

# 1 Quantifying the distribution of feature values over data represented in 2 arbitrary dimensional spaces

3 Enrique R. Sebastian<sup>1#\*</sup>, Julio Esparza<sup>1#\*</sup> and Liset M de la Prida<sup>1\*</sup>

4 Instituto Cajal, CSIC, Madrid 28012, Spain

5 # Co-first authors

6 \* Correspondence: [enrique.rodsebastian@gmail.com](mailto:enrique.rodsebastian@gmail.com) , [esparzaj@cajal.csic.es](mailto:esparzaj@cajal.csic.es) and  
7 [lmprida@cajal.csic.es](mailto:lmprida@cajal.csic.es)

8

## 9 Highlights

- 10 • The Structure Index is a graph-based topological metric
- 11 • It quantifies the distribution of feature values in arbitrary dimensional spaces
- 12 • It can be applied to both scalar and vectorial features
- 13 • When applied to the head-direction neural system, it extracts concordant  
14 information from high- and low-dimensional representations
- 15 • It can be extended to sound and image categorization, expanding the range of  
16 applications

17

## 18 Abstract

19 *Background:* Identifying the structured distribution (or lack thereof) of a given feature over  
20 a point cloud is a general research question. In the neuroscience field, this problem arises  
21 while investigating representations over neural manifolds (e.g., spatial coding), in the  
22 analysis of neurophysiological signals (e.g., auditory coding) or in anatomical image  
23 segmentation.

24 *New method:* We introduce the Structure Index (SI) as a graph-based topological metric to  
25 quantify the distribution of feature values projected over data in arbitrary D-dimensional  
26 spaces (neurons, time stamps, pixels). The SI is defined from the overlapping distribution  
27 of data points sharing similar feature values in a given neighborhood.

28 *Results:* Using model data clouds we show how the SI provides quantification of the  
29 degree of local versus global organization of feature distribution. SI can be applied to both  
30 scalar and vectorial features permitting quantification of the relative contribution of related  
31 variables. When applied to experimental studies of head-direction cells, it is able to  
32 retrieve consistent feature structure from both the high- and low-dimensional  
33 representations. Finally, we provide two general-purpose examples (sound and image  
34 categorization), to illustrate the potential application to arbitrary dimensional spaces.

35 *Comparison with existing methods:* Most methods for quantifying structure depend on  
36 cluster analysis, which are suboptimal for continuous features and non-discrete data  
37 clouds. SI unbiasedly quantifies structure from continuous data in any dimensional space.

38 *Conclusions:* The method provides versatile applications in the neuroscience and data  
39 science fields

## 40 Keywords

41 Neural manifolds; Data science; Image categorization; Temporal series analysis; Machine  
42 learning;

43

## 44 1. Introduction

45 Identifying and quantifying if and how a given feature is structured along a data cloud is a  
46 challenging problem in many fields of science. For instance, in the neurosciences the  
47 temporal evolution of neuronal activity can be pictured as a data cloud on a high-  
48 dimensional space, whose axes are determined by the number of simultaneously recorded  
49 cells or recorded channels (Churchland et al., 2012). Under certain conditions, the high-  
50 dimensional activity can be embedded into 2D or 3D subspaces, where external and  
51 internal variables are reliably visualized (Cunningham and Yu, 2014; Nieh et al., 2021).  
52 Examples include neural manifolds underlying simple motor tasks (Gallego et al., 2017)  
53 and the internal head-direction and grid-cell representational systems (Chaudhuri et al.,  
54 2019; Gardner et al., 2022). In this context, understanding how a given feature is  
55 topologically organized over the manifold sheds light into the representational capacity of  
56 the system under study.

57 Other applications include analysis of multidimensional data that reflect temporal samples,  
58 such as the auditory coding of the spectro-temporal features of natural sounds (Gervain  
59 and Geffen, 2019), image segmentation of multidimensional pixels (Ternes et al., 2022) or  
60 transcriptomic data (Zeisel et al., 2015). In most cases, evaluating the unknown  
61 distribution of a feature over data samples (e.g., the motor reach, spatial and speech  
62 representations, histological categories in an image or neurodevelopmental profiles across  
63 cell-type clusters) relies on the visual inspection of the reduced embedding. Whether the  
64 very same feature had structure in the original high-dimensional space typically remains  
65 unclear.

66 Solving this general-purpose problem can provide solutions for an ample set of scientific  
67 applications. Having the ability to quantify the feature structure in any arbitrary space (i.e.,  
68 that defined by cells, temporal samples, or pixels) may boost applications across fields.  
69 Here, we use the term structure in a loose sense. That is, we say that a variable or feature  
70 is structured along data if it follows some type of non-random distribution in the D-  
71 dimensional representational space.

72 Most methods for structure quantification depend on clustering analysis. However, when  
73 data points do not aggregate in groups or the features do not take discrete or nominal  
74 values, the resulting clusters are not directly interpretable. This renders these methods  
75 suboptimal to many real-world problems where continuous variables and point distributions  
76 are the norm. To overcome this limitation, some studies resort to techniques that depend  
77 on linear correlation metrics, posing limitations for the analysis of more realistic convoluted  
78 distributions (Enns, 2011). Alternative approaches based on decoders tacitly assume that  
79 if a given variable can be decoded from the data cloud, then it must follow some structure.  
80 However, this kind of strategies are highly dependent on the model used, as well as on the  
81 intrinsic dimensionality of the data, being vulnerable to overfitting as sparsity increases  
82 with dimensionality. Crucially, all these approaches provide poor insights about the local  
83 versus global structure of feature representations. It is therefore important to develop a  
84 method that (i) can be applied to non-linear distributions, (ii) generalizes to continuous  
85 features, and (iii) is applicable to arbitrarily high-dimensional spaces.

86 In this paper, we introduce the Structure Index (SI) as a new metric specifically aimed at  
87 quantifying how a given feature is topologically organized along an arbitrary data cloud.  
88 We first demonstrate the principles of our approach with simple model examples and  
89 illustrate how the method can be tuned to quantify the degree of the local/global  
90 organization of feature distribution, as well as its robustness along a broad range of data  
91 characteristics. We show how the SI can be equally applied to vectorial features, in which  
92 more than one variable can be considered. Next, we apply the SI to neural data from  
93 experimental studies of head-direction cells, showing how it can retrieve representation of  
94 different features, which are quantified beyond visual inspection of the neural manifold.

95 Finally, we provide two additional general-purpose examples (sound and image  
 96 categorization), to illustrate the universal application to topological data analysis across  
 97 fields.

98

## 99 **2. Material and Methods**

### 100 **2.1. Definition of the Structure Index (SI)**

101 The SI aims at quantifying the amount of structure present at the distribution of a given  
 102 feature over a point cloud in an arbitrary D-dimensional space. For instance, feature val-  
 103 ues can be distributed in a 2D cloud along a gradient (Fig.1A), or randomly (Fig.1E). Iden-  
 104 tifying such structure without the need for visualization is a major problem in many applica-  
 105 tions, especially for high-dimensional spaces. In the neuroscience field, this problem arises  
 106 for instance when relating the distribution of neuronal activity to external behavioral varia-  
 107 bles projected over the neural manifold.

108 To quantify feature distribution over a point cloud, we first divide the range of values in n-  
 109 equal bins, and then assign each data point to a bin-group according to its feature value  
 110 (Fig. 1A). Note that features can be either categorical (i.e. they may take nominal values  
 111 associated to different categories) or continuous (i.e. they may take values within a scalar  
 112 range). In the case of a discrete feature, each bin-group may correspond to one of the  
 113 possible discrete or nominal values the feature can take.

