



HAL
open science

Local intrinsic dimensionality estimators based on concentration of measure

Jonathan Bac, Andrei Zinovyev

► **To cite this version:**

Jonathan Bac, Andrei Zinovyev. Local intrinsic dimensionality estimators based on concentration of measure. 2020 International Joint Conference on Neural Networks (IJCNN), Jul 2020, Glasgow, United Kingdom. pp.1-8, 10.1109/IJCNN48605.2020.9207096 . hal-02972293

HAL Id: hal-02972293

<https://hal.science/hal-02972293>

Submitted on 20 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Local intrinsic dimensionality estimators based on concentration of measure

Jonathan Bac*

Institut Curie, PSL Research University,
Mines ParisTech, INSERM, U900
Centre de Recherches Interdisciplinaires,
Paris, France
jonathan.bac@cri-paris.org

Andrei Zinovyev†

Institut Curie, PSL Research University,
Mines ParisTech, INSERM, U900
Paris, France
Lobachevsky University
Nizhni Novgorod, Russia
andrei.zinovyev@curie.fr

Abstract—Intrinsic dimensionality (ID) is one of the most fundamental characteristics of multi-dimensional data point clouds. Knowing ID is crucial to choose the appropriate machine learning approach as well as to understand its behavior and validate it. ID can be computed globally for the whole data distribution, or estimated locally in a point. In this paper, we introduce new local estimators of ID based on linear separability of multi-dimensional data point clouds, which is one of the manifestations of concentration of measure. We empirically study the properties of these measures and compare them with other recently introduced ID estimators exploiting various other effects of measure concentration. Observed differences in the behaviour of different estimators can be used to anticipate their behaviour in practical applications.

Index Terms—high-dimensional data, intrinsic dimension, effective dimension, concentration of measure, separability

I. INTRODUCTION

Datasets used in applications of machine learning frequently contain objects characterized by thousands and even millions of features. In this respect, the well-known *curse of dimensionality* is frequently discussed which states that many problems become exponentially difficult in high dimensions [1]. However, in what concerns the application of machine learning methods, the curse of dimensionality is not automatically manifested when the number of features is large: this depends rather on the dataset’s intrinsic dimensionality (ID). If the features of a dataset are correlated in linear or non-linear fashion then the data point cloud can be located close to a subspace of relatively low ID. This makes appropriate the application of dimension reduction algorithms to obtain a lower-dimensional representation of data. By contrast, if the value of ID is high then the data point cloud becomes sparse and its geometrical and topological properties can be highly

non-intuitive, in particular, due to various manifestations of concentration of measure [2].

Therefore, the estimation of ID is crucial to the choice of machine learning methodology and its applications, including validation, explainability and deployment. Indeed, ID determines to a large extent the strategy and feasibility of validating an algorithm, estimating the uncertainty of its predictions and explaining its decisions, which can be even legally required for its use in sensitive applications [3]. There has been recent progress in this area of research, e.g. with several approaches to explain and validate classifiers by looking at various data properties related to ID [4], [5]. Estimators of intrinsic dimension have recently been applied to the question of generating better explanations for a classifier’s behavior [3]. ID can pose fundamental limits on the robustness of classifiers; this has been illustrated recently with adversarial examples, where a minimal perturbation to the input can lead to a misclassification. Recent theoretical results have shown that such adversarial examples are inevitable for a concentrated metric probability space [6]–[10] and algorithms have been proposed that estimate adversarial risk by quantifying concentration [11]. From a more general perspective, theory and algorithms for correcting artificial intelligence-based systems have been developed that exploit properties specific to spaces possessing large ID [12], [13]. In this respect, a complementarity principle has been formulated [14]–[16]: the data space can be split into a low volume (low dimensional) subset, where nonlinear methods are effective, and a high-dimensional subset, characterized by measure concentration and simplicity, allowing the effective application of linear methods.

One important observation is that in real life the ID of the complete dataset might not be equal to the ID of its parts. Therefore, ID can be considered a local characteristic of the data space, defined in each data neighborhood. In this case we refer to it as *local intrinsic dimensionality*. Even when there are no ID variations in a dataset, the relation between global and local intrinsic dimensionality can be non-trivial - for example, one can easily construct examples of datasets which are low-dimensional locally but globally possess large

¹ To whom correspondence should be addressed. This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), the Ministry of education and science of Russia (Project No. 14.Y26.31.0022), the Association Science et Technologie, the Institut de Recherches Internationales Servier and the doctoral school Frontières de l’Innovation en Recherche et Education–Programme Bettencourt.

linear dimension [16]. Methods for estimating global ID can be applied locally in a data neighbourhood, and, vice versa, local ID estimators can be used to derive a global ID estimate. However, local ID estimation presents specific challenges, such as dealing with restricted cardinality, which render many global methods ineffective in practice. These challenges explain the rise in recent works dedicated specifically to developing local estimators [17]–[22].

Well-known methods for estimating global or local ID are based on Principal Component Analysis (PCA) and quantifications of the covariance matrix’s eigenspectrum using various heuristics (thresholding total explained variance, limiting the conditional number, using reference spectra such as broken stick distribution, etc.) [23]–[26]. Other famous examples include maximum likelihood estimation (MLE) and the correlation dimension, based on counting the number of objects in a growing neighbourhood [27], [28]. Since PCA is a linear method, it generally tends to overestimate ID, while MLE and correlation dimension both tend to underestimate larger ID values.

In order to overcome these limitations, several new ID estimators have been recently introduced that exploit concentration of measure [14], [19]–[22], [29], [30]. Increasing ID results in various manifestations of measure concentration that scale differently with respect to dimensionality. For example, linear separability of points increases rapidly, such that 30 dimensions can be already considered large, while the appearance of exponentially large quasi-orthogonal bases or hubness are usually manifested in higher dimensions [1], [31]. These differences lead to different properties of ID estimators, e.g. sensitivity to different dimensions and dependence of the estimated ID on sample size.

