







Relative Intrinsic Dimensionality Is Intrinsic to Learning

Oliver J. Sutton¹(✉) , Qinghua Zhou¹ , Alexander N. Gorban² ,
and Ivan Y. Tyukin¹ 

¹ King's College London, London WC2R 2LS, UK
oliver.sutton@kcl.ac.uk

² University of Leicester, Leicester LE1 7RH, UK

Abstract. High dimensional data can have a surprising property: pairs of data points may be easily separated from each other, or even from arbitrary subsets, with high probability using just simple linear classifiers. However, this is more of a rule of thumb than a reliable property as high dimensionality alone is neither necessary nor sufficient for successful learning. Here, we introduce a new notion of the *intrinsic dimension* of a data distribution, which precisely captures the separability properties of the data. For this intrinsic dimension, the rule of thumb above becomes a law: high intrinsic dimension guarantees highly separable data. We extend this notion to that of the *relative* intrinsic dimension of two data distributions, which we show provides both upper and lower bounds on the probability of successfully learning and generalising in a binary classification problem.

Keywords: Intrinsic dimensionality · Classification problems · High dimensional data

1 Introduction

A *blessing of dimensionality* often ascribed to data sampled from genuinely high dimensional probability distributions is that pairs (and even arbitrary compact subsets) of points may be easily separated from one another with high probability [2, 4–7, 9, 13]. Such a property is naturally highly appealing for Machine Learning and Artificial Intelligence, since it suggests that if sufficiently many attributes can be obtained for each data point, then classification is a significantly easier task.

However, although this provides a useful rule of thumb, it is far from a complete description of the behaviour which may be expected of high dimensional data, and a simple experiment shows that the precise relationship between data dimension and classification performance is more subtle (see also [8], Theorem 5 and Corollary 2). Suppose that data are sampled from two classes, each described by a uniform distribution in a unit ball in \mathbb{R}^d , and that the centres of these balls are at distance $\epsilon \geq 0$ from one another, as shown in Fig. 1. The classifier which offers the optimal (balanced) accuracy in this case is given by the hyperplane




Fig. 1. Two unit balls separated by distance epsilon, and the optimal classifier (dotted) separating the two.

which is normal to the vector connecting the two centres and positioned half way between them. In Fig. 2 we plot the accuracy of this classifier as a function of the distance separating the two centres for data sampled from various different ambient dimensions d . The insight behind the blessing of dimensionality described above is immediately clear: when the data is sampled in high dimensions, for values of ϵ greater than some threshold value $\epsilon_0(d)$ depending on the ambient dimension d , the accuracy of this simple linear classifier is virtually 100%. Yet, what this simplified viewpoint misses is that, for $\epsilon < \epsilon_0(d)$ the probability of correctly classifying a given point sharply drops to close to 50%, demonstrating that raw dimensionality alone is no panacea for data classification¹. On the other hand, data sampled even in 1 dimension may be accurately classified when the centre separation ϵ is sufficiently large: for $\epsilon \geq 2$ (when the two unit balls are disjoint), the two data sets are fully separable in any dimension.

What this simple thought experiment demonstrates is a fact which is not taken into account by previous work, such as [12]:

Determining whether data distributions are separable from each other must depend on a relative property of the two, and even genuine high dimensionality² alone is neither a necessary nor sufficient condition for data separability.

To lay the foundations of our approach, we propose the new concept of the *intrinsic dimension* of a data distribution, based directly on the separability properties of sampled data points.

Definition 1 (Intrinsic dimension). *We say that data sampled from a distribution \mathcal{D} on \mathbb{R}^d has intrinsic dimension $n(\mathcal{D}) \in \mathbb{R}$ with respect to a centre $c \in \mathbb{R}^d$ if*

$$P(x, y \sim \mathcal{D} : (x - y, y - c) \geq 0) = \frac{1}{2^{n(\mathcal{D})+1}}. \tag{1}$$

This definition is designed in such a way that the rule of thumb in the blessing of dimensionality described above becomes a *law of high intrinsic dimension*:

¹ Moreover, standard dimensionality reduction techniques, such as Principle Components Analysis, would not have any effect here since the data are uniformly sampled from d -dimensional balls.

² In the sense that dimensionality reduction techniques cannot be applied to find an equivalent lower dimensional representation of the data.