114 Next, we compute the overlapping between each pair of bin-groups in terms of the k-  
 115 nearest neighbors (Fig.1B). Given two bin-groups,  $\mathcal{U}$  and  $\mathcal{V}$ , we define the overlapping  
 116 score from  $\mathcal{U}$  to  $\mathcal{V}$  ( $OS_{\mathcal{U} \rightarrow \mathcal{V}}$ ) as the ratio of k-nearest neighbors of all the points of  $\mathcal{U}$  that  
 117 belong to  $\mathcal{V}$  in the point cloud space. That is,

$$OS_{\mathcal{U} \rightarrow \mathcal{V}}(k) = \frac{1}{length(\mathcal{U}) \cdot k} \sum_{i=1}^{length(\mathcal{U})} \sum_{j=1}^k \mathcal{O}(i, \mathcal{U}, \mathcal{V}, j) \quad (1)$$

where  $\mathcal{O}(i, \mathcal{U}, \mathcal{V}, j) = \begin{cases} 1, & \text{if } N_j(\mathcal{U}_i, \mathcal{U} \cup \mathcal{V}) \in \mathcal{V} \\ 0, & \text{otherwise.} \end{cases}$

where  $N_j(\mathcal{U}_i, \mathcal{U} \cup \mathcal{V})$  is the  $j_{th}$  nearest neighbor of point  $\mathcal{U}_i$  in the set  $\mathcal{U} \cup \mathcal{V}$

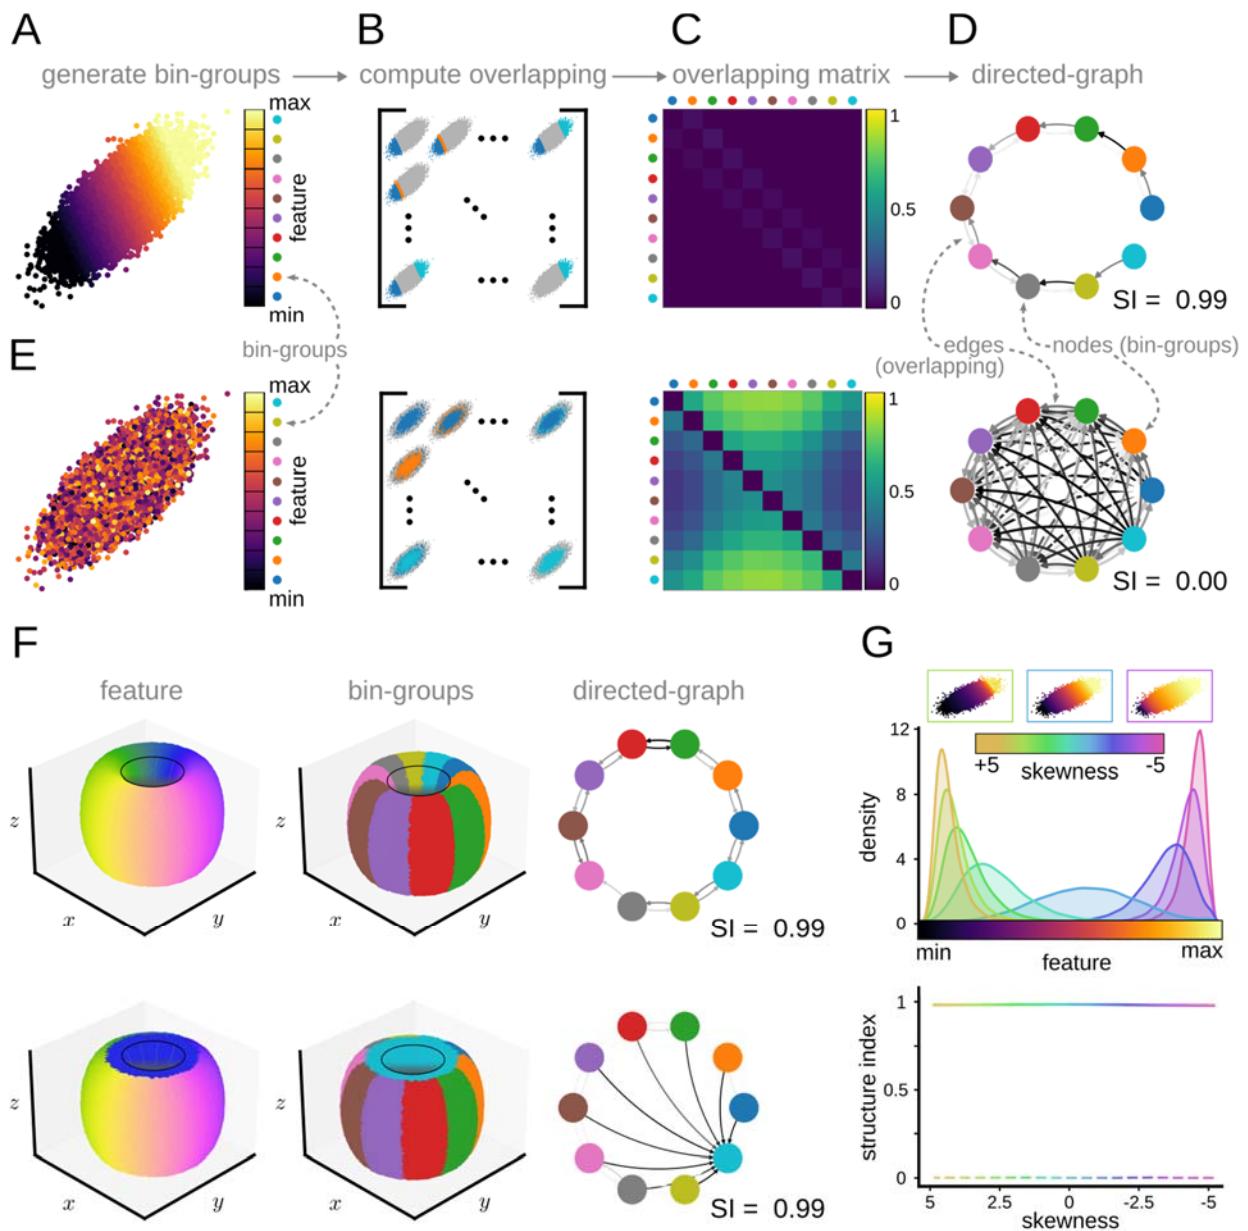
118 Note that the definition of nearest neighbors is determined by the distance metric used  
 119 (i.e., Euclidean distance, geodesic distance, etc.). Computing the overlapping score for  
 120 each pair of bin-groups ( $\mathcal{U}_a$  and  $\mathcal{V}_b$ ) yields an adjacency matrix ( $\mathcal{M}_{nxn}$ ) whose entry  $(a, b)$   
 121 equals  $OS_{\mathcal{U}_a \rightarrow \mathcal{V}_b}$  (Fig. 1C).  $\mathcal{M}$  can be thought as representing a weighted directed graph,  
 122 where each node is a bin-group, and the edges represent the overlap (or connection)  
 123 between them (Fig.1D).

124 Finally, we define the Structure Index as 1 minus the mean weighted degree of the nodes  
 125 after scaling it:

$$SI(\mathcal{M}) = 1 - \left( \frac{2}{n} \sum_i^n \sum_j^n \mathcal{M}_{i,j} \right) \quad (2)$$

126 Under this definition, for a uniform random distribution the overlapping of any two nodes  
 127 would be equal to 0.5 and therefore, the mean degree of the nodes of such distribution

128 would also be 0.5. Thus, the Structure Index would take a value of 0 for a random distribution.  
 129 In contrast, the mean degree of the nodes of a perfectly separated distribution would  
 130 be 0 and thus, the SI would be 1. Therefore, the SI ranges between 0 (random feature dis-  
 131 tribution, fully connected graph) and 1 (maximally separated feature distribution, non-con-  
 132 nected graph; Fig. 1D).



133

**Fig. 1. Illustration of the concepts behind the definition of the Structure Index (SI).**  
**A**, Feature gradient distribution in a 2D-ellipsoid data cloud. Each point in the data cloud is assigned to a group associated with a feature bin value (bin-group). **B, C**, Next, the overlapping matrix between bin-groups is computed according to equation 1. **D**, The overlapping matrix represents a weighted directed graph between bin-groups, where structure (overlapping, clustering, etc..) can be quantified using the SI from 0 (random, equivalent to full overlapping) to 1 (maximal separation, equivalent to zero overlapping between bins). **E**, The case of a feature randomly distributed over a 2D data cloud. **F**, Different feature distribution yielding the same SI but different weighted directed graph. **G**, Lack of effect of the skewness of feature values on the SI.