As noted in [20], most of the local ID estimators use pairwise distances (relationship between two points) or angular information (relationship between three points) [21], [28], [29], whereas others take into account the volume of a d -simplex (i.e., the relationship between $d + 1$ points) [19], [20] or in our case, the separability probability of each point from the others [22]. The common approach is to assume local distributions of data points are close to a uniformly sampled unit n -ball B^n or their scaled vectors to a uniformly sampled unit sphere S^{n-1} . Then, various sample statistics are used whose dependence on n is theoretically established for uniform distributions on B^n or S^{n-1} . If the estimated statistics of a data sample are similar to the theoretical of B^n or S^{n-1} then the dimensionality of the sample is estimated to be n .

For example, DANCo [29] estimates the probability density function of normalized nearest neighbor distance from the center (exploiting concentration of norms), and the parameters of a Von-Mises distribution (exploiting concentration of angles). ESS [19], in its default version, computes simplex skewness, defined as the ratio between the volume of a simplex with one vertex in the centroid and the others in data points, and the volume this simplex would have if edges incident to the centroid were orthogonal. For a d -simplex with $d = 1$, skewness is $\sin(\theta)$, with θ the angle between the two edges

incident to the centroid vertex. The mean sample skewness is compared to the Expected Simplex Skewness for uniformly distributed data on B^n . ANOVA [21] uses a U -statistic for the variance of the angle between pairs of vectors among uniformly chosen points in S^{n-1} .

In our previous work we introduced and benchmarked global estimators of ID based on Fisher separability, using theoretical results by Gorban, Tyukin et al. [12], [22]. The details of this approach are provided below. It appeared that for noisy samples from synthetic manifolds, the method was competitive with other ID estimators. In particular, its behavior was close to two recently introduced ID estimators based on concentration of measure, namely ESS and DANCo. In this work we extend our previous study and introduce two ways to estimate ID locally, in each data point, based on Fisher separability properties. We compare the properties of these local dimensionality estimators with other estimators based on quantifying various manifestations of concentration of measure.

II. LOCAL ID ESTIMATION BASED ON FISHER SEPARABILITY

In the present work, we will follow the notations introduced in the works by A.Gorban, I.Tyukin and their colleagues [12]: we call a data vector $\mathbf{x} \in R^n$ *linearly separable* from a finite set of points $Y \subset R^n$ if there exists a linear functional l such that $l(\mathbf{x}) > l(\mathbf{y})$ for all $\mathbf{y} \in Y$. If for any point \mathbf{x} there exists a linear functional separating it from all other data points, then such a data point cloud is called *linearly separable* or 1-convex. The separating functional l may be computed using the linear Support Vector Machine (SVM) algorithms, the Rosenblatt perceptron algorithm, or other comparable methods. However, these computations may be rather costly for large-scale estimates. Hence, in the works of Gorban, Tyukin and their colleagues it was suggested to use a non-iterative estimate of the linear functional using Fisher’s discriminant which is computationally inexpensive, after a standard pre-processing: [12].

- 1) centering
- 2) performing linear dimensionality reduction by projecting the dataset into the space of k first principal components, where k may be relatively large. In practice, we select the largest k (in their natural ranking) such that the corresponding eigenvalue λ_k is not smaller than λ_1/C , where C is a predefined threshold. Under most circumstances, $C = 10$
- 3) whitening (i.e., applying a linear transformation after which the covariance matrix becomes the identity matrix)

After such normalization of X , it is said that a point $\mathbf{x} \in X$ is Fisher-linearly separable from the cloud of points Y with parameter α , if

$$(\mathbf{x}, \mathbf{y}) \leq \alpha(\mathbf{x}, \mathbf{x})$$

for all $\mathbf{y} \in Y$, where $\alpha \in [0, 1)$. If equation (II) is valid for each point $\mathbf{x} \in X$ such that Y is the set of points $\mathbf{y} \neq \mathbf{x}$

then we call the dataset X Fisher-separable with parameter α . In order to quantify deviation from perfect separability, let us introduce $p_\alpha(\mathbf{y})$, the probability that the point \mathbf{y} is inseparable from all other points.

In order to associate a value of ID to a point \mathbf{y} in the data space, we compare the empirical $p_\alpha(\mathbf{y})$ estimates to the p_α of some reference data distribution whose dimension is known and separability properties can be analytically derived. The simplest such distribution is the uniform distribution of vectors on the surface of a unit n -dimensional sphere. Since this distribution is uniform, p_α does not depend on a data point and equals in any point (see derivation in [1], [12]):

$$p_\alpha \approx \frac{(1 - \alpha^2)^{\frac{n-1}{2}}}{\alpha \sqrt{2\pi n}} \quad (1)$$

By resolving this formula with respect to n , we derive the following value of dimensionality as a function of inseparability probability, for a uniform distribution on the surface of the unit n -dimensional sphere:

$$n_\alpha = \frac{W\left(\frac{-\ln(1-\alpha^2)}{2\pi p_\alpha^2 \alpha^2 (1-\alpha^2)}\right)}{-\ln(1-\alpha^2)} \quad (2)$$

where $W(x)$ is the real-valued branch of the Lambert W function. As a reminder, the Lambert W function solves equation $v = we^w$ with respect to w , i.e. $w = W(v)$. The pseudo-code of the algorithm for computing n_α was provided by us earlier [22].

Here we should make several important notes. Firstly, in order to apply the ID estimate (2) to a dataset X , one should apply an additional scaling step besides the preprocessing steps (1)-(3) described above. The scaling consists in normalizing each vector to the unit length, which corresponds to the projection onto a unit sphere. It means that in practice we do not distinguish an n -ball B^n from n -sphere S^{n-1} , both give $ID=n$ in our case. This can lead to shifting by value 1 the estimations of ID in small dimensions, especially in artificial benchmark examples.

Secondly, as mentioned in [1], the formula (1) is an estimation from above, meaning that the actual empirical value of p_α is strictly less than the right hand side of (1). In particular, for some points, data point density, and values of α one can have empirical estimate $p_\alpha = 0$ which makes (2) inapplicable. The value of α should be adjusted in order to avoid the mean of p_α being too close to zero and in order to avoid too strong finite sampling effects (see [22] and Figure 1). This consideration also provides some theoretical limits on the maximally detectable dimensionalities, as shown hereafter in Figure 2.

