, 2D swiss roll, transformed 2D swiss roll. Colors indicate heat diffusion distance.

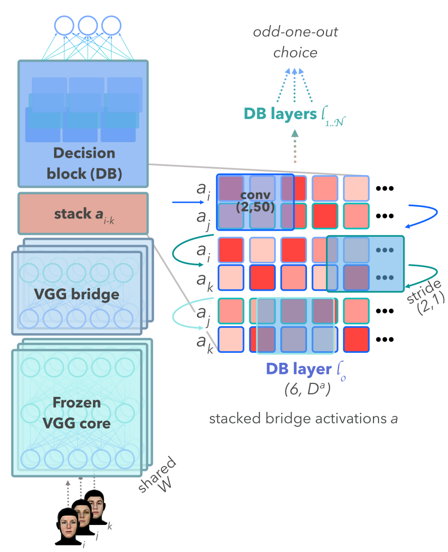

Figure 2: Human-aligned VGG-Face. The network is trained to predict human judgments in a face similarity task. The figure adopted from Anonymous

2.6 SYNTHETIC DATA

To evaluate whether our representational alignment approach, based on a graph framework, can capture local geometric properties while comparing different representations with varying underlying geometries, we generated several synthetic datasets. Specifically, we created a 2D torus dataset and a dataset resembling a 2D version of a swiss roll, as illustrated in Figure1. To examine the extent to which variation in the underlying geometry can be captured by the heat diffusion distance, we transformed the original data using the sigmoid function $S(x) = \frac{1}{1+e^{-x}}$. The data sets are shown in Figure 1. We then constructed a graph representation of each data set as outlined above and computed the pairwise heat diffusion distance. Figure 1 shows the discrepancy captured by the heat diffusion distance, indicating the underlying structural differences between the datasets and their sigmoid compressions. Comparing the distance values within and across data sets shows that the heat diffusion distance could distinguish underlying structures that are more similar within a dataset and its transformation than between datasets or their transforms.

2.7 HUMAN EXPERIMENT DATA

Face stimuli and human similarity judgments Stimulus images were computed using the 3D reconstruction model DECA (Feng et al., 2021) applied to 2D portraits of the Chicago Face Dataset (Ma et al., 2015). Human face similarity judgments (n = 194,261) were acquired in the form of a triplet-odd-one-out task from 1,397 participants (age range 18 - 65, mean age = 31.9 ± 11.2 years) in an online experiment. For more details on the stimulus set and experimental design, we refer to Anonymous.

Human-aligned VGG-Face The pre-trained VGG-Face architecture (Parkhi et al., 2015) was adopted to predict human face similarity judgments in the experiment (Figure 2). First, all layers up to the fully connected layer FC7 were frozen, making their weights non-trainable (*VGG core*). Second, subsequent layers were replaced with one FC layer (*VGG bridge*), which converts a 4,096-dimensional input to a 300-dimensional vector. A *decision block* was added consisting of convolutional layers. This block receives stacked activation maps from the bridge for each input image in a triplet (x_i, x_j, x_k), resulting in a 6x300 matrix $[a_i, a_j, a_i, a_k, a_j, a_k]$. The first convolutional layer in the decision block has 2 filters of size (2, 50) and stride (2, 1), producing an output of size (batch size, 2, 3, 251). After applying a ReLU activation, another convolutional layer with one filter of size (3, 100) and stride (1, 1) is applied, followed by another ReLU. This results in a (batch size, 2, 1, 152) output. Then, the signal was down-sampled to (batch size, 1, 1, 3) using two more convolutional layers (one filter each, kernel sizes: (1, 100) and (1, 51)) with an intermediate ReLU. The resulting 3-length output vector indicates the model's choice, where the highest value identifies the odd-one-out. The architecture was trained using cross-entropy loss with the Adam optimizer, a learning rate of $5e^{-4}$, and a batch size of 16. The data (X: triplet images, Y: human choices) was split into training (70%; $n_{train} = 135,982$), validation (15%; $n_{val} = 29,139$), and test sets (15%; $n_{test} = 29,140$). For more details see Anonymous.

