



OPEN

Data segmentation based on the local intrinsic dimension

Michele Allegra^{1,2}, Elena Facco², Francesco Denti³, Alessandro Laio^{2,4}✉ & Antonietta Mira^{5,6}✉

One of the founding paradigms of machine learning is that a small number of variables is often sufficient to describe high-dimensional data. The minimum number of variables required is called the intrinsic dimension (ID) of the data. Contrary to common intuition, there are cases where the ID varies within the same data set. This fact has been highlighted in technical discussions, but seldom exploited to analyze large data sets and obtain insight into their structure. Here we develop a robust approach to discriminate regions with different local IDs and segment the points accordingly. Our approach is computationally efficient and can be proficiently used even on large data sets. We find that many real-world data sets contain regions with widely heterogeneous dimensions. These regions host points differing in core properties: folded versus unfolded configurations in a protein molecular dynamics trajectory, active versus non-active regions in brain imaging data, and firms with different financial risk in company balance sheets. A simple topological feature, the local ID, is thus sufficient to achieve an unsupervised segmentation of high-dimensional data, complementary to the one given by clustering algorithms.

From string theory to science fiction, the idea that we might be “glued” onto a low-dimensional surface embedded in a space of large dimensionality has tickled the speculations of scientists and writers alike. However, when it comes to multidimensional data such a situation is quite common rather than wild speculation: data often concentrate on hypersurfaces of low *intrinsic dimension* (ID). Estimating the ID of a dataset is a routine task in machine learning: it yields important information on the global structure of a dataset, and is a necessary preliminary step in several analysis pipelines.

In most approaches for dimensionality reduction and manifold learning, the ID is assumed to be constant in the dataset. This assumption is implicit in projection-based estimators, such as Principal Component Analysis (PCA) and its variants¹, Locally Linear Embedding², and Isomap³; and it also underlies geometric ID estimators^{4–6}, which infer the ID from the distribution of between-point distances. The hypothesis of a constant ID complies with simplicity and intuition but is not necessarily valid. In fact, many authors have considered the possibility of ID variations within a dataset^{7–22}, often proposing to classify the data according to this feature. However, the dominant opinion in the community is still that a variable ID is a peculiarity, or a technical detail, rather than a common feature to take into account before performing a data analysis. This perception is in part due to the lack of sufficiently general methods to track ID variations. Many of the methods developed in this field make assumptions that limit their applicability to specific classes of data. Refs.^{7,11–14} use local ID estimators that implicitly or explicitly assume a uniform density. Refs.^{8–10} jointly estimate the density and the ID from the scaling of neighbor distances, by approaches which work well only if the density varies slowly and is approximately constant in a large neighborhood of each point. Ref.¹⁵ requires a priori knowledge on the number of manifolds and their IDs. Refs.^{16–18,20} all require that the manifolds on which the data lay are hyperplanes, or topologically isomorphic to hyperplanes. These assumptions (locally constant density, and linearity in a suitable set of coordinates) are often violated in the case of real-world data. Moreover, many of the above approaches^{7,11–13,15–20} work explicitly with the coordinates of the data, while in many applications only distances between pairs of data points are available. To our knowledge, only Refs.^{21,22} do not make any assumption about the density, as they derive a parametric form of the distance distribution using extreme-value theory, which in principle is valid independently of the form of the underlying density. However, they assume that the low tail of the distance distribution is well approximated by its asymptotic form, an equally non-trivial assumption.

¹Institut de Neurosciences de la Timone UMR 7289, Aix Marseille Université, CNRS, 13385 Marseille, France. ²Scuola Internazionale Superiore di Studi Avanzati, Trieste, Italy. ³University of California, Irvine, CA, USA. ⁴International Centre for Theoretical Physics, Trieste, Italy. ⁵Università della Svizzera italiana, Lugano, Switzerland. ⁶Università dell’Insubria, Como, Italy. ✉email: laio@sissa.it; antonietta.mira@usi.ch

In this work, we propose a method to segment data based on the local ID which aims at overcoming the aforementioned limitations. Building on TWO-NN²³, a recently proposed ID estimator which is insensitive to density variations and uses only the distances between points, we develop a Bayesian framework that allows identifying, by Gibbs sampling, the regions in the data landscape in which the ID can be considered constant. Our approach works even if the data are embedded on highly curved and twisted manifolds, topologically complex and not isomorphic to hyperplanes, and if they are harvested from a non-uniform probability density. Moreover, it is specifically designed to use only the distance between data points, and not their coordinates. These features, as we will show, make our approach computationally efficient and more robust than other methods. Applying our approach to data of various origins, we show that ID variations between different regions are common. These variations often reveal fundamental properties of the data: for example, unfolded states in a molecular dynamics trajectory of a protein lie on a manifold of a lower dimension than the one hosting the folded states. Identifying regions of different dimensionality in a dataset can thus be a way to perform an unsupervised segmentation of the data. At odds with common approaches, we do not group the data according to their density but perform segmentation based on a geometric property defined on a local scale: the number of independent directions along which neighboring data are spread.

Methods

A Bayesian approach for discriminating manifolds with different IDs. The starting point of our approach is the recently proposed TWO-NN estimator²³, which infers the IDs from the statistics of the distances of the first two neighbors of each point. Let the data (x_1, x_2, \dots, x_N) , with N indicating the number of points, be i) independent and identically distributed samples from a probability density $\rho(x)$ with support on a manifold with unknown dimension d , such that ii) for each point x_i , $\rho(x_i)$ is approximately constant in the spherical neighborhood with center i and radius given by the distance between i and its second neighbor. Assumption i) is shared all methods that infer the ID from the distribution of distances (e.g., Refs.^{5,6,9-12,14}. Assumption ii) is the key assumption of TWO-NN: while of course it may not be perfectly satisfied, other methods used to infer the ID from distances require uniformity of the density on a larger scale including the first k neighbors, usually with $k \gg 2$. Let r_{i1} and r_{i2} be the distances of the first and second neighbor of x_i , then $\mu_i \doteq \frac{r_{i2}}{r_{i1}}$ follows the Pareto distribution $f(\mu_i|d) = d\mu_i^{-(d+1)}$. This readily allows the estimation of d from μ_i , $i = 1, 2, \dots, N$. We can write the global likelihood of $\boldsymbol{\mu} \equiv \{\mu_i\}$ as

$$P(\boldsymbol{\mu}|d) = d^N \prod_{i=1}^N \mu_i^{-(d+1)} = d^N e^{-(d+1)V}, \quad (1)$$

where $V \doteq \sum_{i=1}^N \log(\mu_i)$. From Eq. (1), and upon specifying a suitable prior on d , a Bayesian estimate of d is immediately obtained. This estimate is particularly robust to density variations: in²³ it was shown that restricting the uniformity requirement only to the region containing the first two neighbors allows properly estimating the ID also in cases where ρ is strongly varying.

TWO-NN can be extended to yield a heterogeneous-dimensional model with an arbitrarily high number of components. Let $\boldsymbol{x} = \{x_i\}$ be i.i.d samples from a density $\rho(x)$ with support on the union of K manifolds with varying dimensions. This multi-manifold framework is common with many previous works investigating heterogeneous dimension in a dataset^{8-11,14-16,18-20,24,25}. Formally, let $\rho(x) = \sum_{k=1}^K p_k \rho_k(x)$ where each $\rho_k(x)$ has support on a manifold of dimension d_k and $\boldsymbol{p} \doteq (p_1, p_2, \dots, p_K)$ are the a priori probabilities that a point belongs to the manifolds $1, \dots, K$. We shall first assume that K is known, and later show how it can be estimated from the data. The distribution of the μ_i is simply a mixture of Pareto distributions:

$$P(\mu_i|\boldsymbol{d}, \boldsymbol{p}) \doteq \sum_{k=1}^K p_k d_k \mu_i^{-(d_k+1)}. \quad (2)$$