Based on the above definitions, the separability properties of the data point cloud can be globally characterized by the histogram of empirical p_α distribution (probabilities of individual point inseparability) and the profile of intrinsic dimensions n_α (2) for a range of α values (e.g., $\alpha \in [0.6, \dots, 1.0]$).

Let us denote $\bar{p}_\alpha(\mathbf{X})$ the mean value of the distribution of $p_\alpha(\mathbf{x})$ over all data points. We can introduce the global estimate of ID as in our previous work ([22]):

$$n_\alpha^{global} = \frac{W\left(\frac{-\ln(1-\alpha^2)}{2\pi \bar{p}_\alpha^2 \alpha^2 (1-\alpha^2)}\right)}{-\ln(1-\alpha^2)} \quad (3)$$

Now let us specify two localized versions of ID estimate based on Fisher separability. The first one will simply use the formula (2) in order to estimate n_α in a data point \mathbf{y} . We will call this estimate *global pointwise intrinsic dimension*, since in its definition the global separability properties of X are exploited, but the ID is computed in a data point:

$$n_\alpha(\mathbf{y}) = \frac{W\left(\frac{-\ln(1-\alpha^2)}{2\pi p_\alpha(\mathbf{y})^2 \alpha^2 (1-\alpha^2)}\right)}{-\ln(1-\alpha^2)} \quad (4)$$

At the same time we can follow the standard approach to define data neighborhoods and compute the n_α^{global} for these fragments of data. One of the simplest way to define a local neighbourhood is to determine the first k nearest neighbours (kNN) for a data point. This naive approach has well-known drawbacks, such as the edge effect : the ratio between points close to the border of the manifold and points inside it increases with dimensionality, meaning data neighbourhoods will not be uniformly distributed. [29], [32]. Forming kNN using distances of the original space can also create neighborhoods that do not reflect geodesic distance on the manifold. Recent work has suggested new ways to tackle these issues [17], [18]. Nonetheless this basic approach provides an easy way to start applying methods locally. Therefore, we define *local kNN ID* as:

$$n_\alpha^{kNN}(\mathbf{y}) = \frac{W\left(\frac{-\ln(1-\alpha^2)}{2\pi p_\alpha^{kNN}(\mathbf{y})^2 \alpha^2 (1-\alpha^2)}\right)}{-\ln(1-\alpha^2)} \quad (5)$$

where $p_\alpha^{kNN}(\mathbf{y})$ is computed for a dataset comprising k nearest neighbours of \mathbf{y} . In practical applications, k is chosen of few hundreds by the order of magnitude.

III. NUMERICAL RESULTS

A. Estimation of ID for n -dimensional balls

We illustrate our approach for a simple example of n -ball B^n (see Figure 1), sampled by 2000 points. The empirically estimated values of p_α are shown together with the theoretical dependence (1) in Figure 1A for $n = 2 \dots 10$. Figure 1B shows the empirically estimated distributions of p_α for $n = 2 \dots 10$. In Figure 1C we show the empirically estimated value of n as a function of α . One can see that smaller values of α create biased estimates of n . At the same time, maximum possible values of α (for which the mean of p_α is still non-zero) creates unstable estimates suffering from finite sampling effects, starting from $n = 7$. Therefore, for the unique definition of n , our rule of thumb is to select p_α which would equal 0.8 multiplied by the maximally measured p_α (shown by crosses in Figure 1C).

Finite sample effect poses theoretical limits on maximum measurable values of ID from separability analysis as a function of the number of points N . The minimum measurable p_α equals¹ $\frac{1}{N}$: therefore, maximally measurable $n_\alpha(y)$ (4) equals $\frac{W(\frac{-\ln(1-\alpha^2)N^2}{2\pi\alpha^2(1-\alpha^2)})}{-\ln(1-\alpha^2)}$ (shown in In Figure 1D). At the same time, the mean of p_α is non-zero unless the dataset will become completely Fisher-separable. This defines the limit for measuring (3) in $\frac{W(\frac{-\ln(1-\alpha^2)N^4}{2\pi\alpha^2(1-\alpha^2)})}{-\ln(1-\alpha^2)}$ (shown in Figure 1E) since the minimally measurable \bar{p}_α can be estimated as $\frac{1}{N^2}$. This conclusion allows us to estimate the maximally measurable $n_\alpha^{kNN}(y)$ (5) in $\frac{W(\frac{-\ln(1-\alpha^2)k^4}{2\pi\alpha^2(1-\alpha^2)})}{-\ln(1-\alpha^2)}$ which can be read from Figure 1E for cardinality equals k . Notice from Figure 1D,E that the maximum measurable ID quickly saturates with the number of points N but remains relatively high for smaller α .

For the n -ball example, we also studied how well different estimators can characterize the ID of high dimensional samples, depending on the number of points in a sample (see Figure 2). In the case of this simple benchmark, ESS and DANCo show impressive results, giving the exact value of ID even for very small samples (ESS is even able to estimate dimensionality higher than the number of points, which seems to be counter-intuitive at first glance but this is discussed in the paper [19]). Fisher separability analysis tends to overestimate the ID in case of small size samples but work as well as DANCo and ESS for larger number of points. The overestimation is connected to the just discussed point of maximally measurable dimensionality for a given α value. In case of a small sample in very high dimension, the value of α , adjusted using the heuristics described above, becomes small in order to avoid full separability: in this case, the estimate (3 becomes less accurate. Interestingly, many widely used ID estimators, not based on concentration of measure such as Correlation (Fractal) Dimension heavily underestimate the ID, showing saturation at few tens of dimensions. ANOVA estimator showed surprising underestimation of larger ID even for large number of points (Figure 2).