3 RESULTS

3.1 GRAPH-BASED REPRESENTATIONAL SIMILARITY ANALYSIS

To understand the extent to which geometry-informed distance metrics differ from a standard Euclidean distance, we conducted a representational similarity analysis using pairwise distances between data points (faces). We computed similarity matrices using three different metrics for each pair of face representations: Euclidean distance between the representational vectors in the embedding vector space, the geodesic (shortest path) distance on the graphs constructed from the given representations (or similarity judgments), and a curvature-weighted geodesic distance in these graphs. The latter incorporates curvature values assigned to edges as weights, allowing us to evaluate whether the incorporation of local geometric information influences the representational similarity analysis. Figure 3 shows the resulting similarity matrices for the Euclidean case (a), the geodesic distance (b) and the weighted geodesic (c). First, the overall structure is consistent between the Euclidean and geodesic distances. Overall, the geodesic distance seems to capture more structural nuances in the representations. Interestingly, the geodesic distances in the graph constructed

Table 1: Correlation of pairwise distance matrices from different representations with the matrix of human similarity judgments.

Distance Metric	**Aligned VGG-Face**	**VGG-Face**
Euclidean	0.63	0.45
Geodesic	0.60	0.24
Weighted Geodesic	0.54	0.25

from human similarity judgments shows a clearer clustering compared to raw human similarity judgments. This highlights a potential benefit of computing representational similarities from local information.

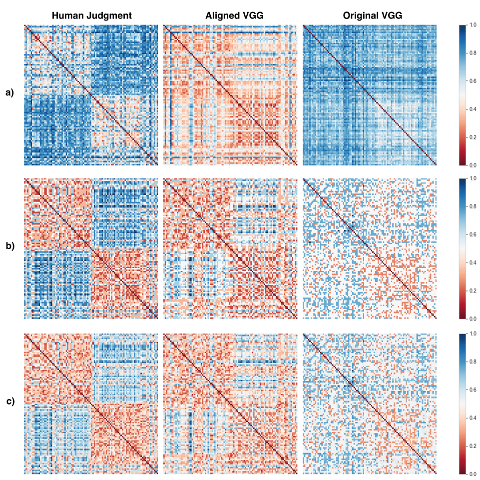

Figure 3: Pairwise distance matrices for Human Judgment, Aligned VGG-Face and original VGG-Face (range [-1, 1]). In a), Euclidean metric and in b) geodesic distances and in c) weighted geodesic distances between points have been computed.

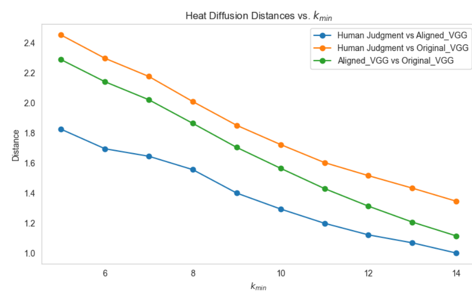

Figure 4: Comparison of heat diffusion distances between pairs of representational graphs across a range of k_{min} parameters for the adaptive graph construction. Each line represent a pairwise comparison (Blue: Human Judgments vs. Aligned VGG-Face, orange: Human Judgment vs. VGG-Face, and green: Aligned VGG-Face vs. VGG-Face).

We also computed the correlation of each representational similarity matrix to the human similarity judgment matrix within the same framework. Table 1 shows that the graph-based results are consistent with the standard Euclidean approach. It also describes adding more information from local geometry decreases the correlation indicating structural difference captured by graph-based distance methods. Further, the graph-based distances indicate a more pronounced difference between the representational geometry of the aligned VGG-Face and the original VGG-Face.

3.2 Heat diffusion distance between representational graphs across different values for k_{min}

To investigate how the mutual distances between representations change by adjusting the parameters of our adaptive graph construction, we conducted a comparative analysis of the heat diffusion distance. Figure 4 shows how the distances between the representations vary as k_{min} range from 5 to 15 and k_{max} range from 15 to 25. Although the absolute magnitude of the distances decreases, the plot reveals that their relative differences remain consistent for most of the range of parameters we tested. Specifically, the smallest distance is observed between Human Judgment and Aligned VGG, followed by Aligned VGG vs Original VGG, and finally Human Judgment vs Original VGG. The comparison across k_{max} values follow the same pattern.