Fig. 2. Accuracy of the best linear classifier separating data uniformly sampled from two balls with unit radius and centres in \mathbb{R}^n separated by distance ϵ for different dimensions n .

points sampled from a distribution with high intrinsic dimension are highly separable. The definition is calibrated so that the uniform distribution $\mathcal{U}(\mathbb{B}_d)$ on a d -dimensional unit ball \mathbb{B}_d satisfies $n(\mathcal{U}(\mathbb{B}_d)) = d$ (see Theorem 1), although alternative normalisations are possible, and by symmetry $n(\mathcal{D}) \geq 0$ for all distributions \mathcal{D} . For $c = 0$, the expression $(x - y, y - c) \geq 0$ in the left-hand side of (1) is simply a statement that x and y are Fisher-separable [8].

Based on the same principle, we further introduce the concept of the *relative intrinsic dimension* of two data distributions, which directly describes the ease of separating data distributions.

Definition 2 (Relative intrinsic dimension). We say that data sampled from a distribution \mathcal{D} on \mathbb{R}^d has relative intrinsic dimension $n(\mathcal{D}, \mathcal{D}') \in \mathbb{R}$ to data sampled from a distribution \mathcal{D}' on \mathbb{R}^d , with respect to a centre $c \in \mathbb{R}^d$, if

$$P(x \sim \mathcal{D}', y \sim \mathcal{D} : (x - y, y - c) \geq 0) = \frac{1}{2^{n(\mathcal{D}, \mathcal{D}') + 1}}. \tag{2}$$

The relative intrinsic dimension is not symmetric, and satisfies $n(\mathcal{D}, \mathcal{D}') \geq -1$, with negative values indicating that \mathcal{D} has lower intrinsic dimension than \mathcal{D}' , and data distributions with a low relative intrinsic dimension may be separated from distributions with a high relative intrinsic dimension.

To illustrate this, consider our previous experiment as an example and let $X = \mathcal{U}(B_1)$ and $Y = \mathcal{U}(B_2)$, where $B_1 = \mathbb{B}_d(1, c_1) \subset \mathbb{R}^d$ and $B_2 = \mathbb{B}_d(1, c_2) \subset \mathbb{R}^d$ are the unit balls centered at c_1 and c_2 respectively, and pick the centre $c = c_1$. When $\epsilon = \|c_1 - c_2\| \geq 2$ (the case when the data distributions are completely separable), we have $n(Y, X) = \infty$. This implies that points y sampled from Y can be separated from points x sampled from X with certainty. The relative intrinsic dimension $n(X, Y)$ is an increasing function of the dimension of the ambient space in which the data is sampled with $n(X, Y) = 0$ in 1 dimension, implying that it becomes easier to separate points in X from points in Y as the dimension increases. These values of the relative intrinsic dimensions suggest that points from Y can easily be separated from points in

X by hyperplanes normal to $y - c_1$, while hyperplanes normal to $x - c_1$ do not separate X from Y .

Although the asymmetry may be slightly surprising at first, it is simply reflecting the asymmetric choice of centre $c = c_1$, which is located at the heart of the X distribution. The relative intrinsic dimensions described above would be reversed for $c = c_2$ and would be equal for $c = \frac{1}{2}(c_1 + c_2)$. A justification for this definition of relative intrinsic dimension is given by Theorem 2, where it is shown (in a slightly generalised setting) that these concepts of intrinsic dimension provide upper and lower bounds on classifier accuracy, indicating that it is indeed necessary and sufficient for learning.

There is a rich history of alternative characterisations of the dimension of a data set, with each contribution typically aimed to solve a particular problem. For example, conventional Principle Components Analysis aims to detect the number of independent attributes which are actually required to represent the data, leading to compressed representations of the same data. However, as discussed above, the representational dimension of a data set does not necessarily give an indication of how easy it is to learn from. Several other notions of dimensionality are captured in the `scikit-dimension` library [3]. Perhaps the most similar notion of dimension to that which we propose here is the Fisher Separability Dimension [1], which is also based on the separability properties of data yet first requires a whitening step to normalise the data covariance to an identity matrix. This whitening step has both advantages and disadvantages: although it brings invariance to the choice and scaling of the basis, it disrupts the intrinsic geometry of the data. The Fisher Separability Dimension also does not address the important question of the *relative* dimension of data distributions and samples, which we argue is a concept fundamental to learning.