Here, like in ref.²³, we assume that the variables μ_i are independent. This condition could be rigorously imposed by restricting the product in Eq. (1) to data points with no common first and second neighbors. In fact, one does not find significant differences between the empirical distribution of $\boldsymbol{\mu}$ obtained with all points and the one obtained by restricting the sample to the independent points (i.e., the empirical distribution of μ , obtained considering entire sample, is well fitted by the theoretical Pareto law). However, restricting the estimation to independent samples generally worsens the estimates, as one is left with a small fraction (typically, $\approx 15\%$) of the original points. Therefore, for the purpose of parametric inference, we decide to include all points in the estimation. In Supplementary Table S1 we show that restricting the estimation to independent points yields results that are qualitatively in good agreement with those obtained using all points, but with generally worse dimension estimates. Following the customary approach²⁶ we introduce latent variables $\boldsymbol{z} \doteq (z_1, z_2, \dots, z_K)$ where $z_i = k$ indicates that point i belongs to manifold k . We have $P(\mu_i|\boldsymbol{d}, \boldsymbol{p}, \boldsymbol{z}) = P(\mu_i|z_i, \boldsymbol{d})P_{pr}(z_i|\boldsymbol{p})$ with with

$$P(\mu_i|z_i, \boldsymbol{d}) = d_{z_i} \mu_i^{-(d_{z_i}+1)}, P_{pr}(z_i|\boldsymbol{p}) = p_{z_i}. \text{ This yields the posterior}$$

$$P_{post}(\boldsymbol{z}, \boldsymbol{d}, \boldsymbol{p}|\boldsymbol{\mu}) \propto P(\boldsymbol{\mu}|\boldsymbol{z}, \boldsymbol{d})P_{pr}(\boldsymbol{z}|\boldsymbol{p})P_{pr}(\boldsymbol{d})P_{pr}(\boldsymbol{p}). \quad (3)$$

We use independent Gamma priors on \boldsymbol{d} , $d_k \sim \text{Gamma}(a_k, b_k)$ and a joint Dirichlet prior on $\boldsymbol{p} \sim \text{Dir}(c_1, \dots, c_K)$. We fix $a_k, b_k, c_k = 1 \forall k$, corresponding to a maximally non-informative prior on the \boldsymbol{p} and an expectation of $\mathbb{E}[d_k] = 1$ for all k . If one has a different prior expectation on \boldsymbol{d} and \boldsymbol{p} , other choices of prior may be more convenient.

The posterior (3) does not have an analytically simple form, but it can be sampled by standard Gibbs sampling²⁷, allowing for the joint estimation of $\mathbf{d}, \mathbf{p}, \mathbf{z}$. However, model (3) has a serious limitation: Pareto distributions with (even largely) different values of d overlap to a great extent. Therefore, the method can not be expected to correctly estimate the z_i : a given value μ_i may be compatible with several manifold memberships. This issue can be addressed by correcting an unrealistic feature of model (3), namely, the independence of the z_i . We assume that the neighborhood of a point is more likely to contain points from the same manifold than from different manifolds. This requirement can be enforced with an additional term in the likelihood that penalizes local inhomogeneity (see also Ref.¹⁴). Consider the q -neighbor matrix $\mathcal{N}^{(q)}$ with nonzero entries $\mathcal{N}_{ij}^{(q)}$ only if $j \neq i$ is among the first q neighbors of i . Let ξ be the probability to sample the neighbor of a point from the same manifold, and $1 - \xi$ the probability to sample it from a different manifold, with $\xi > 0.5$. Define n_i^{in} as the number of neighbors of i with the same manifold membership ($n_i^{in} \equiv \sum_j \mathcal{N}_{ij}^{(q)} \delta_{z_i z_j}$). Then we introduce a probability distribution for $\mathcal{N}^{(q)}$ as:

$$P(\mathcal{N}^{(q)}|\mathbf{z}) = \prod_{i=1}^N \frac{\xi^{n_i^{in}} (1 - \xi)^{q - n_i^{in}}}{\mathcal{Z}(\xi, N_{z_i})}, \quad (4)$$

where \mathcal{Z}_q is a normalization factor that depends also on the sizes of the manifolds (see Supplementary Information for its explicit expression). This term favors homogeneity within a q -neighborhood of each point. With this addition, the model now reads:

$$P_{post}(\mathbf{z}, \mathbf{d}, \mathbf{p}|\boldsymbol{\mu}, \mathcal{N}^{(q)}) \propto P(\boldsymbol{\mu}|\mathbf{z}, \mathbf{d})P(\mathcal{N}^{(q)}|\mathbf{z})P_{pr}(\mathbf{z}|\mathbf{p})P_{pr}(\mathbf{d})P_{pr}(\mathbf{p}). \quad (5)$$

The posterior (5) is sampled with Gibbs sampling starting from a random configuration of the parameters. We repeat the sampling M times starting from different random configurations of the parameters, keeping the chain with the highest maximum log-posterior value. The parameters \mathbf{d} and \mathbf{p} can be estimated by their posterior averages. For all estimations, we only include the last 10% of the points. This ensures that the initial burn-in period is excluded²⁸. Moreover, visual inspection of the MCMC output is performed to ensure that we do not incur in any label switching issues²⁹. Indeed, we observed that once the chains converge to one mode of the posterior distribution they do not transition to different modes. As for the \mathbf{z} , we estimate the value of $\pi_{ik} \equiv P_{post}(z_i = k)$. Point i can be safely assigned to manifold k if $\pi_{ik} > 0.8$, otherwise we will consider its assignment to be uncertain. Note that our definition of “manifold” requires compactness: two spatially separated regions with approximately the same dimension are recognized as distinct manifolds by our approach.

We name our method Hidalgo (Heterogeneous Intrinsic Dimension Algorithm). Hidalgo has three free parameters: the number of manifolds, K ; the local homogeneity range, denoted by q ; the local homogeneity level, denoted by ξ . We fix q and ξ based on preliminary tests conducted on several artificial data sets, which show that the optimal “working point” of the method is at $q = 3, \xi = 0.8$ (Supplementary Fig. S1). As is common for mixture models, the value of K can be estimated by model selection. In particular, we can compare the models with K and $K + 1$ for increasing K , starting from $K = 1$ and stopping when there is no longer significant improvement, as measured by the average log-posterior.

The computational complexity of the algorithm depends on the total number of points N , the number of manifolds K , and the parameter q . As input to the Gibbs sampling, the method requires the neighbor matrix and the $\boldsymbol{\mu}$ vector. To compute the latter, one must identify the first q neighbors of each point, which, starting from the distance matrix, requires $O(qN^2)$ steps (possibly reducible to $O(qN \log N)$ with efficient nearest-neighbor-search algorithms³⁰). The computation of a distance matrix from a coordinate matrix takes as $O(DN^2)$ steps, where D is the total number of coordinates. In all practical examples with $N \lesssim 10^6$ points these preprocessing steps require much less computational time than the Gibbs sampling. The complexity of the Gibbs sampler is dominated by the sampling of the \mathbf{Z} , which scales linearly with N and K , and quadratically with q (in total, KNq^2). The number of iterations required for convergence cannot be simply estimated. In practice, we found this time to be roughly linear in N (Supplementary Fig. S2). Thus, in practice the total complexity of Hidalgo is roughly quadratic in N . The code implementing Hidalgo is freely available at <https://github.com/micheleallegra/Hidalgo>.