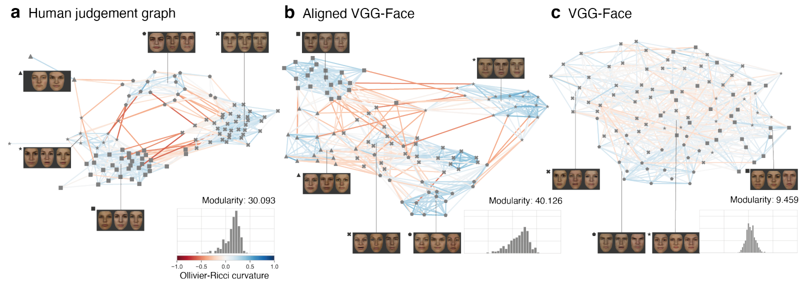

Figure 5: Detected communities (different symbols) by Ricci flow in the three experiments and the corresponding edge curvature value distribution. For each community, we show a random sample of faces drawn from it. The edge colors in graphs indicate curvature values, with red shades representing negative curvature and blue shades indicating positive values. Modularity measurement as a community metric is attributed to each graph representation.

Table 2: Comparison of graph structure based on KL-divergence (KLD) between edge curvature distributions and community metrics derived from Ricc flow: *Conductance*, *Internal edge density* (IED), *Modularity*, and *Average embeddedness* (AE).

Graph	KLD	Communities	Conductance	IED	Modularity	AE
Human judgment graph	0.000	5	0.297	0.563	30.093	0.767
Aligned VGG-Face	0.183	5	0.173	0.550	40.126	0.827
VGG-Face	0.697	4	0.417	0.386	9.459	0.604

3.3 A comparative analysis of node and edge curvatures across models

In our graphs, nodes represented faces, and an edge between two nodes indicated that the corresponding images were considered similar by human participants or in terms of the vector representation within the two ANNs. We first computed ORC on the graphs' edge level and then derived the node curvature by averaging the curvature values across all edges connected to each node.

We observed a notable difference in ORC between models, with the Aligned VGG-Face matching the curvature structure in the Human Judgment graph more closely than the original VGG-Face as shown in Figure 7 in the Appendix. The distribution of edge curvature values for both the Human Judgment graph and Aligned VGG-Face showed similar value ranges with a mean shifted towards positive curvatures and a left-skewed shape, indicating a higher proportion of positive curvature with a sub-population of edges with strongly negative curvature. In contrast, the original VGG-Face displays a more symmetrical pattern distributed around zero. This implies that the data in the aligned network and the human similarity judgments both form local clusters connected by fewer (negatively curved) edges, indicating distinct structural characteristics compared to the original VGG-Face. To quantify the similarity between the distributions, we computed the Kullback-Leibler (KL) divergence between each VGG-Face version and the Human Judgment graph. Indeed, we found a low KL-divergence between the Human Judgment graph and the Aligned VGG-Face, but a higher divergence with the Original VGG-Face (Table 2).

This structural distinction between the VGG-Face and Aligned VGG-Face was also reflected in the node curvature level. VGG-Face qualitatively shows no strong correlation with human judgments (Figure 7b), which is consistent with the differences in edge curvature distribution (Figure 7c, e). Interestingly, while the edge curvatures between the Human Judgment graph and the Aligned VGG-Face were similar on the distribution level, their agreement fluctuated at individual nodes. That is, the node curvature shows a similar tendency (negative vs. positive) for most nodes; however, the node curvature values diverge for some nodes, indicating that these nodes belong to different local neighborhoods, and therefore, their interpretation in terms of cluster membership seemed to differ between humans and ANNs (sample images are shown in Figure 7a and b in the Appendix).

