

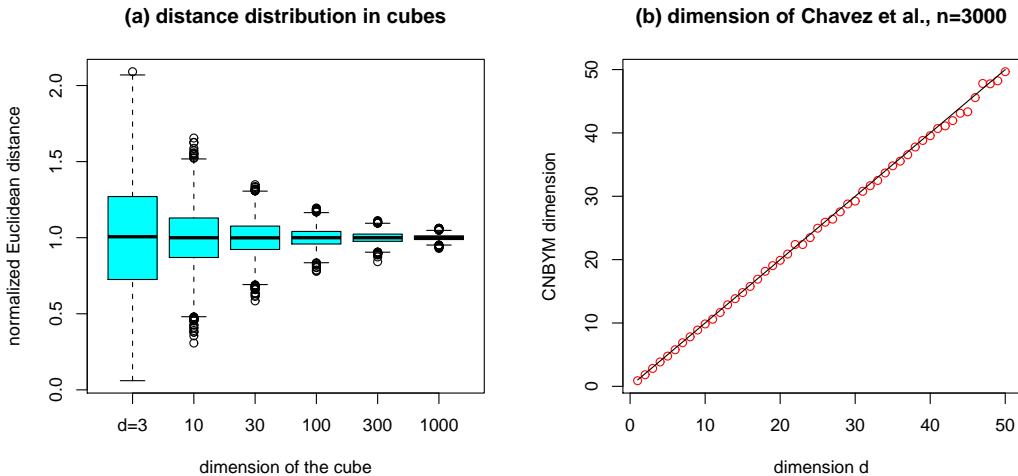
Intrinsic Dimensionality

Vladimir Pestov

Department of Mathematics and Statistics, University of Ottawa, Ottawa, Ontario, Canada

Nature of the concept. Intrinsic dimensionality of data does not refer to a single well-defined parameter, but rather a diverse family of parameters associated to a given similarity workload $(\mathbb{X}, d, \mathbb{U})$ and allowing to gauge the performance of various indexing schemes for similarity-based information retrieval.

Perhaps the best known single feature of a high-dimensional dataset \mathbb{U} is “the empty space paradox”: the average distance, $\mathbb{E}_{\mathcal{E}_{NN}}(X)$, from a query point to the nearest neighbour in \mathbb{U} grows along with the dimension of \mathbb{U} . Asymptotically, $\mathbb{E}_{\mathcal{E}_{NN}}(X)$ approaches the average distance $\mathbb{E}[d(X, Y)]$ between two points of \mathbb{X} (the *characteristic size* of dataset). This is a consequence of the fact that the distance functions $d_x(y) = d(x, y)$, $x \in \mathbb{X}$, concentrate near their median/mean values. The boxplot diagram (a) shows (normalized) pairwise distance distributions of 100 uniformly sampled points in the cube \mathbb{I}^d for various values of d . The lines in the middle of the boxes are median values, the lower and upper sides denote the first and third quartiles respectively, and the small circles mark outliers. For high d , even the outlier distance values approach the median.

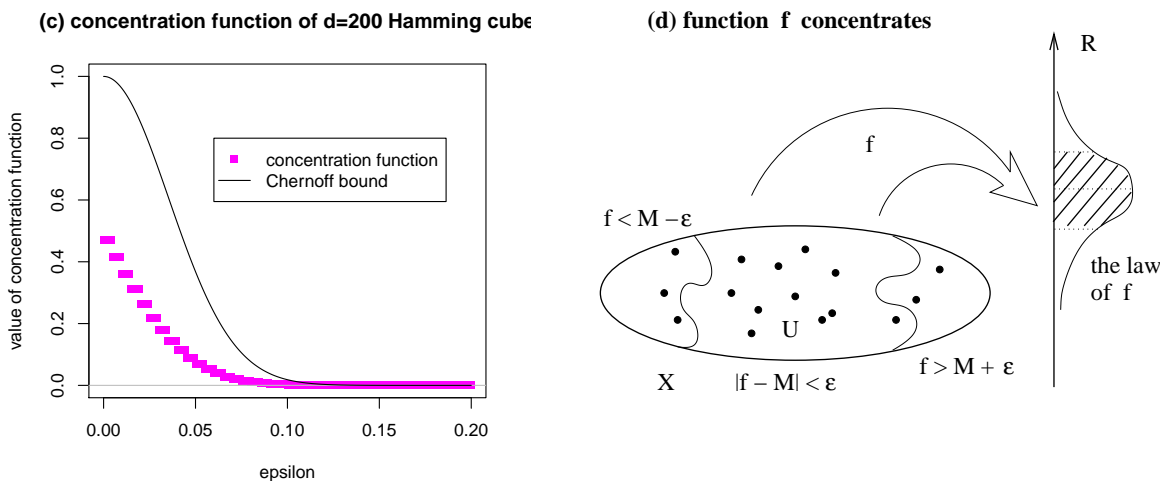


Generally, the intrinsic dimensionality is inversely proportional to some dispersion parameter of data, usually normalized with regard to the characteristic size of \mathbb{U} . The possible choices of a dispersion parameter are many, and they reflect different aspects of the dataset and its indexability.