B. Real-life examples

In Figure 3A,B we show visualization of ID values for a simple synthetic '10 balls' dataset, which represents 10 n -balls of increasing $n = 2..11$ embedded without intersection in 11-dimensional space such that the n -ball shares $n-1$ dimensions with the $n-1$ -ball. From this example, one can see that both pointwise global and local kNN ID based on Fisher separability characterize well the dimensionality of the balls, however, pointwise global estimate is more heterogeneous within a ball and the local kNN estimate suffers from the 'edge effects'. In Figure 3,C on several real-life datasets we visualize the ID values derived from several local dimensionality estimators exploiting the concentration of measure phenomena, including those based on Fisher separability. First conclusion is that the local ID can be highly variable from one data cluster

to another. For example, for the MNIST datasets, local ID highlights some clusters which possess much lower local ID than the others (digit '1' - for the MNIST digit dataset and also '7', letters 'i','l' and 'j' for the MNIST letters dataset, and cluster 'Trousers' for the fashion MNIST dataset). This probably reflects the intrinsic number of degrees of freedom in the distribution of different realizations (e.g., writing or design) of the same object type (e.g., a digit). At the same time, in case of the standard faces dataset (top row), the distribution of local ID estimations is more uniform and close to 3 as expected.

In Figure 4 we demonstrate possible relations between global pointwise and local kNN ID estimates based on Fisher separability. Global pointwise ID estimate is computed based on the separability of a point from the rest of the data point cloud, while local kNN ID estimate is based on the inseparability from the points in the local neighbourhood. These two estimates do not have to be similar: strong deviations of one from another can indicate complex data topology, when distant parts of the data point cloud are co-localized in close data subspaces. Local inseparability usually imposes the global inseparability but not vice versa which leads to characteristic triangular patterns in Figure 4: however, this is not necessarily universal trend as it is the case of isomapFaces and FMINST datasets, which hints on strong data non-linearity or the inadequate choice of the neighbourhood.

C. Dependence of ID estimators on the subsample size

A good property of an ID estimator, both global and local, is ability to quantify a wide range of ID values from a relatively small sample size. In order to study these properties, we took several synthetic and real-life datasets and computed the ID based on four ID estimators, on different subsample sizes. Running these estimations multiple time on different subsamples gave us also an idea about uncertainty in the ID estimation for a given subsample sizes.

The results of this analysis are shown in Figure 5. Notice that for the uniformly sampled S^{10} sphere FisherS determines 11 intrinsic dimensions (as we noticed earlier, FisherS does not distinguish a ball and a sphere). Interestingly, ESS gives global ID=11 while mean local ID equals 10 in the case of uniformly sampled S^{10} . For other uniformly sampled simple distributions (hypercube, multivariate Gaussian), 3 out of 4 estimators provides consistent results but TwoNN significantly underestimate the dimension. For Swiss Roll example, global estimators (DANCo and TwoNN) work well while ESS and FisherS show strong dependence of the mean local ID on the size of the subsample. The mean pointwise global estimate is usually biased towards smaller ID. For all estimators the dependence of the local ID on the sample size can be quite strong and sometimes even not monotonous. Overall, it seems that DANCo gives ID estimates closer to TwoNN while ESS matches better FisherS, even though ESS tends to overestimate ID compared to FisherS in real-life data (for example, in case of ISOMAP Faces dataset, FisherS estimation in 3 or 4 compares with ID \approx 7 for ESS). We can also conclude that