134

135 By definition, the SI is agnostic to the type of structure (e.g., gradient, patchy, etc.) since  
 136 bin-groups do not need to follow any specific arrangement. Instead, the weighted directed  
 137 graph provides additional insights. Fig. 1F shows the example of two different distributions

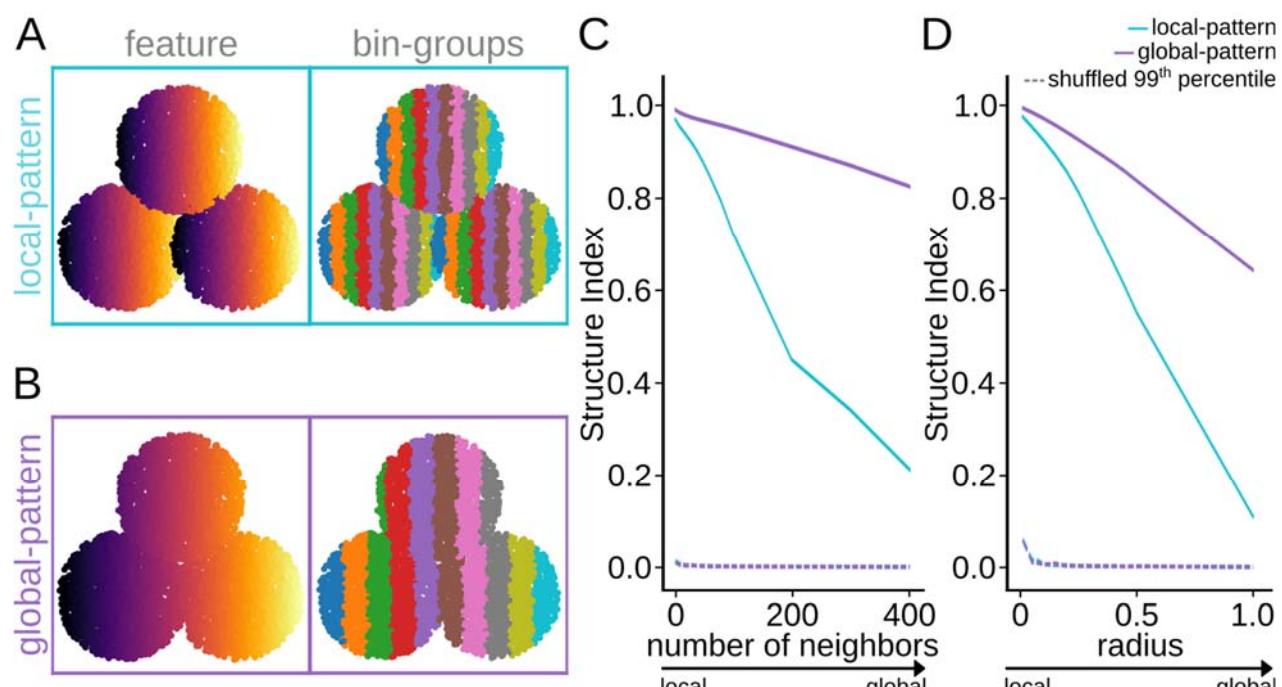
138 with similar SI but different graphs. Of note, the skewness of feature values has little im-  
139 pact on SI, being robust for a wide range of statistical properties (Fig. 1G).

140 Note that this metric can be applied to any type of data represented in arbitrary D-dimen-  
141 sional spaces (cells, time series, pixels). Our approach is not in direct competition with the  
142 many methods that use cluster analysis or topological decoding. Rather, it generalizes at a  
143 class of distributions (i.e., continuous distributions) where clusters typically fail to apply.  
144 Our definition of SI and the equivalent graph makes this metric general enough to ease a  
145 range of application, which we will illustrate along the Results section.

146

## 147 2.2. Parameter dependence of SI on the neighborhood size

148 To compute the overlapping between each pair of bin-groups, the SI looks at the properties  
149 of the  $k$ -nearest neighbors of each point. For a low number of neighbors, the overlapping  
150 is computed in the close vicinity of each point, thus being biased towards the local  
151 distribution. As the number of neighbors increases, the SI tends to better account for the  
152 global structure. This dependence of the SI on the number of neighbors can be exploited  
153 to infer information about the local versus global organization of data features.



154 **Fig. 2. Parametric dependence of SI on the number of neighbors.** **A**, A local pattern is  
155 simulated in 2D by projecting feature values differently along the data cloud (9000 points).  
Note local structure between bin-groups. **B**, The same 2D data cloud exhibiting a global  
distribution of feature values. **C**, Dependency of the SI values as a function of the number  
of neighbors can help to identify the local versus the global distribution trends. Data is  
tested against shuffled distribution of feature values (99<sup>th</sup> percentile). **D**, Same as in C, but  
as a function of the radius.

155

156

157 Figure 2 shows two different feature distributions over the same data cloud. In the local  
158 pattern, feature values replicate at the different regions of the data cloud and so bin-  
159 groups reflect such organization (Fig. 2A). In contrast, in the global pattern, feature values  
160 follow a general trend (Fig. 2B). By evaluating the evolution of the SI as a function of the  
161 number of neighbors, the trade-off between local and global structure can be quantified.

162 For the local pattern, overlapping between bin-groups increases as the number of  
163 neighbors increase, and thus the SI sharply decreases. On the contrary, for the global  
164 pattern, the overlapping is less sensitive to the number of neighbors, and therefore the SI  
165 decreases smoothly (Fig. 2C). As expected, the SI of the shuffled distribution equals 0  
166 independently on the number of neighbors. Thus, by tuning the number of neighbors, one  
167 can effectively change the sensitivity of SI to better detect local or global structures.

168 For data clouds with highly uneven density, the SI presents the option of setting a radius  
169 size ( $r$ ) instead of the number of neighbors. In such a case, the neighbors of a point are  
170 set to be all points that fall within a given distance  $r$ . That is, equation (1) becomes:

$$OS_{\mathcal{U} \rightarrow \mathcal{V}}(r) = \frac{1}{length(\mathcal{U})} \sum_{i=1}^{length(\mathcal{U})} \left( \frac{1}{length(\beta(r)^i)} \sum_{j=1}^{length(\beta(r)^i)} \vartheta(\mathcal{U}, \beta(r)^i, j) \right) \quad (3)$$

where  $\vartheta(\mathcal{U}, \beta(r)^i, j) = \begin{cases} 0, & \text{if } \beta(r)_j^i \in \mathcal{U} \\ 1, & \text{otherwise.} \end{cases}$

where  $\beta(r)^i = \{x \in \mathcal{U} \cup \mathcal{V}: |\mathcal{U}_i - x| < r\}$

171 In such cases, the radius still helps to control for the trade-off between local and global  
172 structure, with smaller values making the SI more sensitive to local, and larger values  
173 being more sensitive to global structure (Fig. 2D).

174

### 175 2.3. Datasets

176 In this study, we used different datasets to evaluate SI performance. For the parameter  
177 study, we created objects (2D-ellipsoids, balls and spheres) using the corresponding  
178 mathematical equations. For the object lamp, we used the model from the ModelNet40  
179 dataset, which is publicly available at <https://github.com/antao97/PointCloudDatasets>.  
180 Different feature value distributions were created over these objects and used to evaluate  
181 SI performance. By default all objects were created with 40,000 data points, except  
182 otherwise reported.