Illustration: the intrinsic dimensionality of Chávez et al. The authors of [1] have proposed and studied a parameter which we denote by the acronym of their last names: $\text{dim}_{\text{CNBYM}}(\mathbb{X}) = \mathbb{E}(d)^2 / 2\text{var}(d)$. Here the distance d is treated as a random variable on $\mathbb{X} \times \mathbb{X}$. The corresponding statistical parameter of a dataset \mathbb{U} has a number of advantages: easy to compute, it makes a good statistical estimator for the dimension of the domain \mathbb{X} . The robustness of “CNBYM dimension” can be seen from diag. (b) where the parameter is calculated for independent random samples of $n = 3,000$ points drawn from the gaussian distribution γ^d on \mathbb{R}^d , $1 \leq d \leq 50$.

The curse of dimensionality. When applied to high-dimensional datasets, many similarity-based information retrieval algorithms slow down to a point where they cannot outperform a simple sequential scan. Even though the phenomenon is well-known to data practitioners, curiously enough there is still no mathematical validation of it, and one of the major theoretical challenges in the field is the so-called *Curse of Dimensionality Conjecture* in the bare-bone setting where $\mathbb{X} = \{0, 1\}^d$ is a Hamming cube [4]. Rigorous lower bounds are only obtained in a small number of cases, some of which we will describe below.

The concentration phenomenon. The curse of dimensionality is closely linked to the *phenomenon of concentration of measure* (cf. e.g. [3], Ch. 3 1/2): on a typical geometric object \mathbb{X} of high dimension not only the distance functions, but all the 1-Lipschitz functions $f: \mathbb{X} \rightarrow \mathbb{R}$ concentrate near their median. (A function is 1-Lipschitz if it does not increase distances: $|fx - fy| \leq d(x, y)$.) The *concentration function* $\alpha_{\mathbb{X}}(\varepsilon)$ of \mathbb{X} gives an upper bound on the probability that the value of a 1-Lipschitz function f at an $x \in \mathbb{X}$ deviates from its median by more than $\varepsilon > 0$. For instance, the concentration function of the d -dimensional Hamming cube $\{0, 1\}^d$ with the normalized Hamming metric and uniform measure satisfies a Chernoff bound $\alpha(\varepsilon) \leq \exp(-2\varepsilon^2d)$, cf. diagram (c). Similar gaussian bounds exist for Euclidean spheres, cubes, and a variety of other



objects. The higher the dimension of \mathbb{X} , the more sharply $\alpha_{\mathbb{X}}(\varepsilon)$ drops off near zero.

Degrading performance of indexing schemes. Many known access methods for *exact* similarity search (including pivot-based schemes, metric trees, etc.) have the following structure. An index-

ing scheme consists of a collection of 1-Lipschitz functions $\{f_a: a \in A\}$ on the domain \mathbb{X} (possibly partially defined). Given a query point $x \in \mathbb{X}$ and the range value $\varepsilon > 0$, calculations of the values $f_{a_i}(x)$, $i = 1, 2, \dots, \ell$ are performed in a recursive way, where the function $f_{a_{i+1}}$ is chosen depending on the previous values $f_{a_1}(x), \dots, f_{a_i}(x)$. If for some i one has $|f_{a_i}(y) - f_{a_i}(x)| > \varepsilon$, the datapoint y cannot possibly belong to the ε -ball around x and so is irrelevant. Datapoints $x_1, \dots, x_k \in \mathbb{U}$ whose irrelevance cannot be certified in this way are returned, and each of them is checked against the condition $d(x, x_i) < \varepsilon$.

If the domain \mathbb{X} has high concentration dimension, then most values of each function f_{a_i} concentrate near the median. Now denote \mathcal{F} the class of all Lipschitz functions used to construct indexing schemes of a given type (e.g. the class of all distance functions for pivot schemes). Assume that \mathcal{F} has a small capacity in the sense of statistical learning theory. (This condition is very natural and verified, e.g., by distance functions on the Euclidean spaces or Hamming cubes.) Assume further that the points of \mathbb{U} are sampled from \mathbb{X} in an i.i.d. fashion. Then the values of each $f \in \mathcal{F}$ at datapoints can be shown to concentrate as well, and only a small proportion of points can be discarded. cf. diag. (d). This argument leads to lower bounds on the performance of pivot indexing schemes which are linear in the size of the dataset [9]. Similar results can be deduced for metric trees.

One is naturally led to *concentration dimension* of \mathbb{X} , defined by $\dim_\alpha(\mathbb{X}) = 1 / \left[2 \int_0^1 \alpha_{\mathbb{X}}(\varepsilon) d\varepsilon \right]^2$ and investigated in [7]. Notice however that while $\dim_\alpha(\mathbb{X}) \geq \dim_\alpha(\mathbb{U})$, the relationship between the two is not understood: the latter value (dimension of the empirical measure) need not be a good statistical estimator for the former.