Results

Validation of the method on artificial data. We first test Hidalgo on artificial data for which the true manifold partition of the data is known. We start from the simple case of two manifolds with different IDs, d_1 and d_2 . We consider several examples, varying the higher dimension d_1 from 5 to 9 while fixing the lower dimension d_2 to 4. On both manifolds $N = 1000$ points are sampled from a multivariate Gaussian with a variance matrix given by $1/d_i, i = 1, 2$ multiplied by the identity matrix of proper dimension. The two manifolds are embedded in a space with dimension corresponding to the highest dimension d_2 , with their centers at a distance of 0.5 standard deviations, so they are partly overlapping. In Fig. 1a,b we illustrate the results obtained in the case of fixed $\xi = 0.5$, equivalent to the absence of any statistical constraint on neighborhood uniformity (note that for $\xi = 0.5$ the parameter q is irrelevant). The estimates of the two dimensions are shown together with the normalized mutual information (NMI) between the estimated assignment of points and the true one. The latter is defined as the MI of the assignment and the ground truth labeling divided by the entropy of the ground truth labeling. As expected, without a constraint on the assignment of neighbors, the method is not able to correctly separate the points and thus to estimate the dimensions of the two manifolds, even in the case of quite different IDs (NMI < 0.001). As soon as we consider $\xi > 0.5$, results improve. A detailed analysis of the influence of the hyperparameters q and ξ is reported in Supplementary Information. Based on such analysis, we identify the optimal parameter choice as $q = 3, \xi = 0.8$. In Fig. 1c,d we repeat the same tests as in Fig. 1a,b but with $q = 3$ and $\xi = 0.8$. Now the NMI between the estimated and ground truth assignment is almost 1 in all cases and, cor-

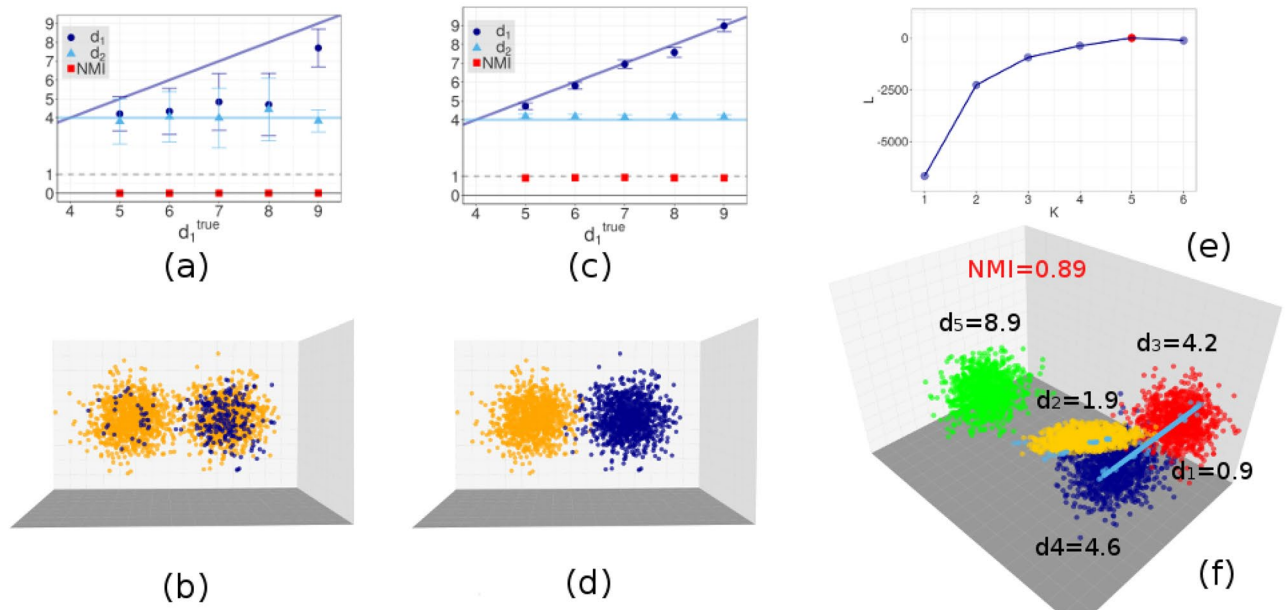


Figure 1. Results on simple artificial data sets. We consider sets of points drawn from mixtures of multivariate Gaussians in different dimensions. In all cases, we perform 10^5 iterations of the Gibbs sampling and repeat the sampling $M = 10$ times starting from different random configurations of the parameters. We then consider the sampling with the highest maximum average of the log-posterior value. Panels (a,b) Points are drawn from two Gaussians in different dimensions. The higher dimension varies from $d_1 = 5$ to $d_1 = 9$, the lower dimension is fixed at $d_2 = 4$. $N = 1000$ points are sampled from each manifold. We fix $K = 2$. We show results obtained with $\xi = 0.5$, namely, without enforcing neighborhood uniformity (here $q = 1$, but since $\xi = 0.5$ the value of q is irrelevant). In panel (a) we plot the estimated dimensions of the manifolds (dots: posterior means; error bars: posterior standard deviations) and the MI between our classification and the ground truth. In panel (b) we show the assignment of points to the low-dimensional (blue) and high-dimensional (orange) manifolds for the case $d_1 = 9$ (points are projected onto the first 3 coordinates). Similar figures are obtained for other values of d_1 . Panels (c,d) The same setting as in panel (a,b), but now we enforce neighborhood uniformity, using $\xi = 0.8$ and $q = 3$. Points are now correctly assigned to the manifolds whose ID is properly estimated. Panels (e,f) Points drawn from five Gaussians in dimensions $d_1 = 1, d_2 = 2, d_3 = 4, d_4 = 5, d_5 = 9$. $N = 1000$ points are sampled from each manifold. Some pairs of manifolds are intersecting, as their centers are one standard deviation apart. The analysis is performed assuming $\xi = 0.8, q = 3$, and with different values of K . In panel (e) we show the average log-posterior value \mathcal{L} as a function of K . The maximum \mathcal{L} corresponds to the ground truth value $K = 5$. In panel (f) we show the assignment of points to the five manifolds in different colors (points were projected onto the first 3 coordinates).

respondingly, the estimation of d_1 and d_2 is accurate. To verify whether our approach can discriminate between more than two manifolds ($K > 2$), we consider a more challenging scenario consisting of five Gaussians with unitary variance in dimensions 1, 2, 4, 5, 9 respectively. Some of the Gaussians have similar IDs, as in the case of dimensions 1 and 2, or 4 and 5: Moreover, they can be very close to each other: the centers of those in dimensions 4 and 5 are only half a variance far from each other, and they are crossed by the Gaussian in dimension 1. To analyze such dataset we again choose the hyperparameters $q = 3$ and $\xi = 0.8$. We do not fix the number of manifolds K to its ground truth value $K = 5$, but we try to let the method estimate K without relying on a priori information. We perform the analysis with different values of $K = 1, \dots, 6$ and compute an estimate of the average of the logarithm of the posterior value \mathcal{L} for each K . Results are shown in Fig. 1e. We see that \mathcal{L} increases up to $K = 5$, and then decreases, from which we infer that the optimal number of manifolds is $K = 5$. In Fig. 1f we illustrate the final assignment of points to the respective manifolds together with the estimated dimensions, upon setting the number of manifolds to $K = 5$. The separation of the manifolds is very good. Only a few points of the manifold with dimension 1 are incorrectly assigned to the one with dimension 2 and vice versa. The values of normalized mutual information between the ground truth and our classification is 0.89.

Finally, we show that our method can handle data embedded in curved manifolds, or with complex topologies. Such data should not pose additional challenges to the method: Hidalgo is sensitive only to the local structure of data. In terms of the distribution of the first and second nearest-neighbor distances, a 2-d torus looks exactly like a 2-d plane, and thus we should be able to correctly classify these objects, that are so topologically different, as 2-d objects regardless of their shape. We generated a dataset with 5 Gaussians in dimension 1, 2, 4, 5, 9 (Fig. 2a). Each Gaussian is then embedded on a curved nonlinear manifold: a 1-dimensional circle, a 2-dimensional torus, a 4-dimensional Swiss roll, a 5-dimensional sphere, and a 9-dimensional sphere. None of these manifolds is topologically isomorphic to a hyperplane, except for the Swiss roll. Moreover, the manifolds are intersecting.