Our approach may appear reminiscent of Kernel Embeddings, through which nonlinear kernels are used to embed families of data distributions into a Hilbert space structure [11]. Although Kernel Embeddings and our work are motivated by very different classes problems, the common fundamental focus is on understanding the properties of a data distribution through the evaluation of (nonlinear) functionals of the distribution. Here we demonstrate how a single, targeted, property appears to encode important information about the separability properties of data.

An interesting question which arises from this work is how well the (relative) intrinsic dimension can be estimated from data samples directly. If it can be, then this could provide a new tool for selecting appropriate feature mappings for data and shine a new light on the training of neural networks. We briefly investigate this in Sect. 4, where we show that high order polynomial feature maps can actually be detrimental to the separability of data.

2 Separability of Uniformly Distributed Data

We investigate the separability properties of data sampled from a uniform distribution in the unit ball in various dimensions. This provides the basis for our definition of intrinsic dimension.

To simplify the presentation of our results, we introduce the following geometric quantities related to spheres in high dimensions. The volume of a ball with radius r in d dimensions is denoted by

$$V_d^{\text{ball}}(r) = \frac{\pi^{d/2} r^d}{\Gamma(\frac{d}{2} + 1)},$$

and the surface area of the same ball is denoted by

$$S_d^{\text{ball}}(r) = \frac{d\pi^{d/2} r^{d-1}}{\Gamma(\frac{d}{2} + 1)}.$$

Similarly, the volume of the spherical cap with height h of the same sphere (i.e. the set of points $\{x \in \mathbb{R}^d : \|x\| \leq r \text{ and } x_0 \geq r - h\}$) is given by $V_d^{\text{cap}}(r, h) = V_d^{\text{ball}}(r)W_d^{\text{cap}}(r, h)$, where

$$W_d^{\text{cap}}(r, h) = \begin{cases} 0 & \text{for } h \leq 0, \\ \frac{1}{2}I_{(2rh-h^2)/r^2}(\frac{d+1}{2}, \frac{1}{2}) & \text{for } 0 < h \leq r, \\ 1 - W_d^{\text{cap}}(r, 2r - h) & \text{for } r < h \leq 2r, \\ 1 & \text{for } 2r < h, \end{cases}$$

represents the fraction of the volume of the unit ball contained in the spherical cap. The function $I_x(a, b) = B(a, b)^{-1} \int_0^x t^{a-1}(1-t)^{b-1} dt$ denotes the regularised incomplete beta function, where $B(a, b) = B(1; a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$ is the standard beta function.




Fig. 3. The behaviour of $f_\theta(d)$, formally extended to non-integer values of d , for various values of θ . The function is only invertible for $-1 \leq \theta \leq 0$, and we note the asymptote of $\frac{1}{2}$ as $d \rightarrow 0$ when $\theta = 0$ and as $d \rightarrow \infty$ when $\theta = -1$

Theorem 1 (Separability of uniformly sampled points). *Let $\theta \in \mathbb{R}$, let d be a positive integer and suppose that $x, y \sim \mathcal{U}(\mathbb{B}_d(1, c))$, define*

$$R_\theta(t) = \max \left\{ \frac{t^2}{4} - \theta, 0 \right\}^{\frac{1}{2}}, \quad a_\theta(t) = \frac{1 - R_\theta^2(t)}{t} - \frac{t}{4}, \tag{3}$$

and

$$b_\theta(t) = 1 - a_\theta(t) - \frac{t}{2}, \tag{4}$$

and let

$$f_\theta(d) = \int_0^1 dt^{d-1} (W_d^{\text{cap}}(1, b_\theta(t)) + R_\theta^d(t)W_d^{\text{cap}}(R_\theta(t), R_\theta(t) + a_\theta(t))) dt. \tag{5}$$

Then

$$P(x, y : (y - x, x - c) \geq \theta) = f_\theta(d), \tag{6}$$

and, in particular,

$$P(x, y : (y - x, x - c) \geq 0) = \frac{1}{2^{d+1}}. \tag{7}$$

Furthermore, f_θ may be simplified in the following cases as

$$f_\theta(d) = \begin{cases} 1 & \text{for } \theta \leq -2, \\ \frac{1}{2^{d+1}} & \text{for } \theta = 0, \\ \int_{2\theta^{1/2}}^1 dt^{d-1} \left(\frac{t^2}{4} - \theta\right)^{d/2} dt & \text{for } 0 < \theta < \frac{1}{4}, \\ 0 & \text{for } \frac{1}{4} \leq \theta. \end{cases} \tag{8}$$

and $f_\theta(d) \geq \frac{1}{2}$ for $\theta \leq -1$.