183 To study neural manifold representations, a publicly available head-direction dataset was  
184 used (<http://crcns.org/data-sets/thalamus/th-1>; doi:10.6080/K0G15XS1) (Peyrache et al.,  
185 2015). We chose this dataset because the neural manifold organization of head direction  
186 angles was recently validated (Chaudhuri et al., 2019), excluding any confounding in the  
187 ability of the SI to extract structure. Moreover, as we will show in the Results section, using  
188 these data allowed us to illustrate the capacity of the SI to quantify structure in the original  
189 space, which was not tested in the aforementioned reference due to lack of  
190 computationally efficient available methods. We used all data available to build the 3D  
191 neural manifold as reported in (Chaudhuri et al., 2019) using Isomap (Tenenbaum et al.,  
192 2000). We also built the representations in the original space using single cell data (one  
193 axis per cell), yielding different high-dimensional spaces per mouse (n=6; mouse12 -  
194 120806: 37 cells; mouse17-130130: 29 cells; mouse20-130520: 11 cells; mouse24-  
195 131216: 10 cells; mouse25-140130: 10 cells; mouse28-140313: 22 cells). Information  
196 about brain states was used to separate the neural manifolds from awake and sleeping  
197 periods, with sleep classified as Slow-Wave Sleep (SWS) or Rapid Eyes Movement (REM)  
198 sleep (Chaudhuri et al., 2019).

199 To evaluate the application of SI to temporal data, we opted to use musical notes given  
200 their similarities with electroencephalographic waveforms (Baier et al., 2007). An additional

201 advantage is that musical notes are directly interpretable, allowing us to focus in  
202 evaluating SI performance. We chose the NSynth dataset (Engel et al., 2017), which  
203 contains over 300,000 musical notes produced by around 1000 different acoustic,  
204 electronic or synthetic instruments, including the human voice. This dataset is available on  
205 the TensorFlow Magenta project at <https://magenta.tensorflow.org/datasets/nsynth>.  
206 Different features (source, instrument family, pitch and velocity) characterize each musical  
207 note. Each note consists in 4 seconds of monophonic 16 kHz audio snippets at five  
208 different velocities. For analysis, we downsampled the original snippets to 1.2 kHz  
209 resulting in 4800 time-stamps, which were used to build the high-dimensional space (one  
210 point per note). To comply with the Nyquist–Shannon sampling theorem, audio snippets  
211 with an associated pitch higher than 73 MIDI (equivalent to 554 Hz) were discarded.  
212 Binary features were not included in the analysis. For statistical testing, the dataset was  
213 divided in 5 equivalent batches, which were analyzed both in the original and the 3D-  
214 reduced space using Uniform Manifold Approximation and Projection (UMAP) (McInnes et  
215 al., 2018).

216 Finally, to provide examples of image analysis using SI we resorted to the bird species  
217 problem, given its application in fine-grained image recognition. Similar as above, this  
218 allowed us to focus in evaluating the performance of SI without requiring any particular  
219 interpretation. To this purpose, we used SI the 100-bird species dataset created by Gerald  
220 Piosenka, which is hosted on the Kaggle platform  
221 (<https://www.kaggle.com/datasets/gpiosenka/100-bird-species>; date of download July 28<sup>th</sup>,  
222 2022). The dataset consists of more than 70000 RGB images of 450 bird species. Images  
223 are 224 x 224 pixels x 3 color (jpg format) annotated by species name. For analysis, we  
224 downsampled images to 56x56 pixels x 3 colors, resulting in a 9408-dimensional space,  
225 where each axis is the value of a pixel (one point per image). To expand the number of  
226 features associated to each image, we performed an automated data scraping from  
227 Wikipedia, so that for each bird species we also extracted information about geographical  
228 distribution (continents), as well as the scientific order and family, using the Python library  
229 'Wikipedia' (<https://pypi.org/project/wikipedia/>). The dataset was divided in 2 equivalent  
230 batches.

231 Statistical analysis of different instances of each dataset was performed using one- or two-  
232 way ANOVAs followed by Student t-tests or equivalent. Spearman correlation was used to  
233 evaluate relationship between variables, which were fitted by exponential curves.

234

## 235 **2.4. Computational resources**

236 All simulations and analysis were performed in Python 3.8.13 using personal computer  
237 workstations (Intel Xeon CPU E5-2620 v4 @ 2.10GHz processor with 16 cores, 64GB  
238 RAM memory, GeForce GTX 1080 Ti GPU with 11GB memory and 0.355 TFlops for  
239 double precision). Whenever required, the supercomputer cluster Artemisa  
240 (<https://artemisa.ific.uv.es/web/content/nvidia-tesla-volta-v100-sxm2>) was used to  
241 accelerate calculation and parametric analysis (NeuroDIM Project).

242

## 243 **2.5. Data and code availability**

244 All data used in this study is publicly available (see section 2.3). Code is deposited at  
245 [https://github.com/PridaLab/structure\\_index](https://github.com/PridaLab/structure_index).

246

247

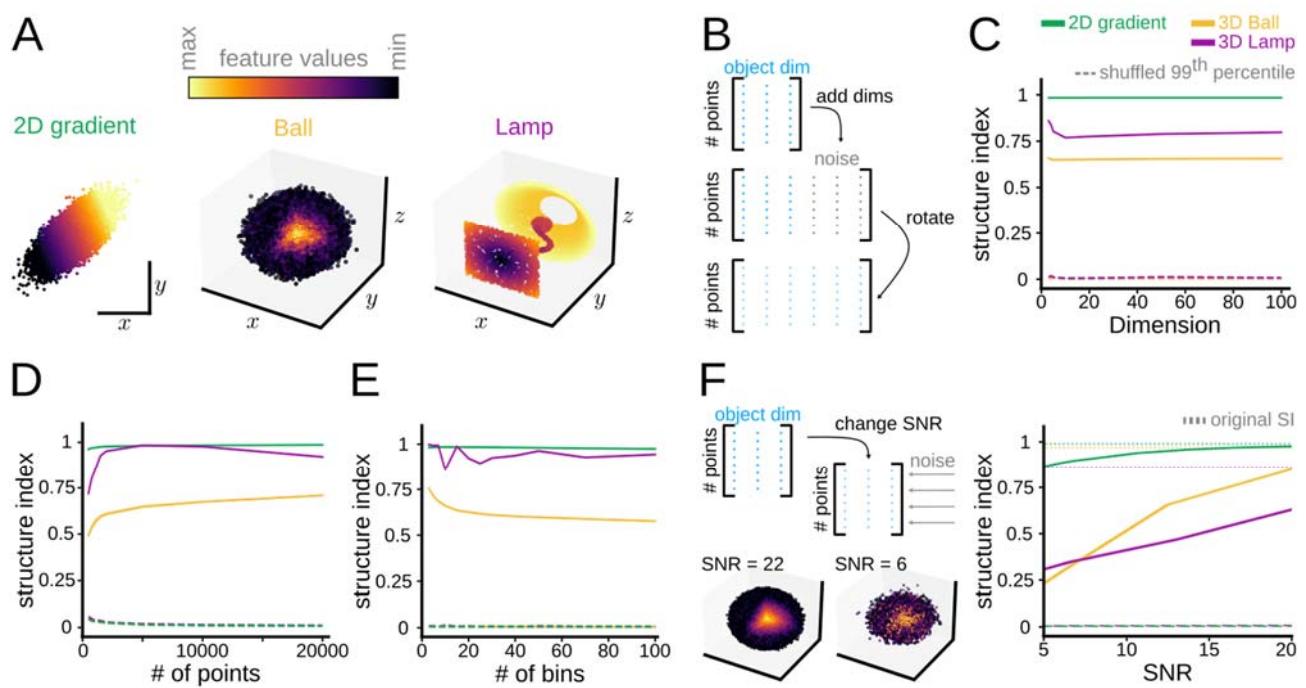
## 248 3. Results

### 249 3.1. SI quantifies the topological distribution of scalar feature values

250 Before applying our method to the study of neural data, we used toy model data to  
251 illustrate its performance and robustness to a wide range of point cloud characteristics. We  
252 generated 3 independent toy-models, including a 2D linear gradient (40,000 points), a 3D  
253 solid ball (40,000 points) where the feature was distributed along the radius, and a 3D  
254 lamp (32,000 points) whose feature varies in terms of the three axes (Fig. 3A).

255 To test for the stability of the method, we began by expanding these models in an  
256 increasing number of dimensions while adding white noise and then rotating the object in  
257 the extended space (Fig. 3B). By doing so, we maintained the intrinsic dimension of the  
258 object but spread the information along all dimensions. The SI showed a consistent  
259 response while increasing dimensionality (Fig. 3C). Importantly, the SI performed smoothly  
260 for a wide range of points in the cloud when examined in 2 dimensions (Fig. 3D).