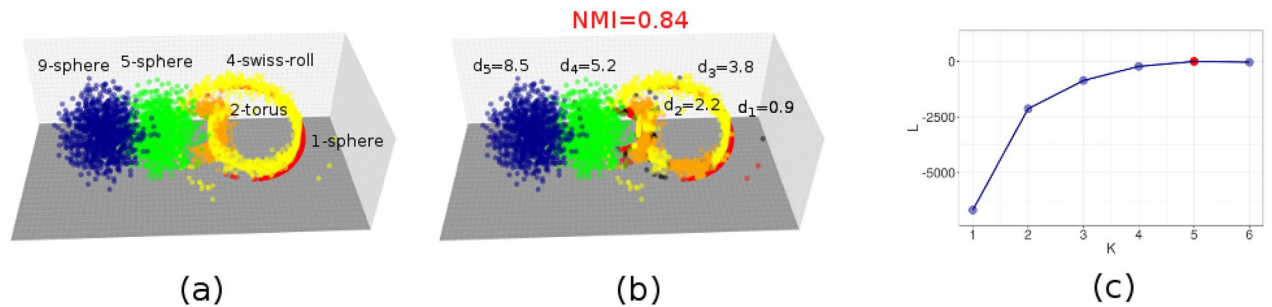


Figure 2. Results on data sets on curved manifolds. We consider sets of points drawn from mixtures of multivariate Gaussians embedded in curved nonlinear manifolds: a 1-dimensional sphere, a 2-dimensional torus, a 4-dimensional Swiss roll, a 5-dimensional sphere, and a 9-dimensional sphere. All the manifolds are embedded in a 9-dimensional space. $N = 1000$ points are sampled from each manifold. Some pairs of manifolds are intersecting. In all cases, we perform 10^5 iterations of the Gibbs sampling and repeat the sampling $M = 10$ times starting from different random configurations of the parameters. We consider the sampling with the highest maximum log-posterior value. The analysis is performed with $\xi = 0.8$, $q = 3$, and with different values of K . In panel (a) we show the ground-truth assignment of points to the five manifolds in different colors (points were projected onto the first 3 coordinates). In panel (b) we show the assignment of points as given by Hidalgo with $K = 5$. In panel (c) we show the average log-posterior value \mathcal{L} as a function of K . The maximum \mathcal{L} corresponds to the ground truth value $K = 5$.

As shown in Fig. 2b, our approach can distinguish the five manifolds. The dimensions are correctly retrieved (Fig. 2b), points are assigned to the right manifold with an accuracy corresponding to a value of the normalized mutual information (NMI) = 0.84. Only some points of the 3d-torus are misassigned to the 4d-swiss roll. As in the previous example, \mathcal{L} increases up to $K = 5$, and then decreases, from which we infer that the optimal number of manifolds is $K = 5$ (Fig. 2c).

ID variability in a protein folding trajectory. As a first real application of Hidalgo, we address ID estimation for a dynamical system. Asymptotically, dynamical systems are usually restrained to a low-dimensional manifold in phase space, called an attractor. Much effort has been devoted to characterizing the ID of such attractor⁴. However, in the presence of multiple metastable states an appropriate description of the visited phase space may require the use of multiple IDs. Here, we consider the dynamics of the villin headpiece (PDB entry: 2F4K). Due to its small size and fast folding kinetics, this small protein is a prototypical system for molecular dynamics simulations. Our analysis is based on the longest available simulated trajectory of the system from Ref.³³. During the simulated $125 \mu\text{s}$, the protein performs approximately 10 transitions between the folded and the unfolded state. We expect to find different dimensions in the folded and unfolded state, since these two states are metastable, and they would be considered as different attractors in the language of dynamical systems. Moreover, they are characterized by different chemical and physical features: the folded state is compact and dry in its core, while the unfolded state is swollen, with most of the residues interacting with a large number of water molecules. We extract the value of the 32 key observables (the backbone dihedral angles) for all the $N = 31,000$ states in the trajectory and apply Hidalgo to this data of extrinsic dimension $D = 32$. We obtain a vector of estimated intrinsic dimensions \mathbf{d} and an assignment of each point i to one of the K manifolds. We find four manifolds, three low-dimensional ones ($d_1 = 11.8$, $d_2 = 12.9$, $d_3 = 13.2$) and a high-dimensional one ($d_4 = 22.7$). We recall that two spatially separated regions with approximately the same dimension (in this case, d_2 and d_3) are recognized as distinct manifolds by our approach. To test whether this partition into manifolds is related to the separation between the folded and the unfolded state we relate the partition to the fraction of native contacts Q , which can be straightforwardly estimated on each configuration of the system. Q is close to one only if the configuration is folded, while it approaches zero when the protein is unfolded. In Fig. 3 we plot the probability distribution of Q restricted to the four manifolds. We find that the vast majority of the folded configurations ($Q > 0.8$) are assigned to the high-dimensional manifold. Conversely, the unfolded configurations ($Q \leq 0.8$) are most of the times assigned to one of the low-dimensional manifolds. This implies that a configuration belonging to the low dimensional manifolds is almost surely unfolded. The normalized mutual information between the manifold assignment and the folded/unfolded classification is 0.60. Thus, we can essentially identify the folded state using the intrinsic dimension, a purely topological observable unaware of any chemical detail.

ID variability in time-series from brain imaging. In the next example, we analyze a set of time-series from functional resonance imaging (fMRI) of the brain, representing the BOLD (blood oxygen-level dependent) signal of each voxel, which captures the activity a small part of the brain³⁴. fMRI time-series are often projected on a lower dimension through linear projection techniques like PCA³⁵, a step that assumes a uniform ID. However, the gross features of the signal (e.g., power spectrum and entropy) are often highly variable in different parts of the brain, and also non-uniformities in the ID may well be present. Here, we consider a single-subject fMRI recording containing $D = 202$ images collected while a subject was performing a visuomotor task^{31,36}. From the images we extract the $N = 29851$ time series corresponding to the BOLD signals of each voxel. Apply-

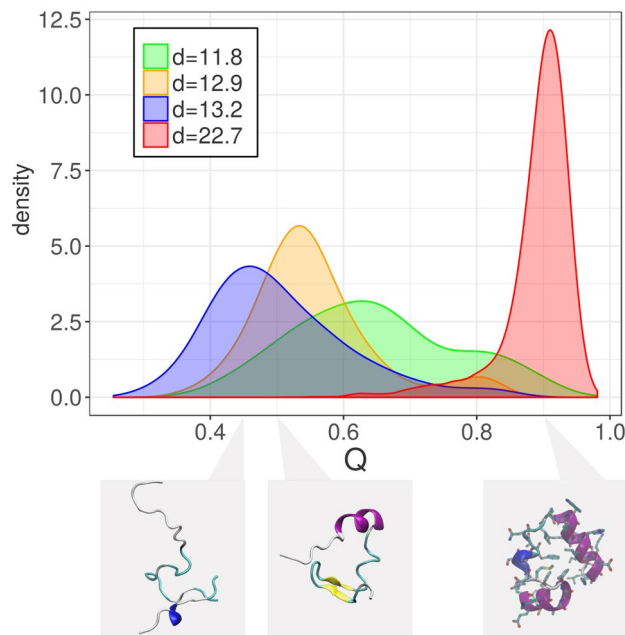


Figure 3. Protein folding trajectory. We consider $N \sim 31,000$ configurations of a protein undergoing successive folding/unfolding cycles. For each configuration, we extract the value of the $D = 32$ backbone dihedral angles. Applying Hidalgo to these data, we detect four manifolds, of intrinsic dimensions 11.8, 12.9, 13.2, and 22.7. For each configuration, we also compute the fraction of native contacts, Q , which measures to which degree the configuration is folded. The figure shows the probability distribution of Q in each manifold. Most of the folded configurations belong to the high-dimensional manifold: the analysis essentially identifies the folded configurations as a region of high intrinsic dimension. Results are obtained with $q = 3$ and $\xi = 0.8$. The distance between each pair of configurations was computed by the Euclidean metric with periodic boundary conditions on the vectors of the dihedral angles.

ing Hidalgo, we find two manifolds with very different dimensions $d_1 = 16.2$, $d_2 = 31.9$. Again, we relate the identified manifolds to a completely independent quantity, the clustering frequency Φ introduced in^{31,32}, which measures the temporal coherence of the signal of a voxel with the signals of other voxels in the brain. Voxels with non-negligible clustering frequency ($\Phi > 0.2$) are likely to belong to brain areas involved in the cognitive task at hand. In Fig. 4 we show the probability distribution of Φ restricted to the two manifolds. We find that the “task-related” voxels ($\Phi > 0.2$) almost invariably belong to the manifold with high dimensionality. These voxels appear concentrated in the occipital, parietal, and temporal cortex (Fig. 4 bottom), and belong to a task-relevant network of coherent activity³¹. The normalized mutual information between the manifold assignment and the task-relevant/task-irrelevant classification is 0.24; note that this value is lower than for the protein folding case, because the low-dimensional manifold is “contaminated” by many points with low Φ . This result finds an interesting interpretation: the subset of “relevant” voxels gives rise to patterns that are not only coherent but also characterized by a larger ID than the remainder of the voxels. On the contrary, the incoherent voxels exhibit a lower ID, hence a reduced variability, which is consistent with the fact that the corresponding time series are dominated by low-dimensional noise. Again, this feature emerges from the global topology of the data, revealed by our ID analysis, without exploiting any knowledge of the task that the subject is performing.