¹We neglect small difference between N and $N-1$

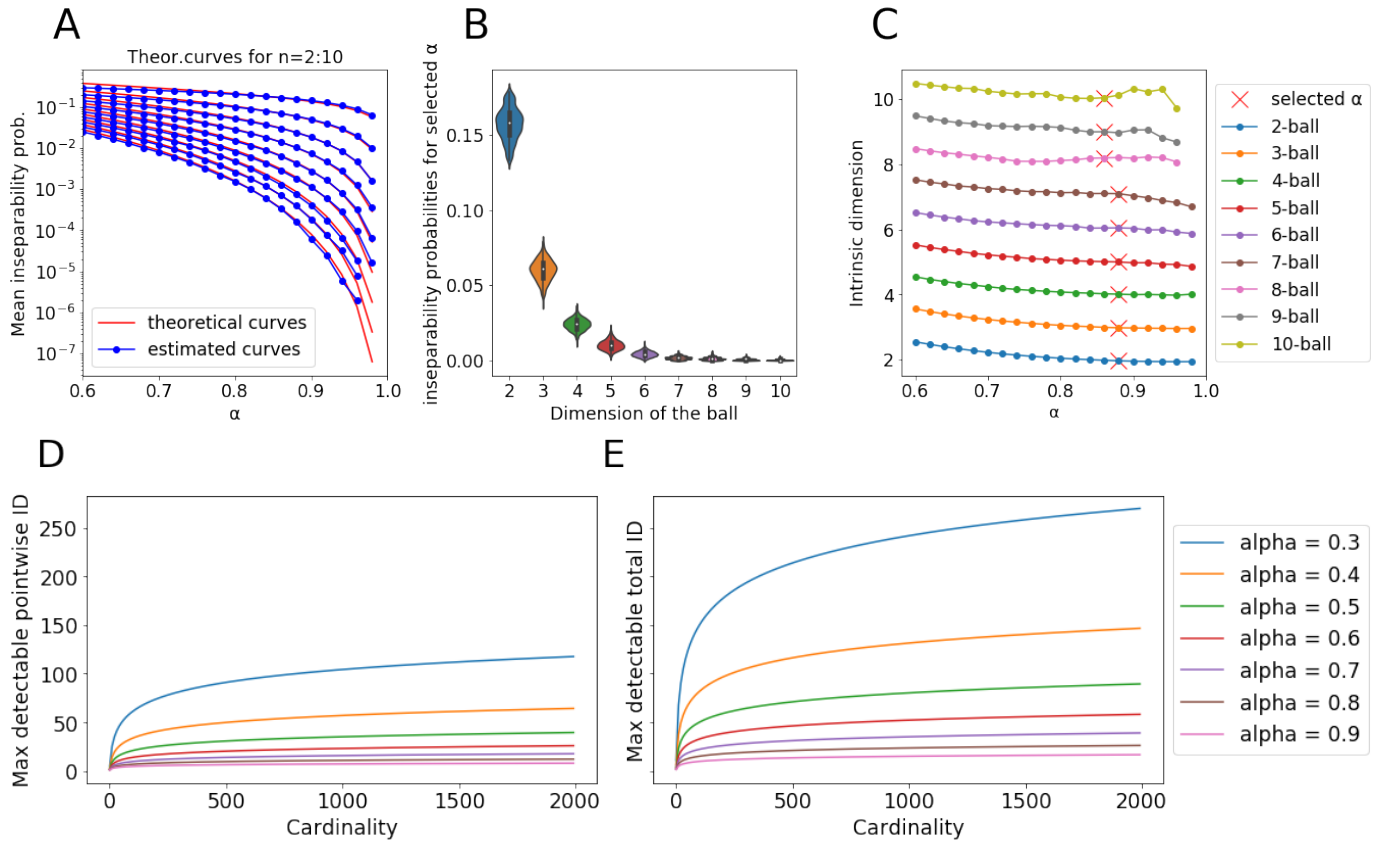


Fig. 1. **Top row**, left to right: theoretical vs estimated mean inseparability probability for balls of increasing dimension; complete inseparability probability histograms for selected $\alpha = .88$; intrinsic dimension estimated from the mean inseparability probability. Due to measure concentration, inseparability probability sharply decreases as a function of dimension. **Bottom row**: maximum detectable pointwise intrinsic dimension, maximum detectable mean pointwise intrinsic dimension

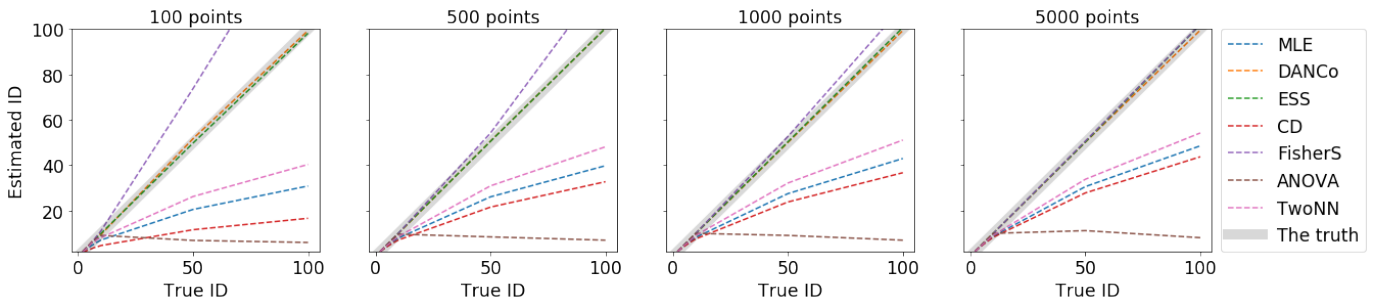


Fig. 2. Ability of various estimators to quantify high intrinsic dimensionality from different sample size, tested on the simplest n -ball example. Here CD stands for the Correlation Dimension, FisherS is the global ID estimation based on Fisher separability, MLE stands for Maximum Likelihood Estimation of Intrinsic Dimension [28]. Other estimator names are referenced in the Introduction.

the local ID estimates match better between the methods than the global ones.

IV. IMPLEMENTATION

We provide a Python 3 package and Jupyter notebooks replicating the analyses provided in this study at the following address: <https://github.com/j-bac/id-concentration>. In the future we plan to provide a Python library implementing most of the published intrinsic dimension estimators in a uniform fashion.

V. CONCLUSION

We characterized several properties of local intrinsic dimension estimators based on concentration of measure effects. Despite they all rely on mathematically related phenomena, their properties appear to be different from one to another even in the simplest benchmarks. In continuation of our previous work, we introduced two new local ID estimators, based on quantifying the separability properties of a dataset and comparing these statistics to a reference uniform distribution