261 In terms of the number of bin-groups used when computing the SI, there are two potential  
262 cases. For discrete or nominal features, the number of bin-groups is determined by the  
263 unique values the feature can take, so that there is a bin-group per discrete value. When  
264 dealing with continuous feature values, the number of bin-groups becomes a heuristic  
265 choice, which can be informed by statistical analysis. It should be large enough so that the  
266 continuity of the feature values is fully captured, but small enough so that there is a  
267 reasonable number of points assigned to each bin-group. While the SI performs  
268 consistently for a range of bin-groups (Fig. 3E), the topological characteristics of the data  
269 cloud may have different impacts that should be examined for each application.



270

**Fig. 3. Robustness of SI under a wide range of data-cloud characteristics.** **A**, Three toy models used to evaluate performance of SI (40,000 points for the gradient and the ball; 32,000 points for the lamp). **B**, Objects in A were embedded into spaces of increasing dimensionality by adding noise and then rotating. **C**, Dependence of the SI on the embedded dimensionality for the three toy models. **D**, Effects of the number of points in the data cloud as examined in a 2D space. **E**, Effects of the number of bin-groups on the SI for the three toy models. **F**, Effect of different levels of signal-to-noise ratio on SI.

271

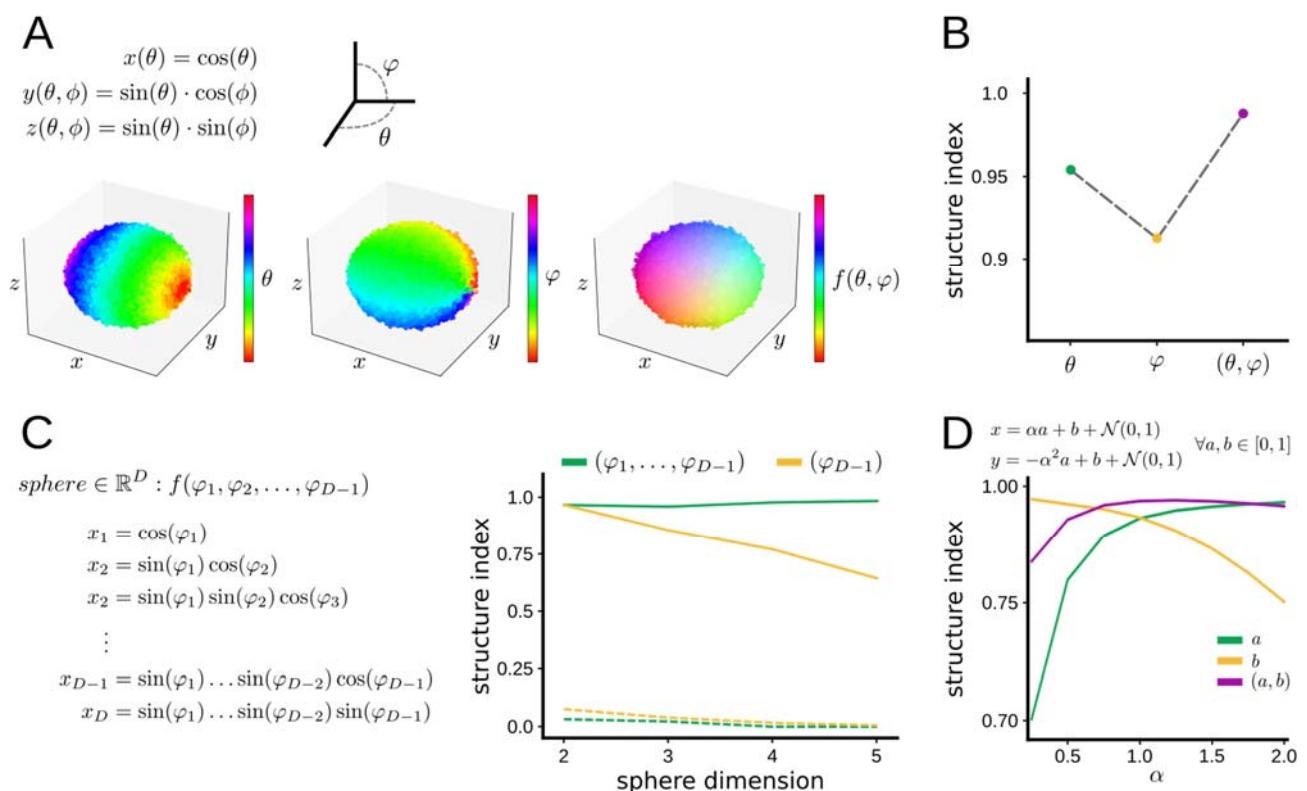
272

273 Finally, we studied the sensitivity of the SI to different levels of noise in terms of the Signal  
 274 to Noise Ratio (SNR), as defined in (Zeng et al., 2019). To this purpose, we introduced  
 275 Gaussian noise across all existing dimensions (Fig. 3F, left). While noise has effect in  
 276 structure, the SI was able to capture the trends even when introducing high levels of noise  
 277 into the point clouds (Fig. 3F, right). This renders the SI suitable for testing a wide variety  
 278 of experimental data sets.

279

280 **3.2. Evaluating the structure distribution of vectorial features**

281 The definition of bin-groups used in the SI can be extended to vectorial features which  
 282 integrate values from several characteristics. For example, a feature vector  $(A, B)$  can be  
 283 created from two scalar features,  $A$  and  $B$ , taking values along a continuous scale. In such  
 284 case, bin-groups can be defined by the upper and lower bound for both  $A$  and  $B$ . Thus, a  
 285 point  $(p)$  in the cloud will fall within the bin-group  $\mathcal{U}$  if and only if both entries of the  
 286 associated feature vector fall within the common range.



287

**Fig. 4. Evaluating structure of vectorial features.** **A**, 3D sphere defined by trigonometric equations depending on angles  $\theta$  and  $\varphi$  (40,000 points). Feature values can be defined for each angle independently,  $\theta$  or  $\varphi$ , and for both together in a vectorial form  $(\theta, \varphi)$ . **B**, SI for each individual angle values and for the vectorial angle. **C**, A  $D$ -dimensional sphere is defined by trigonometric equations depending on  $D-1$  angles ( $8^{D \times N}$  points, with  $N=40,000$  points to keep cloud density over  $D$ -dimensional spaces). The plot at right shows the dependence of SI on the sphere dimension, computed for the  $D-1$  angle alone, and for all angles in vectorial form. Dashed lines indicate results from shuffled distribution values (99<sup>th</sup> percentile). **D**, Behavior of SI for a feature defined in 2D according to the equation shown (20,000 points).

288

289

290 To illustrate the case, we generated a point cloud sampled from a sphere of unitary radius  
 291 using two angles  $\theta, \varphi$ , with added Gaussian noise in 3D. Mathematically, the  $x$ -coordinate  
 292 of a sphere is defined by the cosine of  $\theta$ , while the  $y$ - and  $z$ -coordinates follow

293 trigonometric relationships between both  $\theta$ ,  $\varphi$  (Fig. 4A). Thus, a feature defined by  $\theta$  and  $\varphi$   
294 independently will distribute differently along the sphere than a vectorial one integrating  
295 both angles (Fig. 4A). By definition, the structure of each angle separately should be lower  
296 than the vectorial angle ( $\theta$ ,  $\varphi$ ). Moreover, given that the  $x$ -coordinate is completely defined  
297 by  $\theta$ , we would expect more structure for  $\theta$  than for  $\varphi$ . Consistently, the SI behaved as  
298 expected, with the lowest SI value obtained for  $\varphi$ , then for  $\theta$ , and the highest value for both  
299 angles as a vectorial feature ( $\theta$ ,  $\varphi$ ) (Fig. 4B).