ID variability in financial data. Our final example is in the realm of economics. We consider firms in the well-known Compustat database ($N = 8309$). For each firm, we consider $D = 31$ balance sheet variables from the fiscal year 2016 (for details, see Supplementary Table S2). Applying Hidalgo we find four manifolds of dimensions $d_1 = 5.4$, $d_2 = 6.3$, $d_3 = 7.0$ and $d_4 = 9.1$. To understand this result, we try to relate our classification with common indexes showing the type and financial stability of a firm. We start by relating our classification to the Fama-French classification³⁷, which assigns each firm to one of twelve categories depending on the firm’s trade. In Fig. 5(top) we separately consider firms belonging to the different Fama-French classes, and compute the fraction of firms assigned to each of the four manifolds identified by Hidalgo. The two classifications are not independent, since the fractions for different Fama-French classes are highly non-uniform. More precisely, the normalized mutual information between the two classifications is 0.19, rejecting hypothesis of statistical independence (p -value $< 10^{-5}$). In particular, firms in the utilities and energy sector show a preference for low dimensions (d_1 and d_2), while firms purchasing products (nondurables, durables, manufacturing chemicals, equipment, wholesale) are concentrated in the manifold with the highest dimension d_4 . The manifold with intrinsic dimension d_3 mostly includes firms in the financial and health care sectors. Different dimensions are not only related to the classification of the firm, but also their financial robustness. We consider the S&P quality ratings for the firms assigned to each manifold (also from Compustat; ratings are available for 2894 firms). In Fig. 5(bottom)

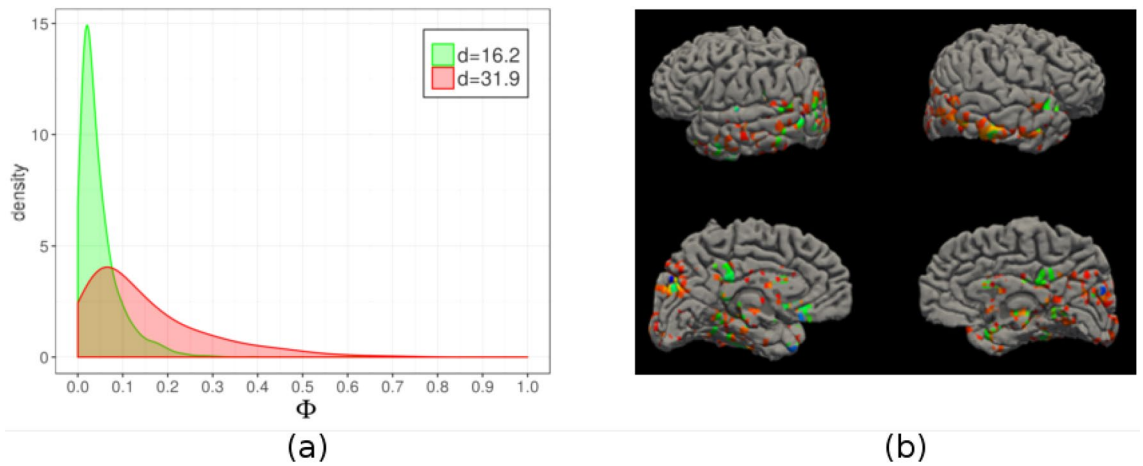


Figure 4. Neuroimaging data. We consider the BOLD time series of $N \sim 30,000$ voxels in an fMRI experiment with $D = 202$ scans. Hidalgo detects a low-dimensional manifold ($d = 16.2$) and a high-dimensional one ($d = 31.9$). We compute the clustering frequency Φ , which measures the participation of each voxel to coherent activation patterns and is a proxy for voxel involvement in the task³¹. Panel (a) shows the probability distribution of Φ in the two manifolds. Strongly activated voxels ($\Phi > 0.2$) are consistently assigned to the high-dimensional manifold. Panel (b) shows the rendering of the cortical surface (left: left hemisphere; right: right hemisphere). Blue voxels have high clustering frequency ($\Phi > 0.2$), red voxels are those assigned to the high-dimensional manifold, and green voxels satisfying both criteria. Almost all voxels with high clustering frequency are assigned to the high-dimensional manifold, and are concentrated in the occipital, temporal, and parietal cortex. Results are obtained with $q = 3$ and $\xi = 0.8$. The distance between two time series is computed by a Euclidean metric, after standard pre-processing steps³².

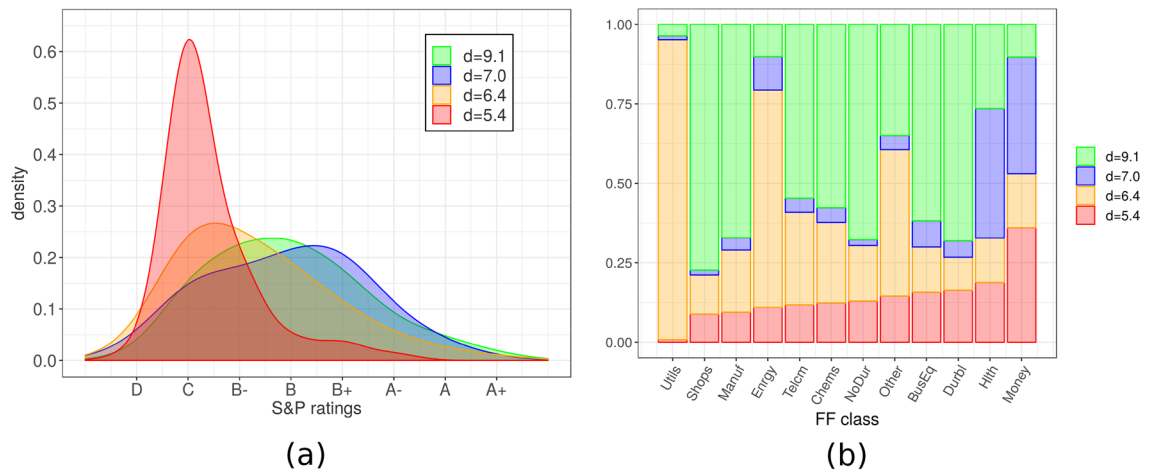


Figure 5. Financial data. For $N \sim 8,000$ firms selected from the COMPUSTAT database, we compute $D = 31$ variables from their yearly balance sheets. Hidalgo finds four manifolds of intrinsic dimensions 5.4, 6.4, 7.0, and 9.1. Panel (a) shows the fractions of firms assigned to the four manifolds for each type of firm, according to the Fama–French classification. The four manifolds contain unequal proportions of firms belonging to different classes, implying that some classes of firms are preferentially assigned to manifolds of high versus low dimension. Panel (b) shows the probability distribution of the S&P ratings of the firms assigned to each manifold. Firms with low ratings preferentially belong to low-dimensional manifolds. Results are obtained with $q = 3$ and $\xi = 0.8$. To correct for firm size, we normalize the variable vector of each firm by its norm, and then applied standard Euclidean metric.

we show the distribution of ratings for the different manifolds. These distributions appear to be different. We cannot predict the rating based on the manifold assignment (the normalized mutual information between the manifold assignment and the rating is 0.04), but companies with worse ratings show a preference to belong to low-dimensional manifolds: by converting ratings into numerical values scores from 1 to 8, we found a Pearson correlation of 0.22 between the local ID and the rating ($p < 10^{-10}$). We suggest a possible interpretation for this phenomenon: a low ID may imply a more rigid balance structure, which may entail a higher sensitivity to market shocks which, in turn, may trigger domino effects in contagion processes. This result shows that information

about the S&P rating can be derived using only the topological properties of the data landscape, without any in-depth financial analysis. For example, no information on the commercial relationship between the firms or the nature of their business is used.