Concentration versus complexity. Heuristically, indexability of a workload $(\mathbb{X}, d, \mathbb{U})$ amounts to the existence of sufficiently many 1-Lipschitz functions f on \mathbb{U} which at the same time *dissipate* (as opposed to *concentrate*) and have a low computational complexity. Thus, indexing means finding the right balance between concentration and complexity.

Here is an example of a “success story” where smallness of values of an intrinsic dimension guarantees the existence of an efficient access method.

Assouad dimension. The *Assouad (doubling) dimension* of a metric space \mathbb{X} is the minimum value $\rho \geq 0$ such that every set A in X can be covered by 2^ρ balls of half the diameter. (The *diameter* of a set A is the supremum of $d(x, y)$, $x, y \in A$.) To appreciate that $\dim_{Assouad}(X)$ is indeed an inverse quantity to some parameter of dispersion of data, notice that if \mathbb{X} is equipped with a probability measure and $\dim_{CNBYM}(\mathbb{X})$ is high, then the size of balls as a function of radius grows fast in the vicinity of $\mathbb{E}(d)$, and so $\dim_{Assouad}(X)$ has to be large.

As shown in [5], a low Assouad dimension of a workload allows for the construction of a simple and efficient indexing scheme which is essentially a metric tree. The quantity 2^ρ bounds from above the number of leaf nodes (degree) of such a tree.

Concluding remarks and problems. • A thorough survey of intrinsic dimensionality in the context of similarity search is [2], though even this authoritative source is not comprehensive.

• It is safe to predict that the mainstream research direction in the field will be a discovery of further “dimensionality parameters of positive type”, such as Assouad dimension. More and more real datasets that currently appear high-dimensional will turn out to have low, tractable intrinsic dimensions.

- (Ciaccia, Pestov). Gromov’s reconstruction theorem ([3], p. 120) says that a metric space with measure can be recovered from distribution laws of all n -point metric subspaces, for all n . (Compare with the CNBYM dimension, determined by the distribution law of 2-point subspaces.) How much can one infer about indexability from the distribution law of 3-point subspaces?
- Notice that fractal-type dimensions (cf. [8]) fit into a general paradigm of intrinsic dimensionality outlined by us at the beginning: they reflect the rate of growth of balls/boxes. So does a version of intrinsic dimension proposed in [6], the *disorder dimension*, whose aim is to quantify the situation where $\mathbb{E}(\varepsilon_{NN})$ is close in value to $\mathbb{E}(d)$.
- We assumed above that queries follow the same underlying distribution as datapoints. Now suppose we have two underlying measures, μ modelling data and ν modelling query centers. Can one build a theory of concentration and intrinsic dimension for such metric bi-measure spaces?

References

- [1] E. Chávez, G. Navarro, R. Baeza-Yates, and J. L. Marroquín, *Searching in metric spaces*, ACM Computing Surveys **33** (2001), 273–321.
- [2] K.L. Clarkson, *Nearest-neighbor searching and metric space dimensions*, In: Nearest-Neighbor Methods for Learning and Vision: Theory and Practice, MIT Press, 2006, pp. 15–59.
- [3] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Progress in Mathematics **152**. Birkhauser Verlag, 1999.
- [4] P. Indyk, *Nearest neighbours in high-dimensional spaces*, In: J.E. Goodman, J. O’Rourke, Eds., Handbook of Discrete and Computational Geometry, Chapman and Hall/CRC, Boca Raton–London–New York–Washington, D.C., 2004, pp. 877–892.
- [5] R. Krauthgamer and J.R. Lee, *Navigating nets: simple algorithms for proximity search*, Proc. 15th ACM-SIAM symp. on Discrete algorithms (SODA’04), New Orleans, 2004, pp. 798–807.
- [6] Y. Lifshits, *Combinatorial framework for similarity search*, Proc. 2nd Int. Workshop on Similarity Search and Applications (SISAP 2009), Prague, Czech Republic, 2009, pp. 39-46.
- [7] V. Pestov, *Intrinsic dimension of a dataset: what properties does one expect?* - In: Proc. 20th Int. Joint Conf. on Neural Networks, Orlando, FL, 2007, pp. 1775–1780; *An axiomatic approach to intrinsic dimension of a dataset.* - Neural Networks **21**, 2-3 (2008), 204-213.
- [8] C. Traina, Jr., A.J.M. Traina, and C. Faloutsos, *Distance exponent: A new concept for selectivity estimation in metric trees*, Technical Report CMU-CS-99-110, Computer Science Department, Carnegie Mellon University, 1999.
- [9] I. Volnyansky and V. Pestov, *Curse of dimensionality in pivot-based indexes.* - Proc. 2nd Int. Workshop on Similarity Search and Applications (SISAP 2009), Prague, Czech Republic, 2009, pp. 39-46.