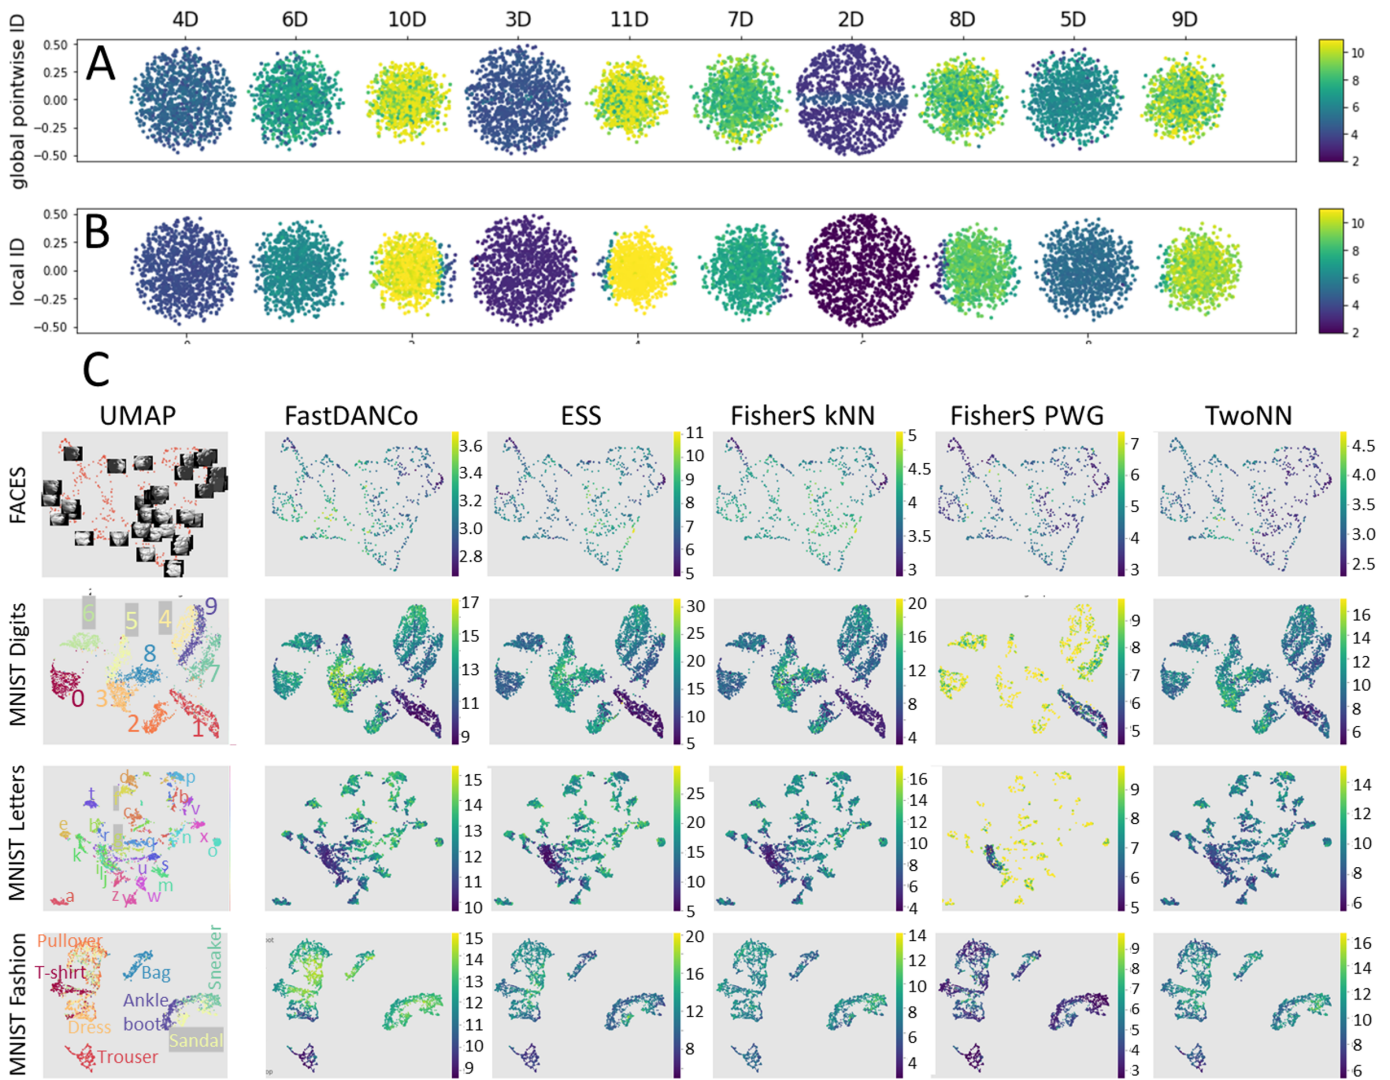


Fig. 3. Visualization of local ID measures on 2D layouts of synthetic (A,B) and real (C) datasets. For the MNIST datasets, the class labels are annotated directly on the UMAP image with colored labels located close to the corresponding clusters.

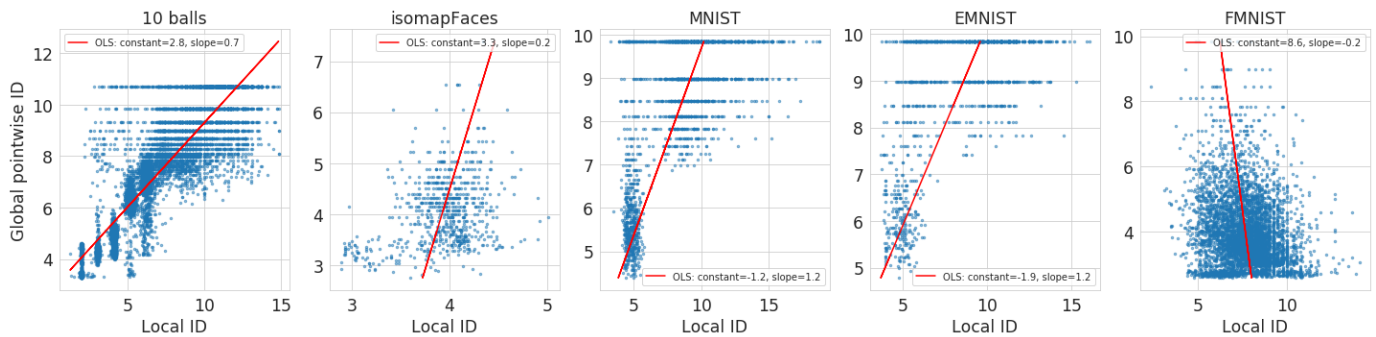


Fig. 4. Relation between global pointwise ID and local ID computed based on Fisher separability analysis. Linear regression line is shown by red.

on a unit n -dimensional sphere.

While concentration of measure affects many different statistics, current estimators introduce and exploit only one to produce their estimates. A notable exception is DANCo,

which combines both angle and norm concentration to provide a better final estimate. It remains an open question whether this approach can be extended to integrate the various statistics used by other estimators and produce a unified way of