Comparison with other methods. We have compared our method with two state-of-the-art methods in the literature, considering the synthetic datasets as well as the protein folding dataset (for which the folded/unfolded classification yields an approximate “ground truth”). The first method we consider is the “sparse manifold clustering and embedding” (SMCE, Ref.¹⁹). The method creates a graph by connecting neighboring points, supposedly lying on the same manifold, and then uses spectral clustering on this graph to retrieve the manifolds. As an output, it returns a manifold assignment for each point, together with an ID estimate. We resort to the implementation provided by Ehsan Elhamifar (<http://khoury.neu.edu/home/eelhami/codes.htm>). For SMCE, we can compute the NMI between the manifold assignment and the ground-truth manifold label. Furthermore, we can report the estimated local ID for points with a given ground-truth assignment, identified with the ID of the manifold to which they are assigned. SMCE requires to specify the number of manifolds K as input and depends on a free parameter λ . To present a fair comparison, we fixed K at its ground truth value and explored different values in the range $0.01 \leq \lambda \leq 20$, and we report the results corresponding to the highest NMI with the ground truth. Furthermore, we repeated the estimation $M = 10$ times and kept the results with the highest NMI with the ground truth. The second method is the local ID (LID) estimation by Carter et al.¹⁰, which combines ID estimation restricted to a (large) neighborhood and local smoothing to produce an integer ID estimate for each point. We used the implementation provided by Kerstin Johnsson³⁸ (<https://github.com/kjohnsson/intrinsicDimension>) using a neighborhood size $k = 50$ which is a standard value. Note that the LID estimation is deterministic, so we do not need to repeat the estimation several times. For LID, we can compute the NMI between the (integer) label given by the local ID estimate and the ground truth; moreover, we can compute the average ID, $\langle d \rangle$ for the points with a given ground truth assignment.

The results of the comparison are presented in Table 1. For the examples with two manifolds, SMCE¹⁹ is able to correctly retrieve the manifolds with high accuracy ($\text{NMI} \geq 0.94$). The estimate of the lower dimension $d_2 = 4$ is always correct, while the highest dimension is underestimated for $d_1 \geq 7$. Overall, SMCE performs very well (even better than Hidalgo) on the datasets with two manifolds, as the two manifolds are well separated, matching well the assumptions of the model. On the examples with five manifolds, the performance of SMCE is worse, as the method cannot correctly retrieve intersecting manifolds. In the example with five linear manifolds, SMCE cannot retrieve the manifold of dimension $d_1 = 1$, which is intersecting the manifolds of dimension $d_3 = 4$ and $d_4 = 5$. Moreover, the ID estimates are quite poor. In the example with five curved manifolds, SMCE merges the manifolds with $d_4 = 5$ and $d_5 = 9$, and cannot retrieve the manifolds with $d_1 = 1$ and $d_3 = 4$, which get split. Again, the dimension estimates are poor. For the protein folding data, where regions with different IDs are probably not well separated, SMCE finds a single manifold with $d = 11$, and thus it is not able to detect any ID variation.

For the examples with two manifolds, the LID method¹⁰ is generally able to correctly retrieve regions with different IDs with high accuracy ($\text{NMI} \geq 0.92$), except for the case with $d_1 = 5$, $d_2 = 4$ where results are inaccurate. The estimate of the lower dimension $d_2 = 4$ is always precise, while the higher dimension is always underestimated. The degree of underestimation increases with increasing d_1 . This is expected, since the ID estimates are produced by assuming a uniform density in a large neighborhood of each point, an assumption that fails in these data, especially for points on the manifold with higher dimension. Overall, LID performs quite well (comparably with Hidalgo) in separating regions with different IDs in the datasets with two manifolds, but gives worse ID estimates than Hidalgo. In the examples with five manifolds, the performance of LID is slightly worse. In the example with five linear manifolds LID merges the manifolds of dimension $d_3 = 4$ and $d_4 = 5$, and misassigns some of their points to the manifold of dimension 1. This is because relatively large neighborhoods are used for the estimation of the local ID: this leads to a difficulty in discriminating close regions with similar ID ($d_3 = 4$ and $d_4 = 5$) and to a “contamination” of results by points of the manifold with $d_1 = 1$ intersecting the two. In the example with five curved manifolds, LID can correctly identify the manifolds with $d_2 = 2$, $d_4 = 5$, and $d_5 = 9$, even if it underestimates the ID in the latter two cases (high dimension, large density variations). Many points in the manifold of dimension $d_3 = 4$ are mis-assigned to the manifold of dimension $d_1 = 1$, due to the fact that the two manifolds are highly intersecting. Finally, for the protein folding data, where the density is highly variable, LID shows a tendency for unfolded configurations to have a higher local ID than unfolded ones. However, it cannot discriminate between folded and unfolded configurations as Hidalgo. The ID of the folded configurations is highly underestimated compared to the Hidalgo estimate, probably because of density variations.

In Fig. 6 we compare the results of the three methods (Hidalgo, SMCE, and LID) for the dataset with five curved manifolds.

Discussion

The increasing availability of a large amount of data has considerably expanded the opportunities and challenges for unsupervised data analysis. Often data come in the form of a completely uncharted “point cloud” for which no model is at hand. The primary goal of the analyst is to uncover some structure within the data. For this purpose, a typical approach is dimensionality reduction, whereby the data are simplified by projecting them onto a low-dimensional space.

The appropriate *intrinsic dimension (ID)* of the space onto which one should project the data is not constant everywhere. In this work, we developed an algorithm (Hidalgo) to find manifolds of different IDs in the data. Applying Hidalgo, we observed large variations of the ID in datasets of diverse origin (a molecular dynamics simulation, a set of time series from brain imaging, a dataset of firm balance sheets). This finding suggests that a

Dataset	Hidalgo d	NMI	Ref. ¹⁹ d	NMI	Ref. ¹⁰ d	NMI
2 Gaussians	$d_1 = 4.7$ $d_2 = 4.2$	0.91	$d_1 = 5$ $d_2 = 4$	0.96	$\langle d \rangle = 4.0$ $\langle d \rangle = 4.7$	0.44
2 Gaussians	$d_1 = 5.8$ $d_2 = 4.2$	0.93	$d_1 = 6$ $d_2 = 4$	0.94	$\langle d \rangle = 4.0$ $\langle d \rangle = 5.1$	0.95
2 Gaussians	$d_1 = 6.9$ $d_2 = 4.1$	0.94	$d_1 = 6$ $d_2 = 4$	0.97	$\langle d \rangle = 4.0$ $\langle d \rangle = 6.0$	0.96
2 Gaussians	$d_1 = 7.6$ $d_2 = 4.2$	0.92	$d_1 = 7$ $d_2 = 4$	0.94	$\langle d \rangle = 4.0$ $\langle d \rangle = 6.7$	0.93
2 Gaussians	$d_1 = 9.0$ $d_2 = 4.1$	0.92	$d_1 = 7$ $d_2 = 4$	0.95	$\langle d \rangle = 4.0$ $\langle d \rangle = 7.2$	0.92
5 Gaussians (linear)	$d_1 = 0.9$ $d_2 = 1.9$ $d_3 = 4.2$ $d_4 = 4.6$ $d_5 = 8.9$	0.89	$d_1 = 7, d_4 = 9, d_5 = 10$ $d_2 = 10$ $d_3 = 9$ $d_4 = 9$ $d_5 = 10$	0.73	$\langle d \rangle = 1.0$ $\langle d \rangle = 2.0$ $\langle d \rangle = 3.5$ $\langle d \rangle = 4.3$ $\langle d \rangle = 7.3$	0.81
5 Gaussians (curved)	$d_1 = 0.9$ $d_2 = 2.2$ $d_3 = 2.2$ $d_4 = 3.8$ $d_5 = 5.2$ $d_5 = 8.5$	0.84	$d_1 = 6, d_3 = 7, d_4 = 8, d_5 = 10$ $d_2 = 10$ $d_3 = 10$ $d_4 = 10$ $d_5 = 10$	0.61	$\langle d \rangle = 1.0$ $\langle d \rangle = 2.0$ $\langle d \rangle = 2.0$ $\langle d \rangle = 4.6$ $\langle d \rangle = 7.2$	0.77
Protein folding	$d_1 = 11.8$ $d_2 = 12.9$ $d_3 = 13.2$ $d_4 = 22.3$	0.60	$d_1 = 11$	0	$\langle d \rangle = 11.4$ $\langle d \rangle = 13.7$	0.22

Table 1. Performances of Hidalgo, restricting analysis to points with non-overlapping first and second neighbor. For each data set we show the number of points in the original data set (N_0), the number of points upon restriction N , the estimated value of the d_k , and NMI between the assignment and the ground truth. QLL results were obtained with $\xi = 0.8, q = 3$, repeating the sampling $M = 10$ times and considering the sampling with the highest average log-posterior.