3.4 Ricci Flow analysis provides insights into structural properties of representational geometries

To detect and compare community structures based on our curvature measurements, we computed the discrete Ricci flow (Ni et al., 2019) in each representational graph. The Ricci flow defines the community structure of a graph by deforming space similarly to heat diffusion: regions with a large positive curvature contract, while those with strong negative curvature expand (see Methods for details). We then analyzed the structural properties of the processed graphs, as summarized in Table 2. In particular, we calculated *conductance* (community separation by the between-and-within edge ratio), *internal edge density* (edges within the community vs. maximum possible), *modularity* (community strength by comparing edge densities with a random graph), and *average embeddedness* (number of shared neighbors for pairs of nodes within a community). These results (Table 2) clearly show that the Human Judgment graph and the graph representation of the Aligned VGG-Face have very similar properties in terms of community structure. In contrast, the Original VGG-Face, despite being trained on face images, does not show a similarly high degree of structure.
Next, we wanted to visually assess the community structure and its relationship to the edge curvatures. As shown in Figure 5, the Human Judgment graph and the Aligned VGG-Face exhibited similar community patterns overall and both show similar tendencies to cluster, e.g., male and female face images (as indicated by the three randomly chosen samples per community). However, the cluster structure in the Aligned VGG-Face is less clearly defined compared to the Human Judgment graph. In line with our other findings, the VGG-Face graph does not show a clear and coherent structure that we could relate to the Human Judgment graph.

3.5 Shared image features across representational geometries

Finally, to better understand the identified clusters in the representational geometries of human judgments, aligned VGG-Face and VGG-Face, we compared facial features that are specific to each cluster. To do this, we performed an ANOVA on the communities detected within each representational space. Each data point (image) was assigned a feature vector representing various facial characteristics, such as eye size, face width, and nose length. This analysis identifies which characteristics contribute to the formation of communities within each representation. Features with a p-value less than 0.05 were selected and the top ten most significant characteristics were identified. Subsequently, the selected features were compared with respect to their overlap between the different representations. Figure 6 shows the overlapping features between the representational geometries. Overall, all three representations share a set of common features, probably reflecting universal characteristics that are essential for clustering faces (e.g., nose shape). Aligned VGG-Face and Human Judgment share the highest number of intersecting features. Interestingly, there is a clear pattern in the significant features that are shared only between the human reference and the Aligned VGG-Face: all three features describe the distances between the cheeks and the chin. This gives us a more detailed view of how facial feature representations might systematically differ between representations in artificial neural network and human perception.

4 Discussion

Our study explored how Ollivier Ricci curvature and Ricci flow reveal the geometric structure of neural representations beyond conventional distance-based approaches like Representational Similarity Analysis. By analyzing curvature at edge level, we captured fine-grained local geometry, while Ricci flow allowed us to uncover higher-level structural patterns, including community structures in representation graphs. Our results demonstrate that aligning an artificial neural network with human behavior reshapes its representational structure in measurable ways. The human-aligned VGG-Face model exhibited curvature properties that more closely mirrored those found in human similarity judgments, particularly in local clustering and connectivity patterns. However, key differences remained, suggesting that behavioral alignment does not fully bridge the gap between cognitive and artificial representations. These findings highlight the potential of discrete geometric tools for studying neural representations and raise important questions about their broader applicability. By moving beyond traditional similarity metrics, we provide a new perspective on how internal representations evolve and align across different systems. We hope this approach sparks further research into the underlying geometry of representations in both artificial and biological neural networks.

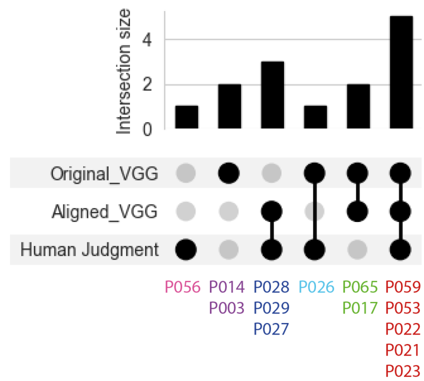

Figure 6: Comparison between the significant features across representations. Feature coding: P056: 'eye size',P014: 'face width', P003: 'nose length', P028: 'distance left cheek and chin', P029: 'average distance left and right cheek to chin', P027: 'distance right cheek to chin', P026: 'distance bottom lip to chin', P065: 'face width-to-height ratio', P017: 'upper face length', P059: 'chin length divided by face length', P053: 'nose shape', P022: 'distance left pupil to lip', P021: 'distance right pupil to lip', P023: 'average distance left and right pupil to lip'.

References

Anonymous. This article has been removed for the review.