300 To evaluate the generalization of this behavior to vectorial features of any dimension, we  
301 generated point-clouds sampled from D-dimensional spheres according to the equation  
302 shown in Fig. 4C (left). For each point cloud in D-dimensional space, we computed the SI  
303 for both the  $D - 1$  angle used to generate the sphere and all angles together as a feature  
304 vector (Fig. 4C, right). As predicted, the SI obtained when introducing all angles as a  
305 vector remained stable for all D-dimensional spheres. However, when only the  $D - 1$   
306 angle is considered, the SI declined as the dimensionality of the sphere increased. This  
307 reflects the fact that as the dimensions become larger, a lower percentage of coordinates  
308 depend on the  $D-1$  angle, and thus the position of a given point is less dependent on it.

309 This property of the SI can be exploited to examine the interdependence between distinct  
310 interrelated features. For instance, we created a 2D cloud where the position of each point  
311 depends on two features:  $a$ ,  $b$  (Fig. 4D; see equation). While the impact of  $b$  in the position  
312 of the points was constant, the impact of  $a$  could be tuned by increasing or decreasing the  
313 parameter  $\alpha$ . We proceeded by computing the SI of the scalar features  $a$ ,  $b$ , and for the  
314 vector ( $a$ ,  $b$ ) using a range of  $\alpha$  values (Fig. 4D). The maximum SI( $b$ ) was obtained for  $\alpha$   
315 equal to zero (as the position of the points was completely defined by  $b$ ) and decreased  
316 consistently as  $\alpha$  increased. In contrast, SI( $a$ ) increased with  $\alpha$  as expected. Interestingly,  
317 SI( $a$ ,  $b$ ) was lower than SI( $b$ ) for low  $\alpha$  values (as the points are completely defined by  $b$ ,  
318 the structure of ( $a$ ,  $b$ ) is lower than that of  $b$ ). However, SI( $a$ ,  $b$ ) rapidly increased with  $\alpha$   
319 reaching a plateau at maximum structure around 1 when both  $a$  and  $b$  equally contributed  
320 to the position of points.

321 These examples illustrate the capability of the SI to capture the structure of vectorial  
322 features, opening new avenues to study the relative impact and dependency between  
323 mathematically or experimentally related variables.

324

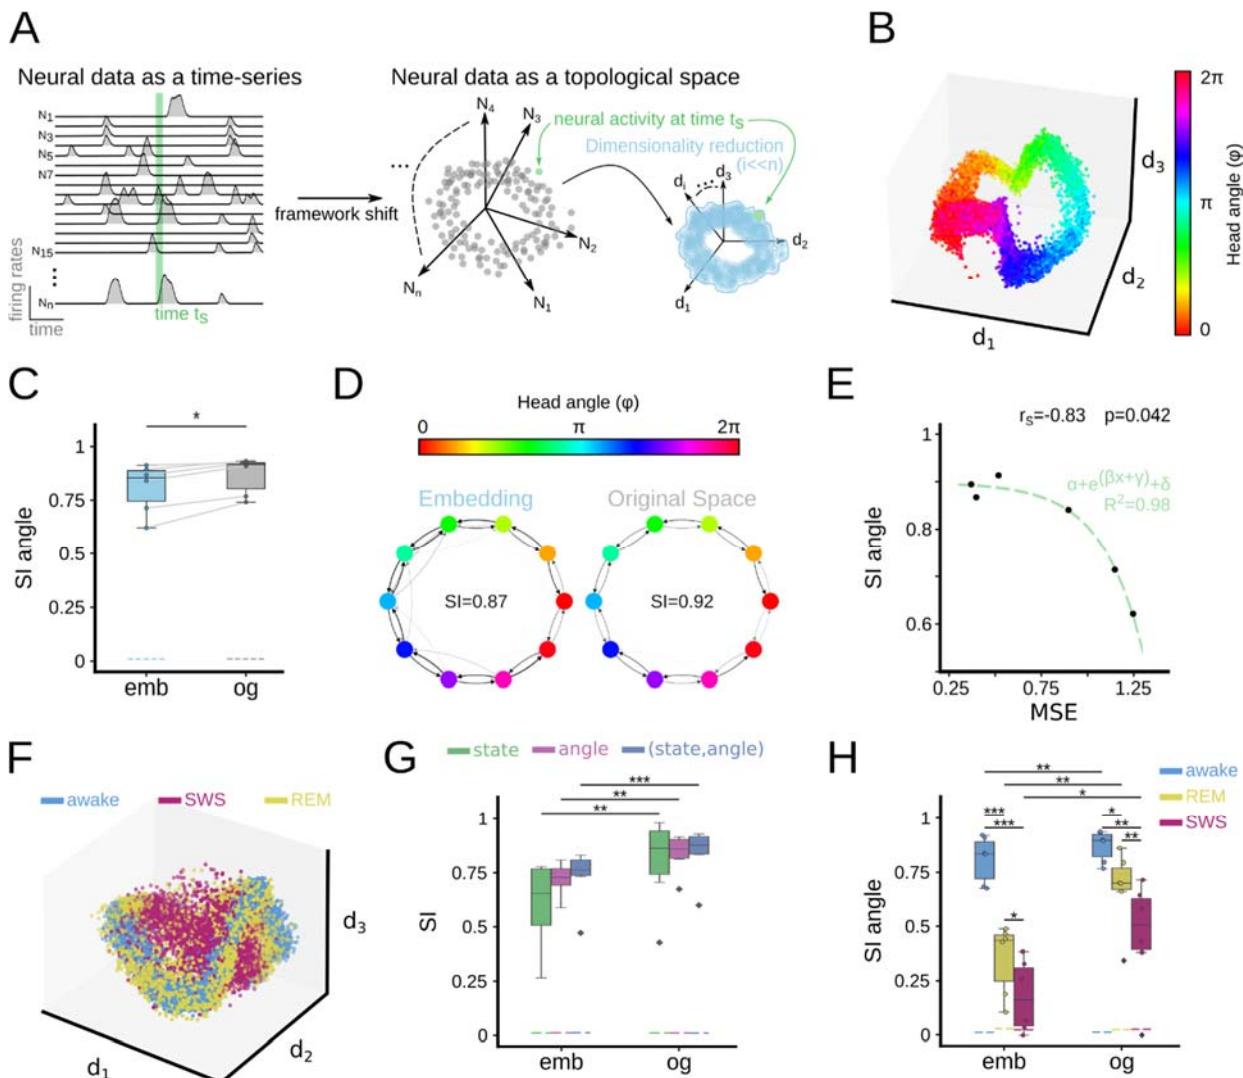
### 325 **3.3. Application to neural manifold representations**

326 Having established the main readouts expected from the SI metric, we sought to apply it to  
327 the study of neural manifolds. To illustrate the effectiveness of the approach, we chose a  
328 public dataset of extracellular recordings from multi-site silicon probes in the anterodorsal  
329 thalamic nucleus (ADn) of freely moving mice (Peyrache et al., 2015). This dataset has  
330 been recently used to demonstrate the intrinsic attractor manifold of the mammalian head-  
331 direction system (Chaudhuri et al., 2019), permitting direct testing of the ability of SI to  
332 extract feature structure.

333 In their study, Chaudhuri et al. showed that neural activity of N-simultaneously recorded  
334 ADn neurons of mice foraging in an open environment was constrained to a ring-shaped  
335 3D manifold (Fig. 5A, right;  $n=6$  mice), which they visualized in 3D using Isomap (Fig. 5B).  
336 Therefore, structure was implicitly expected at least in the low dimensional representation.

337 When computing the SI of the head-direction angle over the neural manifold (3 neighbors),  
338 we obtained a high structure concordant with visual inspection of the embedding (Fig. 5C;  
339 blue). Importantly, in Chaudhuri et al. the analysis was mainly restricted to the 3D reduced  
340 space. Since the SI can be applied to an arbitrarily high-dimensional data cloud, we also  
341 evaluated the structure of the head-direction angle over the original N-dimensional neural  
342 space. Interestingly, the SI for all animals in the original space was slightly higher than in

343 the low dimensional embedding (Fig. 5C; grey, paired sample t-test  $p=0.011$ ). Visualization  
 344 of individual weighted directed graphs from the high- and the low-dimensional  
 345 representations confirm similar organization (Fig.5D).