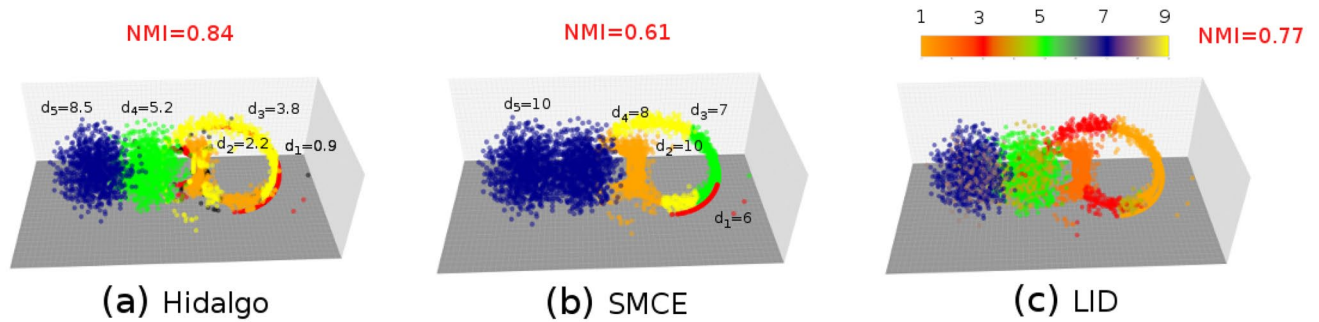


Figure 6. Results of other variable-ID analysis methods on topologically complex artificial data. We consider sets of points drawn from mixtures of multivariate Gaussians embedded in curved nonlinear manifolds, as detailed in Fig. 2, and compared the results of Hidalgo (a), SMCE (b), LID (c). In panel (a) we plot again, for ease of comparison, the results of Fig. 2b. In panel (b), we show the local ID estimates given by LID¹⁰, considering $k = 50$ nearest neighbors for local ID estimation and smoothing. The method assigns an integer dimension to each point: in the plot, the dimension of each point is represented by its color, according to the color bar shown. The method correctly retrieves the manifolds of dimension 5 and 9, even though the points in the manifold with dimension 5 are assigned a local ID that oscillates between 4 and 5, and points in the manifold with dimension 9 are assigned a local ID that oscillates between 7 and 8. Also, the manifold of dimension 2 is well identified. The manifolds of dimension 1 and 4, however, cannot be correctly discriminated. Points of the manifold of dimension 4 are assigned ID estimates of 1 or 3, without a clear separation of the manifolds. The NMI between the ground truth and the integer local ID is 0.77. In panel (c) we show the assignment of points as given by SMCE¹⁹. The manifolds of dimension 5 and 9 are merged. The manifold of dimension 2 is identified, but it is also contaminated by points from the manifold of dimension 4. SMCE cannot correctly retrieve the manifolds of dimension 1 and 4. The NMI between the ground truth and the assignment is 0.61. ID estimates obtained according to the prescription given in¹⁹ are largely incorrect.

highly non-uniform ID is not an oddity, but a rather common feature. Generally speaking, ID variations can be expected whenever the system under study can be in various “macrostates” characterized by a different number of degrees of freedom due, for instance, to differences in the constraints. As an example, the folded state the protein is able to locally explore the phase space in many independent directions, while in the unfolded state it only performs wider movements in fewer directions. In the case of companies, a financially stable company may have more degrees of freedom to adjust its balance sheet.

In the cases we have analyzed, regions characterized by different dimensions were found to host data points differing in important properties. Thus, variations of the ID within a dataset can be used to classify the data into different categories. It is remarkable that such classification is achieved by looking at a simple topological property, the local ID. Let us stress that ID-based segmentation should not be considered an alternative to clustering methods. In fact, in most cases (e.g., most classical benchmark sets for clustering) different clusters do not correspond to significantly different intrinsic dimensions - rather, they correspond to regions of high density of points within a single manifold of well-defined ID. In such cases, clusters cannot be identified as different manifolds by Hidalgo. Conversely, when the ID is variable, regions of different IDs do not necessarily correspond to clusters in the standard meaning of the word: they may or may not correspond to regions of high density of points. A typical example could be that of protein folding: while the folded configurations are quite similar to one another, and hence constitute a cluster in the traditional sense, the unfolded configurations may be very heterogeneous, hence quite far in data space and then not a “cluster” in the standard meaning of the word.

The idea that ID may vary in the same data is not new. In fact, many works have discussed the possibility of a variable ID and developed methods to estimate multiple IDs^{7–18,20–22}. Our method builds on these previous contributions but is designed with the specific goal of overcoming technical limitations of other available approaches and make ID-based segmentation a general-purpose tool. Our scheme uses only the distances between the data points, and not their coordinates, which significantly enlarges its scope of applicability. Moreover, the scheme uses only the distances between a point and its q nearest neighbors, with $q \leq 5$. We thus circumvent the notoriously difficult problem of defining a globally meaningful metric³, only needing a consistent metric on a small local scale. Finally, Hidalgo is not hindered by density variations or curvature. For these reasons, Hidalgo is competitive with other manifold learning and variable ID estimation methods and, in particular, can yield better ID estimates and manifold reconstruction. We have compared our method with two state-of-the-art methods in the literature, the local ID method (LID¹⁰), and the sparse manifold clustering and embedding (SMCE¹⁹) method. Both methods show issues in the cases of intersecting manifolds and variable densities and yield worse ID estimates than Hidalgo, especially for large IDs.

Hidalgo is computationally efficient and therefore suitable for the analysis of large data sets. For example, it takes ≈ 30 mins to perform the analysis of the neuroimaging data, which includes 30000 data points, on a standard computer using a single core. Implementing the algorithm with parallel Gibbs sampling³⁹ may considerably reduce computing time and then yield a fully scalable method. Obviously, Hidalgo has some limitations. Some are intrinsic to the way the data are modeled. Hidalgo is not suitable to cover cases in which the ID is a continuously varying parameter²¹, or in which sparsity is so strong that points cannot be assumed to be sampled from a

continuous distribution. In particular, Hidalgo is not suitable for discrete data for which the basic assumptions of the method are violated. For example, ties in the sample could lead to null distances between neighbors, jeopardizing the computation of μ . Others are technical issues related with the estimation procedure, and, at least in principle, susceptible to improvement in refined versions of the algorithm: for instance, one may improve the estimation with suitable enhanced sampling techniques. Finally, let us point out some further implications of our work. Our findings suggest a caveat with respect to common practices of dimensionality reduction, which assume a uniform ID. In the case of significant variations, a global dimensionality reduction scheme may become inaccurate. In principle, the partition in manifolds obtained with Hidalgo may be the starting point for using standard dimensionality reduction schemes. For example, one can imagine to apply PCA¹ or Isomap³, or sketchmap⁴⁰ separately to each manifold. However, we point out that a feasible scheme to achieve this goal does not come as an immediate byproduct of our method. Once a manifold with given ID is identified, it is highly nontrivial to provide a suitable parameterization thereof, especially because the manifolds may be highly nonlinear, and even topologically non-trivial. How to suitably integrate our approach with a dimensionality reduction scheme remains a topic for further research. Another implication is that a simple topological invariant, the ID, can be very a powerful tool for unsupervised data analysis, lending support to current efforts at characterizing topological properties of the data^{41,42}.

Data availability

Figures 1, 2, and 6 have raw simulated data. Figures 3, 4 and 5 use real data. All data used are publicly available in the following github repository: <https://github.com/micheleallegra/Hidalgo>

Code availability

Code is publicly available in the following github repository: <https://github.com/micheleallegra/Hidalgo>.