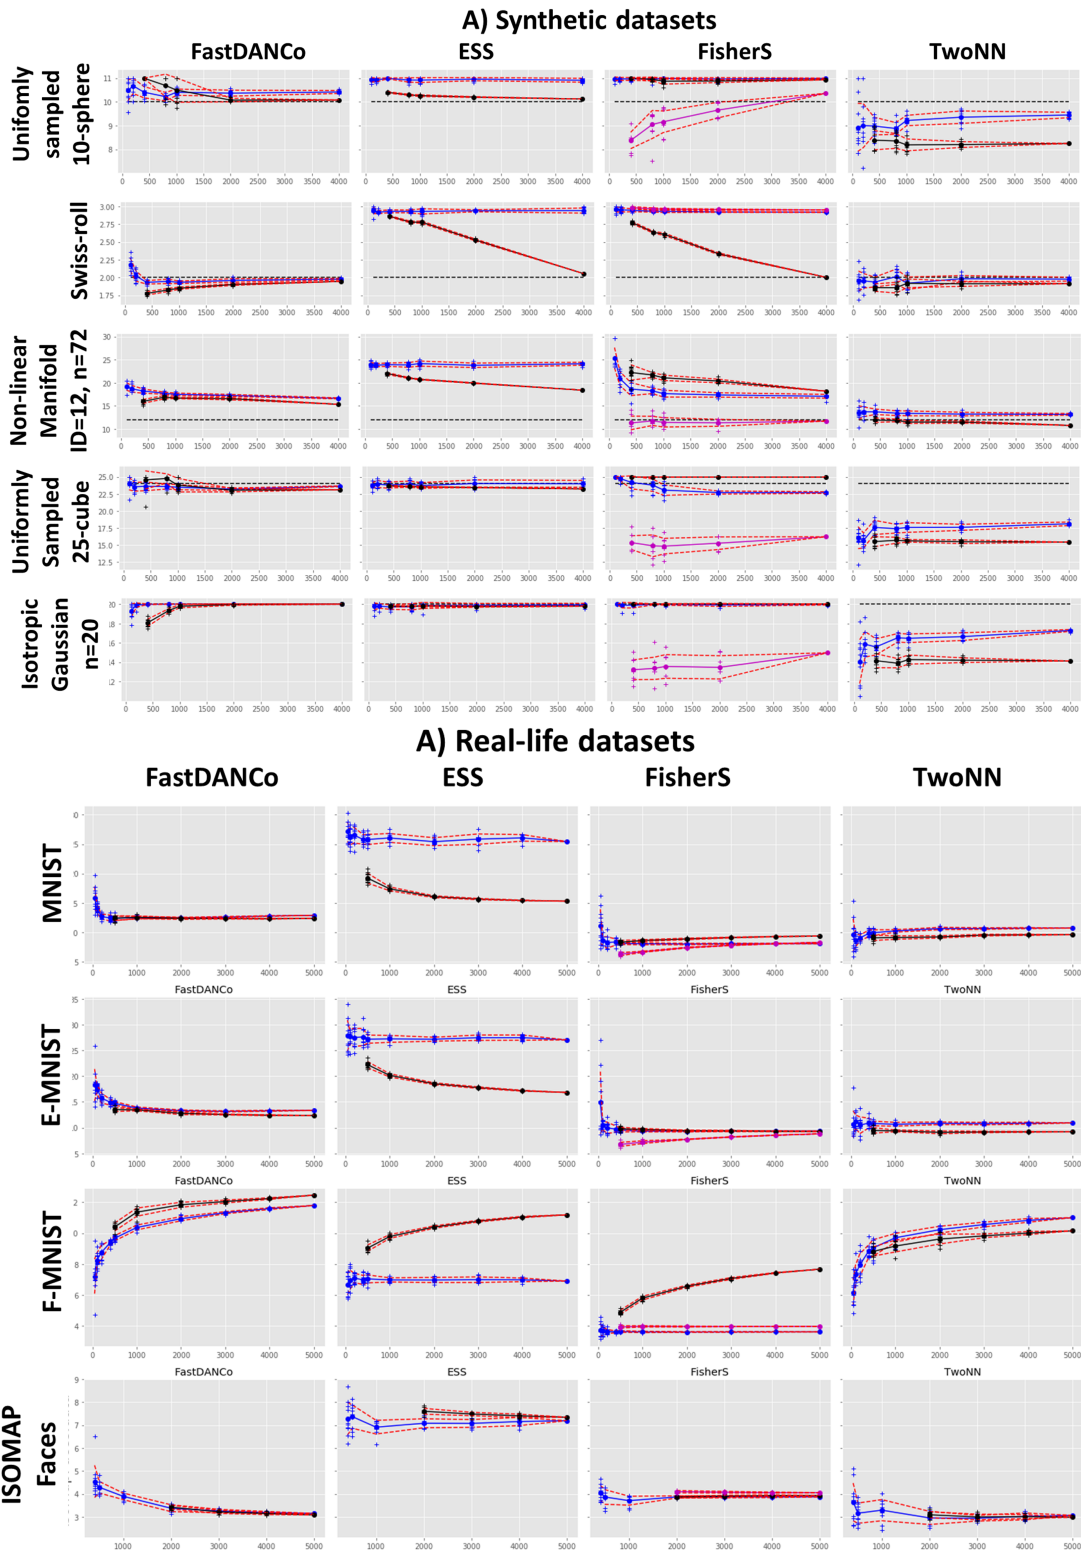


Fig. 5. Dependence of four ID estimators on subsample size. A selection of synthetic and real-life datasets used in this study is shown, the rest of the plots are available from <https://github.com/j-bac/id-concentration>. Blue line shows the global ID estimate, while the black one shows the mean local ID. Red dashed lines show confidence intervals. In case of FisherS, the purple line shows the mean global pointwise ID estimation.

characterizing intrinsic dimension from measure concentration effects. From our benchmarks, it seems that the local estimates

of ID tend to provide better consensus than the global ID estimate.

We can conclude that more efforts are needed to study the properties of intrinsic dimensionality, both global and local, estimated through different approaches. New generation of concentration of measure-based estimators provide relevant values of ID from small to very high dimensions: however, they need to be assessed and improved in order to deal with real-life datasets, characterized by non-uniformity and variability in the local ID values.