346

347

**Fig. 5. Using SI to evaluate neural representations.** **A**, In the neural manifold framework, firing rates from  $N$ -neurons at a given time ( $t_s$ ) are represented in an  $N$ -dimensional Euclidean space. Activity is constrained in a subspace which can be retrieved using dimensionality reduction approaches ( $d$ -dimension). **B**, 3D neural manifold computed from the head direction system by Chaudhuri et al., with the head direction angle projected over the data cloud. **C**, SI of the head direction angle in the original (og) and 3D-embedded representations (emb). Dashed lines indicate results from shuffled distribution values (99<sup>th</sup> percentile). **D**, Example of the weighted directed graphs from the same mouse (mouse12 - 120806) in the original and in the low-dimensional embedding. Note similar organization. **E**, Relationship between the SI and the mean square error (MSE) of the decoder trained by Chaudhuri et al. in the reduced space. Fitting curve parameters:  $\alpha = -0.12$ ;  $\beta = 4.47$ ;  $\gamma = -4.7$  and  $\Delta = 0.9$ ; tested significant at  $p < 0.05$ . **F**, Head direction data plotted over the 3D embedding for awake, REM, and SWS states separately. **G**, SI for states and angles separately, and for both features together as expressed in a vectorial form. Results are shown for both the original (og) and the reduced space (emb). ANOVA effects for space ( $F(2,1) = 8.2$ ,  $p = 0.007$ ) but not for feature nor interaction. Post-hoc tests: \*,  $p < 0.05$ ; \*\*,  $p < 0.01$ ; \*\*\*,  $p < 0.001$ . **H**, SI of the head direction angle for each state separately both in the original (og) and the reduced space (emb). ANOVA effects for state ( $F(2,1) = 83.5$ ,  $p < 0.0001$ ), space ( $F(2,1) = 25.7$ ,  $p < 0.0001$ ) and interaction. Post-hoc tests: \*,  $p < 0.05$ ; \*\*,  $p < 0.01$ ; \*\*\*,  $p < 0.001$ .

348

349 In their original work, the author parametrized the manifold with splines of matching  
350 topology and used them to decode the represented latent variable (head-direction angle).  
351 We thus tested how the decoder performance (measured as the mean square error of  
352 predictions per mice) related to the head-direction information structured in the data. We  
353 found that the SI correlated with the decoder error (Spearman correlation -0.83,  $p=0.042$ ),  
354 following an exponential decay relationship ( $R^2=0.98$ ; Fig.5E). That is, manifolds with lower  
355 decoding errors had higher head-direction structure as measured by SI.

356 Given the nature of the data, we wondered whether the head direction representation can  
357 be retrieved during REM as well as in SWS states (Senzai and Scanziani, 2022). To tackle  
358 this question, we resorted to the same dataset but used all neural data to compute the 3D  
359 manifold as reported in Chaudhuri et al. Indeed, when points of the manifold were color-  
360 coded according to the state (awake, REM, nREM) we noted some stratification, which  
361 could be quantified using the SI (Fig. 5F). The SI returned structure for both the animal  
362 state and the head-direction angle, with higher values in the original than in the reduced  
363 space (Fig.5G; ANOVA effects for space,  $F(2,1)=8.2$ ,  $p=0.007$ , but not for feature nor  
364 interaction). Interestingly, structure was higher when using a vectorial feature consisting on  
365 the state and the head-direction angle together, indicating that there may be some  
366 interdependency between them (Fig. 5G).

367 Finally, we computed the SI of the head-direction angle for each state separately (Fig. 5H).  
368 SI was maximal in awake conditions. In general, data represented in the original space  
369 provided more structure than in the manifold embedding (ANOVA effects for state  
370 ( $F(2,1)=83.5$ ,  $p<0.0001$ ), space ( $F(2,1)=25.7$ ,  $p<0.0001$ ) and interaction). Moreover,  
371 whereas REM and SWS yielded a low SI in the manifold, it was significantly higher in the  
372 original space, indicating that information was lost while reducing dimension. Thus, being  
373 able to evaluate neural activity in the original space using the SI might provide new  
374 insights into the representative capacity during multiple brain states.

375

### 376 **3.4. Application to arbitrary D-dimensional spaces (temporal samples and images)**

377 Finally, we applied the SI to two additional general-purpose examples: sound (temporal  
378 data) and image categorization (pixels), illustrating the usefulness of the SI metric for  
379 analysis of different types of data and across fields.

380 For temporal data, we resorted to musical notes given similarities with electro-  
381 encephalographic waveforms (Baier et al., 2007). In addition, using this dataset allowed us  
382 to focus in evaluating SI performance directly, given direct interpretability of musical notes.  
383 Data consisted on 4 seconds of musical notes of different pitch and velocity downsampled  
384 as 4800 time stamps. They were produced by different instruments (including the human  
385 voice) using acoustic, electronic or synthetic sources. Instruments are annotated as  
386 belonging to different families. Notes were represented in the 4800-dimensional space  
387 (Fig.S1A), and the SI was calculated both locally (using 3 neighbors) and globally (60  
388 neighbors). In general, data showed a higher local than global structure (Fig.S1B; ANOVA  
389 effects  $F(3,1)=4.0$ ,  $p<0.0001$ ). We found that the pitch provided maximal structure,  
390 followed by the source and family, as confirmed by the weighted directed graph returned  
391 by the overlapping of instrument families (Fig.S1C; ANOVA effects for features  
392  $F(3,1)=12.0$ ,  $p<0.0001$ ). Reducing data to 3D allowed for visualization of these trends  
393 (Fig.S1D), and provided similar SI figures as for the original space (Fig.S1E).

394 For image analysis, we chose using images of one hundreds bird species, typically  
395 exploited in fine-grained recognition problems. Data consist on RGB images (56x56x3  
396 pixels) so that each image was represented as a point in a 9408-dimensional space  
397 (Fig.S2A). Birds were classified as belonging to different species, continent, scientific order

398 and family. SI was maximal for bird species, followed by family and order (Fig.2SB). We  
399 noted that continents provided the lower structure, potentially reflecting migratory habits  
400 and/or species diversification. Visualization of images showing maximal and minimal  
401 overlapping values confirmed that the SI successfully captured the underlying structure of  
402 the data (Fig. S2C,D).

403 These examples illustrate how the SI method successfully operates in arbitrary D-  
404 dimensional spaces, allowing for a range of multidisciplinary applications in neuroscience,  
405 as well as across other research fields.

406

#### 407 4. Discussion

408 With the development of a graph-based topological metric (SI), we have enabled accurate  
409 quantification of the structure of feature distributions. The approach is not constrained by  
410 the dimensionality of the space and is robust to a wide range of data and feature  
411 characteristics. Importantly, the SI not only quantifies the “amount” of structure of scalar  
412 feature values represented over a point cloud, but it can also provide insights into the  
413 topological distribution of the feature by looking at the overlapping directed graph.  
414 Importantly, the SI is not limited to Euclidean spaces, as one can define the  $k$ -closest  
415 neighbors in terms of different distance metrics. For instance, the SI allows for the use of  
416 geodesic distance and cosine distance among others.

417 A common issue in current dimensionality reduction methods is being able to capture the  
418 global structure without deforming local relationships. Indeed, most dimensionality  
419 reduction methods have a parameter to control that tradeoff (e.g., the number of  
420 neighbors). Here, we demonstrated that the SI can be tuned to better detect local vs global  
421 structure by changing the number of neighbors (or equivalently the radius) used to  
422 compute overlapping between bin-groups. Thus, the SI can be used not only to quantify  
423 the structure in the original space, but also to evaluate the quality of the dimensionality  
424 reduction by looking at how much structure has been preserved both locally and globally  
425 while reducing from the high- to the low-dimensional representations.

426 As demonstrated above, the SI can be extended to vectorial features, expanding the range  
427 of applications. Note that a vectorial label can be created by grouping multiple  
428 scalar/categorical features, or by integrating several related variables. By doing so, the SI  
429 allows for the study of how different features interact with each other, allowing for a deeper  
430 understanding of how data structure is determined. This may ease data-driven discoveries  
431 of latent interaction between experimental features, which cannot be established a priori.