Received: 14 February 2020; Accepted: 13 July 2020

Published online: 05 October 2020

References

- Jolliffe, I. T. (ed) Principal component analysis and factor analysis. In *Principal Component Analysis*, 115–128 (Springer, Berlin, 1986).
- Roweis, S. T. & Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**, 2323–2326 (2000).
- Tenenbaum, J. B., De Silva, V. & Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000).
- Grassberger, P. & Procaccia, I. Measuring the strangeness of strange attractors. In *The Theory of Chaotic Attractors* (eds Hunt, B. R., Li, T.-Y., Kennedy, J. A., & Nusse, H. E.), 170–189 (Springer, Berlin, 2004).
- Levina, E. & Bickel, P. J. Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems 17* (eds Saul L. K., Weiss, Y., & Bottou, L.) (MIT Press, 2005).
- Rozza, A., Lombardi, G., Ceruti, C., Casiraghi, E. & Campadelli, P. Novel high intrinsic dimensionality estimators. *Mach. Learn.* **89**, 37–65 (2012).
- Barbará, D. & Chen, P. Using the fractal dimension to cluster datasets. In *Proceedings of the sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 260–264 (ACM, London, 2000).
- Gionis, A., Hinneburg, A., Papadimitriou, S. & Tsaparas, P. Dimension induced clustering. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 51–60 (ACM, London, 2005).
- Costa, J. A., Girotra, A. & Hero, A. Estimating local intrinsic dimension with k-nearest neighbor graphs. In *2005 IEEE/SP 13th Workshop on Statistical Signal Processing*, 417–422 (IEEE, 2005).
- Carter, K. M., Raich, R. & Hero, A. O. III. On local intrinsic dimension estimation and its applications. *IEEE Trans. Signal Process.* **58**, 650–663 (2010).
- Campadelli, P., Casiraghi, E., Ceruti, C., Lombardi, G. & Rozza, A. Local intrinsic dimensionality based features for clustering. In *International Conference on Image Analysis and Processing*, 41–50 (Springer, Berlin, 2013).
- Johnsson, K., Sonesson, C. & Fontes, M. Low bias local intrinsic dimension estimation from expected simplex skewness. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(1), 196–202 (2015).
- Mordohai, P. & Medioni, G. G. Unsupervised dimensionality estimation and manifold learning in high-dimensional spaces by tensor voting. In *IJCAI*, 798–803 (2005).
- Haro, G., Randall, G. & Sapiro, G. Translated poisson mixture model for stratification learning. *Int. J. Comput. Vis.* **80**, 358–374 (2008).
- Souvenir, R. & Pless, R. Manifold clustering. In *Tenth IEEE International Conference on Computer Vision, 2005. ICCV 2005*, vol. 1, 648–653 (IEEE, 2005).
- Wang, Y., Jiang, Y., Wu, Y. & Zhou, Z.-H. Multi-manifold clustering. In *Pacific Rim International Conference on Artificial Intelligence*, 280–291 (Springer, 2010).
- Goh, A. & Vidal, R. Segmenting motions of different types by unsupervised manifold clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR'07*, 1–6 (IEEE, 2007).
- Vidal, R. Subspace clustering. *IEEE Signal Process. Mag.* **28**, 52–68 (2011).
- Elhamifar, E. & Vidal, R. Sparse manifold clustering and embedding. *Advances in Neural Information Processing Systems 24*, (eds Shawe-Taylor, J., Zemel, R. S., Bartlett, P. L., Pereira, F., & Weinberger, K. Q.) 55–63 (NIPS, 2011).
- Elhamifar, E. & Vidal, R. Sparse subspace clustering: algorithm, theory, and applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 2765–2781 (2013).
- Amsaleg, L. *et al.* Estimating local intrinsic dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 29–38 (ACM, 2015).
- Faranda, D., Messori, G. & Yiou, P. Dynamical proxies of north atlantic predictability and extremes. *Sci. Rep.* **7**, 41278 (2017).
- Facco, E., d’Errico, M., Rodriguez, A. & Laio, A. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Sci. Rep.* **7**, 12140 (2017).
- Xiao, R., Zhao, Q., Zhang, D. & Shi, P. Data classification on multiple manifolds. In *2010 20th International Conference on Pattern Recognition (ICPR)*, 3898–3901 (IEEE, 2010).
- Goldberg, A., Zhu, X., Singh, A., Xu, Z. & Nowak, R. Multi-manifold semi-supervised learning. In *Artificial Intelligence and Statistics* 169–176 (2009).

26. Richardson, S. & Green, P. J. On bayesian analysis of mixtures with an unknown number of components (with discussion). *J. R. Stat. Soc. Ser. B Stat. Methodol.* **59**, 731–792 (1997).
27. Casella, G. & George, E. I. Explaining the gibbs sampler. *Am. Stat.* **46**, 167–174 (1992).
28. Diebolt, J. & Robert, C. P. Estimation of finite mixture distributions through bayesian sampling. *J. R. Stat. Soc. Ser. B (Methodol.)* **56**, 363–375 (1994).
29. Celeux, G. Bayesian Inference for Mixture: The Label Switching Problem. In *Compstat* 227–232. https://doi.org/10.1007/978-3-662-01131-7_26 (1998).
30. Preparata, F. P. & Shamos, M. I. *Computational Geometry: An Introduction* (Springer, Berlin, 2012).
31. Allegra, M. *et al.* Brain network dynamics during spontaneous strategy shifts and incremental task optimization. *NeuroImage* **116**854 (2020).
32. Allegra, M. *et al.* fmri single trial discovery of spatio-temporal brain activity patterns. *Hum. Brain Map.* **38**, 1421–1437 (2017).
33. Lindorff-Larsen, K., Piana, S., Dror, R. O. & Shaw, D. E. How fast-folding proteins fold. *Science* **334**, 517–520 (2011).
34. Huettel, S. A. *et al.* *Functional Magnetic Resonance Imaging* Vol. 1 (Sinauer Associates, Sunderland, 2004).
35. Poldrack, R. A., Mumford, J. A. & Nichols, T. E. *Handbook of Functional MRI Data Analysis* (Cambridge University Press, Cambridge, 2011).
36. Schuck, N. W. *et al.* Medial prefrontal cortex predicts internally driven strategy shifts. *Neuron* **86**, 331–340 (2015).
37. Fama, E. F. & French, K. R. Industry costs of equity. *J. Financ. Econ.* **43**, 153–193 (1997).
38. Johnsson, K. *Structures in High-Dimensional Data: Intrinsic Dimension and Cluster Analysis* (Centre for Mathematical Sciences, Lund University, Lund, 2016).
39. Gonzalez, J., Low, Y., Gretton, A. & Guestrin, C. Parallel Gibbs sampling: from colored fields to thin junction trees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* 324–332 (2011).
40. Ceriotti, M., Tribello, G. A. & Parrinello, M. Simplifying the representation of complex free-energy landscapes using sketch-map. *Proc. Natl. Acad. Sci.* **108**, 13023–13028 (2011).
41. Carlsson, G. Topology and data. *Bull. Am. Math. Soc.* **46**, 255–308 (2009).
42. Zomorodian, A. & Carlsson, G. Computing persistent homology. *Discrete Comput. Geom.* **33**, 249–274 (2005).

Acknowledgements

M.A. was supported by Horizon 2020 FLAG-ERA (European Union), grant ANR-17-HBPR-0001-02 during completion of this work. A.M. acknowledges the financial support of SNF 100018_172892. We thank Giovanni Barone Adesi and Julia Reynolds (Institute of Finance, USI, Lugano, Switzerland) for helping us with the Compustat dataset and offering precious suggestions for its analysis. We thank Giulia Sormani (SISSA, Trieste, Italy) for suggesting, and providing us with the protein dynamics data. We thank Alex Rodriguez (SISSA, Trieste, Italy) for helping with the analysis of protein data. M.A. thanks the SISSA community at large for the substantial moral and intellectual support received, which has been critical for completion of this work.

Author contributions

M.A., E.F., F.D., A.L. and A.M. designed research; M.A., E.F., F.D., A.L. and A.M. performed research; M.A., E.F., F.D., A.L. and A.M. analyzed data; M.A., E.F., F.D., A.L. and A.M. wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-72222-0>.

Correspondence and requests for materials should be addressed to A.L. or A.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020