SueYeon Chung and Larry F Abbott. Neural population geometry: An approach for understanding biological and artificial neural networks. *Current opinion in neurobiology*, 70:137–144, 2021.

Marin Dujmović, Jeffrey S Bowers, Federico Adolfi, and Gaurav Malhotra. Some pitfalls of measuring representational similarity using representational similarity analysis. *bioRxiv*, pp. 2022–04, 2022.

Hamza Farooq, Yongxin Chen, Tryphon T Georgiou, Allen Tannenbaum, and Christophe Lenglet. Network curvature as a hallmark of brain structural connectivity. *Nature communications*, 10(1): 4937, 2019.

Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021.

Carl Friedrich Gauss. *Disquisitiones generales circa superficies curvas*. Dieterich, 1828.

Adam Gosztolai and Alexis Arnaudon. Unfolding the multiscale structure of networks with dynamical ollivier-ricci curvature. *Nature Communications*, 12(1):4561, 2021.

Richard S Hamilton. Three-manifolds with positive ricci curvature. *Journal of Differential geometry*, 17(2):255–306, 1982.

David K Hammond, Yaniv Gur, and Chris R Johnson. Graph diffusion distance: A difference measure for weighted graphs based on the graph laplacian exponential kernel. pp. 419–422, 2013.

Jürgen Jost. *Bernhard Reimann: On the Hypotheses which Lie at the Bases of Geometry*. Springer, 2016.

Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.

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.

Tong Lin and Hongbin Zha. Riemannian manifold learning. *IEEE transactions on pattern analysis and machine intelligence*, 30(5):796–809, 2008.

Debbie S. Ma, Joshua Correll, and Bernd Wittenbrink. The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47(4):1122–1135, 2015. doi: 10.3758/s13428-014-0532-5.

Yalda Mohsenzadeh, Caitlin Mullin, Benjamin Lahner, and Aude Oliva. Emergence of visual center-periphery spatial organization in deep convolutional neural networks. *Scientific reports*, 10(1): 4638, 2020.

Chien-Chun Ni, Yu-Yao Lin, Feng Luo, and Jie Gao. Community detection on networks with ricci flow. *Scientific reports*, 9(1):9984, 2019.

Yann Ollivier. Ricci curvature of metric spaces. *Comptes Rendus Mathematique*, 345(11):643–646, 2007.

Omkar Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015.

Grisha Perelman. The entropy formula for the ricci flow and its geometric applications. *arXiv preprint math/0211159*, 2002.

Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.

Areejit Samal, RP Sreejith, Jiao Gu, Shiping Liu, Emil Saucan, and Jürgen Jost. Comparative analysis of two discretizations of ricci curvature for complex networks. *Scientific reports*, 8(1): 8650, 2018.

Romeil Sandhu, Tryphon Georgiou, Ed Reznik, Liangjia Zhu, Ivan Kolesov, Yasin Senbabaoglu, and Allen Tannenbaum. Graph curvature for differentiating cancer networks. *Scientific reports*, 5 (1):12323, 2015.

Ilia Sucholutsky, Lukas Muttenthaler, Adrian Weller, Andi Peng, Andreea Bobu, Been Kim, Bradley C Love, Erin Grant, Iris Groen, Jascha Achterberg, et al. Getting aligned on representational alignment. *arXiv preprint arXiv:2310.13018*, 2023.

A APPENDIX

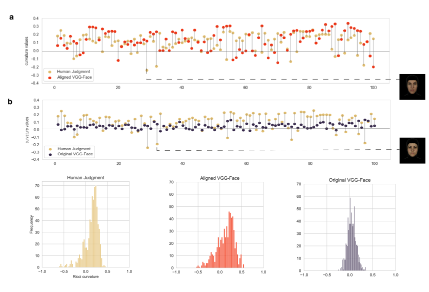

Figure 7: Node curvature distribution (a, b) and edge curvature distribution (c, d, e) across different graphs. (a) and (b) compare node curvatures of the Human Judgment graph (yellow) with Aligned VGG-Face (red) and Original VGG-Face (blue), highlighting divergences. The x-axis represents nodes (images), and the y-axis shows curvature values. Edge curvature distributions are shown for Human Judgment (c), Aligned VGG-Face (d), and Original VGG-Face (e).