REFERENCES

- [1] A. N. Gorban, V. A. Makarov, and I. Y. Tyukin, "High-Dimensional Brain in a High-Dimensional World: Blessing of Dimensionality," *Entropy* 2020, Vol. 22, Page 82, vol. 22, no. 1, p. 82, jan 2020.
- [2] M. Gromov, "Isoperimetry of waists and concentration of maps," *Geom. Funct. Anal.*, vol. 13, pp. 178–215, 2003.
- [3] Y. Jia, J. Bailey, K. Ramamohanarao, C. Leckie, and M. E. Houle, "Improving the quality of explanations with local embedding perturbations," 2019.
- [4] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should i trust you?: Explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 2016, pp. 1135–1144.
- [5] H. Jiang, B. Kim, M. Guan, and M. Gupta, "To trust or not to trust a classifier," in *Advances in neural information processing systems*, 2018, pp. 5541–5552.
- [6] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.
- [7] J. Gilmer, L. Metz, F. Faghri, S. S. Schoenholz, M. Raghu, M. Wattenberg, and I. Goodfellow, "Adversarial spheres," *arXiv preprint arXiv:1801.02774*, 2018.
- [8] A. Fawzi, H. Fawzi, and O. Fawzi, "Adversarial vulnerability for any classifier," in *Advances in Neural Information Processing Systems*, 2018, pp. 1178–1187.
- [9] A. N. Bhagoji, D. Cullina, and P. Mittal, "Lower bounds on adversarial robustness from optimal transport," in *Advances in Neural Information Processing Systems*, 2019, pp. 7496–7508.
- [10] S. Mahloujifar, D. I. Diochnos, and M. Mahmoody, "The curse of concentration in robust learning: Evasion and poisoning attacks from concentration of measure," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 4536–4543.
- [11] S. Mahloujifar, X. Zhang, M. Mahmoody, and D. Evans, "Empirically measuring concentration: Fundamental limits on intrinsic robustness," *arXiv preprint arXiv:1905.12202*, 2019.
- [12] A. Gorban, A. Golubkov, B. Grechuk, E. Mirkes, and I. Tyukin, "Correction of ai systems by linear discriminants: Probabilistic foundations," *Information Sciences*, vol. 466, pp. 303 – 322, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0020025518305607>
- [13] I. Y. Tyukin, A. N. Gorban, A. McEwan, and S. Meshkinfamfard, "Bringing the blessing of dimensionality to the edge," in *2019 1st International Conference on Industrial Artificial Intelligence (IAI)*, July 2019, pp. 1–5.
- [14] A. N. Gorban and I. Y. Tyukin, "Blessing of dimensionality: mathematical foundations of the statistical physics of data," *Philos Trans A Math Phys Eng Sci*, vol. 376, no. 2118, Apr 2018.
- [15] A. Gorban, V. Makarov, and I. Tyukin, "Symphony of high-dimensional brain. Reply to comments on "The unreasonable effectiveness of small neural ensembles in high-dimensional brain"," *Phys. Life Rev.*, vol. 29, pp. 115–119, 2019.
- [16] J. Bac and A. Zinovyev, "Lizard Brain: Tackling Locally Low-Dimensional Yet Globally Complex Organization of Multi-Dimensional Datasets," *Frontiers in Neurobotics*, vol. 13, jan 2020. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fnbot.2019.00110/full>
- [17] L. Amsaleg, O. Chelly, M. E. Houle, K.-i. Kawarabayashi, M. Radovanović, and W. Treeratanajaru, "Intrinsic dimensionality estimation within tight localities," in *Proceedings of the 2019 SIAM International Conference on Data Mining*. SIAM, 2019, pp. 181–189.
- [18] O. Chelly, M. E. Houle, and K.-i. Kawarabayashi, "Enhanced estimation of local intrinsic dimensionality using auxiliary distances," Technical report, National Institute of Informatics, Tech. Rep., 2016.
- [19] K. Johnsson, "Structures in high-dimensional data: Intrinsic dimension and cluster analysis," Ph.D. dissertation, Faculty of Engineering, LTH, 8 2016.
- [20] D. R. Wissel, "Intrinsic dimension estimation using simplex volumes," Ph.D. dissertation, 2018. [Online]. Available: <http://hss.ulb.uni-bonn.de/2018/4951/4951.htm>
- [21] M. Díaz, A. J. Quiroz, and M. Velasco, "Local angles and dimension estimation from data on manifolds," 2018.
- [22] L. Albergante, J. Bac, and A. Zinovyev, "Estimating the effective dimension of large biological datasets using fisher separability analysis," in *Proc. IEEE IJCNN 2019 - International Joint Conference on Neural Networks*. Budapest Hungary: IEEE, 2019, pp. N–19 814.
- [23] K. Fukunaga and D. Olsen, "An Algorithm for Finding Intrinsic Dimensionality of Data," *IEEE Transactions on Computers*, vol. C-20, no. 2, pp. 176–183, feb 1971. [Online]. Available: <http://ieeexplore.ieee.org/document/1671801/>
- [24] J. Bruske and G. Sommer, "Intrinsic dimensionality estimation with optimally topology preserving maps," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 20, no. 5, pp. 572–575, 1998.
- [25] M. Fan, N. Gu, H. Qiao, and B. Zhang, "Intrinsic dimension estimation of data by principal component analysis," Tech. Rep., 2010. [Online]. Available: <https://pdfs.semanticscholar.org/7d4d/936df2f628550123cef40996277628468e61.pdf>
- [26] A. V. Little, J. Lee, Y.-M. Jung, and M. Maggioni, "Estimation of intrinsic dimensionality of samples from noisy low-dimensional manifolds in high dimensions with multiscale SVD," in *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*. IEEE, 2009, pp. 85–88.
- [27] P. Grassberger and I. Procaccia, "Measuring the strangeness of strange attractors," *Physica D: Nonlinear Phenomena*, vol. 9, no. 1-2, pp. 189–208, oct 1983. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0167278983902981>
- [28] E. Levina and P. J. Bickel, "Maximum Likelihood estimation of intrinsic dimension," pp. 777–784, 2004. [Online]. Available: <https://dl.acm.org/citation.cfm?id=2976138>
- [29] C. Ceruti, S. Bassis, A. Rozza, G. Lombardi, E. Casiraghi, and P. Campadelli, "DANCo: An intrinsic dimensionality estimator exploiting angle and norm concentration," *Pattern Recognition*, vol. 47, no. 8, pp. 2569–2581, aug 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S003132031400065X>
- [30] E. Facco, M. D'Errico, A. Rodriguez, and A. Laio, "Estimating the intrinsic dimension of datasets by a minimal neighborhood information," *Scientific Reports*, vol. 7, no. 1, p. 12140, dec 2017. [Online]. Available: <http://www.nature.com/articles/s41598-017-11873-y>
- [31] A. N. Gorban, I. Y. Tyukin, D. V. Prokhorov, and K. I. Sofeikov, "Approximation with random bases: Pro et Contra," *Information Sciences*, vol. 364-365, pp. 129–145, oct 2016.
- [32] P. J. Verveer and R. P. W. Duin, "An evaluation of intrinsic dimensionality estimators," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 17, no. 1, pp. 81–86, 1995.