432 In this context, we have applied the SI to study the representation capability of the head-  
433 directional system (Peyrache et al., 2015). By using data from a previous study that  
434 demonstrated low dimensional representations of the head-direction angle (Chaudhuri et  
435 al., 2019), we have shown that the SI captures structure both in the lower dimensional  
436 manifold and in the original space. Moreover, by applying the metrics to awake, REM, and  
437 nREM states, we showed that the head-direction representation is preserved during REM,  
438 providing additional interpretation. This is consistent with recent data supporting mental  
439 replay of head-direction angles during REM sleep (Senzai and Scanziani, 2022).

440 The SI method can be applied to the study of temporal data expressed in high-dimensional  
441 spaces. By representing temporal events in the space built from the individual time  
442 stamps, electrophysiological signals can be analyzed with state space methods (Durbin  
443 and Koopman, 2001; Gervasoni et al., 2004; Reichinnek et al., 2010). Applying the SI to  
444 these representations may thus allow for new strategies for the analysis of the spectro-  
445 temporal organization of brain oscillations and/or perception (Valero et al., 2017; Lopes-  
446 Dos-Santos et al., 2018; Gervain and Geffen, 2019; Navas-Olive et al., 2020; Douchamps

447 et al., 2022). Similarly, the SI permits image quantification and categorization in the service  
448 for fine-grained image recognition problems applicable to several research fields.

449 As topological and high-dimensional analysis become the norm in the neuroscience field,  
450 we expect that the SI will be a powerful tool to shed light into a wide range of questions.  
451 Here we have provided several examples, from high-dimensional geometrical analysis to  
452 sound and image categorization, expanding the applicability of the tool across fields.

453

#### 454 **CRediT authorship contribution statement**

455 ERS: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation,  
456 Writing. JE: Conceptualization, Methodology, Software, Validation, Formal analysis,  
457 Investigation, Writing. LMP: Conceptualization, Writing, Supervision Resources, Project  
458 Administration, Funding acquisition.

459

#### 460 **Declaration of Competing Interest**

461 The authors declare that they have no competing interests.

462

#### 463 **Acknowledgments**

464 This work is supported by a grant from Fundación La Caixa (LCF/PR/HR21/52410030;  
465 DeepCode) to LMP. JE received the support of a PhD fellowship from "la Caixa"  
466 Foundation (ID 100010434; LCF/BQ/DR22/11950026). Access to supercomputer cluster  
467 Artemisa (NeuroDIM) is co-funded by the European Union through the 2014-2020 FEDER  
468 Operative Programme of Comunitat Valenciana, project IDIFEDER/2018/048.

469

#### 470 **Data availability**

471 All data used in this study is publicly available (see section 2.3).

472

#### 473 **References**

474 Baier G, Hermann T, Stephani U (2007) Event-based sonification of EEG rhythms in real  
475 time. *Clin Neurophysiol* 118:1377–1386.

476 Chaudhuri R, Gercek B, Pandey B, Peyrache A, Fiete I (2019) The intrinsic attractor  
477 manifold and population dynamics of a canonical cognitive circuit across waking and  
478 sleep. *Nat Neurosci* 22:1512–1520.

479 Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy  
480 K V (2012) Neural population dynamics during reaching. *Nature* 487:51–56.

481 Cunningham JP, Yu BM (2014) Dimensionality reduction for large-scale neural recordings.  
482 *Nat Neurosci* 17:1500–1509.

483 Douchamps V, di Volo M, Torcini A, Battaglia D, Goutagny R (2022) Hippocampal gamma  
484 oscillations form complex ensembles modulated by behavior and learning. *bioRxiv*  
485 doi: 10.11.

486 Durbin J, Koopman S (2001) Time series analysis by space state methods. Oxford  
487 University Press.

488 Engel J, Resnick C, Roberts A, Dieleman S, Eck D, Simonyan K, Norouzi M (2017) Neural  
489 Audio Synthesis of Musical Notes with WaveNet Autoencoders. *arXiv* 1704.01279.

490 Enns RH (2011) It's a nonlinear world. Springer.

491 Gallego JA, Perich MG, Miller LE, Solla SA 4 (2017) Neural Manifolds for the Control of  
492 Movement. *Neuron* 94:978–984.

493 Gardner RJ, Hermansen E, Pachitariu M, Burak Y, Baas NA, Dunn BA, Moser MB, Moser  
494 EI (2022) Toroidal topology of population activity in grid cells. *Nat* 2022 602:123–128.

495

496 Gervain J, Geffen MN (2019) Efficient Neural Coding in Auditory and Speech Perception.  
497 *Trends Neurosci* 42:56–65.

498 Gervasoni D, Lin S-C, Ribeiro S, Soares ES, Pantoja J, Nicolelis MAL (2004) Global  
499 Forebrain Dynamics Predict Rat Behavioral States and Their Transitions. *J Neurosci*  
500 24:11137–11147.

501 Lopes-Dos-Santos V, van de Ven GM, Morley A, Trouche S, Campo-Urriza N, Dupret D  
502 (2018) Parsing Hippocampal Theta Oscillations by Nested Spectral Components  
503 during Spatial Exploration and Memory-Guided Behavior. *Neuron* 100:940–952.e7.

504 McInnes L, Healy J, Melville J (2018) UMAP: Uniform Manifold Approximation and  
505 Projection for Dimension Reduction. Available at: <http://arxiv.org/abs/1802.03426>  
506 [Accessed January 16, 2020].

507 Navas-Olive A, Valero M, de Salas A, Jurado-Parras T, Averkin RG, Gambino G, Cid E, de  
508 la Prida LM (2020) Multimodal determinants of phase-locked dynamics across deep-  
509 superficial hippocampal sublayers during theta oscillations. *Nat Commun* 11:2217.

510 Nieh EH, Schottdorf M, Freeman NW, Low RJ, Lewallen S, Koay SA, Pinto L, Gauthier JL,  
511 Brody CD, Tank DW (2021) Geometry of abstract learned knowledge in the  
512 hippocampus. *Nature* 595:8084.

513 Peyrache A, Lacroix MM, Petersen PC, Buzsáki G (2015) Internally organized  
514 mechanisms of the head direction sense. *Nat Neurosci* 18:569–575.

515 Reichinnek S, Künsting T, Draguhn A, Both M (2010) Field potential signature of distinct  
516 multicellular activity patterns in the mouse hippocampus. *J Neurosci* 30:15441–15449.

517 Senzai Y, Scanziani M (2022) A cognitive process occurring during sleep is revealed by  
518 rapid eye movements. *Science* (80- ) 377:999–1004.

519 Tenenbaum J, de Silva V, Langford J (2000) A global geometric framework for nonlinear  
520 dimensionality reduction. *Science* (80- ) 290:2319–2323.

521 Ternes L, Dane M, Gross S, Labrie M, Mills G, Gray J, Heiser L, Chang YH (2022) A multi-  
522 encoder variational autoencoder controls multiple transformational features in single-  
523 cell image analysis. *Commun Biol* 5:255.

524 Valero M, Averkin RG, Fernandez-Lamo I, Aguilar J, Lopez-Pigozzi D, Brotons-Mas JR,  
525 Cid E, Tamas G, Menendez de la Prida L (2017) Mechanisms for Selective Single-Cell  
526 Reactivation during Offline Sharp-Wave Ripples and Their Distortion by Fast Ripples.  
527 *Neuron* 94.

528 Zeisel A, Moz-Manchado AB, Codeluppi S, Lönnerberg P, Manno G La, Juréus A, Marques  
529 S, Munguba H, He L, Betsholtz C, Rolny C, Castelo-Branco G, Hjerling-Leffler J,  
530 Linnarsson S (2015) Cell types in the mouse cortex and hippocampus revealed by  
531 single-cell RNA-seq. *Science* (80- ) 347:1138–1142.

532 Zeng J, Cheung G, Ng M, Pang J, Yang C (2019) 3D Point Cloud Denoising using Graph  
533 Laplacian Regularization of a Low Dimensional Manifold Model. *arXiv* 1803.07252.

534