Fig. 4. The shaded area is the volume computed in the proof of Theorem 1. The two different shading colours indicate the two spherical caps used in the proof.

Proof. Without loss of generality, we suppose that $c = 0$, and consider points $x, y \sim \mathcal{U}(\mathbb{B}_d)$. Rearranging terms, we observe that

$$(y - x, x) = \frac{1}{4} \|y\|^2 - \|x - \frac{y}{2}\|^2,$$

and therefore, for fixed y , the set of x satisfying $(y-x, x-c) \geq \theta$ may be similarly described as those points x contained within the ball

$$\|x - \frac{y}{2}\|^2 \leq R(\|y\|) = \max\left\{\frac{1}{4}\|y\|^2 - \theta, 0\right\}.$$

Combining this with the condition that $x \sim \mathbb{B}_d(1, 0)$, we find that x belongs to the intersection of the balls

$$\{x \in \mathbb{R}^d : \|x\| \leq 1\} \cap \left\{x \in \mathbb{R}^d : \|x - \frac{y}{2}\|^2 \leq R_\theta(\|y\|)\right\}. \tag{9}$$

This may be expressed as the union of two spherical caps, as depicted in Fig. 4. Comparing the triangles O, p, q and $\frac{y}{2}, p, q$ shows that the lengths a and b in the Figure are exactly those defined in (3) with $t = \|y\|$. Since y only appears through its norm, we deduce that

$$\begin{aligned} P(x : (y-x, x) \geq \theta \mid \|y\|) &= P(x : (y-x, x) \geq \theta \mid y) \\ &= \frac{V_d^{\text{cap}}(R_\theta(\|y\|), R_\theta(\|y\|) + a_\theta(\|y\|)) + V_d^{\text{cap}}(1, b_\theta(\|y\|))}{V_d^{\text{ball}}(1)}, \end{aligned}$$

The result (6) follows by applying the law of total probability, which implies

$$P(x, y : (y-x, x) \geq \theta) = \int_0^1 P(x : (y-x, x) \geq \theta \mid \|y\| = t) p_{\|y\|}(t) dt,$$

where $p_{\|y\|}(t) = \frac{S_d^{\text{ball}}(t)}{V_d^{\text{ball}}(1)}$ is the density associated with $\|y\|$ for $y \sim \mathcal{U}(\mathbb{B}_d)$.

When $\theta \geq 0$, the ball centered at $\frac{y}{2}$ is entirely contained within \mathbb{B}_d , and so

$$\begin{aligned} P(x, y : (y-x, x) \geq \theta) &= \int_0^1 \frac{S_d^{\text{ball}}(t) V_d^{\text{ball}}(R_\theta(t))}{(V_d^{\text{ball}}(1))^2} dt \\ &= \int_0^1 dt^{d-1} \max\left\{\frac{t^2}{4} - \theta, 0\right\}^{d/2} dt. \end{aligned}$$

Since the integrand is zero for $t \leq 2\theta^{1/2}$, for $\theta \in (0, \frac{1}{4})$ we have

$$P(x, y : (y-x, x) \geq \theta) = \int_{2\theta^{1/2}}^1 dt^{d-1} \left(\frac{t^2}{4} - \theta\right)^{d/2} dt.$$

Moreover, $P(x, y : (y-x, x) \geq \theta) = 0$ for $\theta \geq \frac{1}{4}$, and in the simplest case of $\theta = 0$

$$P(x, y : (y-x, x) \geq 0) = \frac{d}{2^d} \int_0^1 t^{2d-1} dt = \frac{1}{2^{d+1}}.$$

On the other hand, for $\theta \leq -2$ we have $\sqrt{R_\theta(t)} \geq 1 + \frac{1}{2}t$ for all t , implying that the intersection (9) is the entirety of \mathbb{B}_d , and hence

$$P(x, y : (y-x, x) \geq \theta) = 1.$$

□

The behaviour of $f_\theta(d)$ is illustrated in Fig. 3 for various values of the separation threshold θ . Heuristically, we observe the following limiting behaviour:

$$\lim_{d \rightarrow \infty} f_\theta(d) = \begin{cases} 1 & \text{for } \theta < -1, \\ \frac{1}{2} & \text{for } \theta = -1, \\ 0 & \text{for } \theta > -1, \end{cases}$$

which may be explained by the fact that when $\theta = -1$, the surfaces of the ball \mathbb{B}_d and the ball centered at $\frac{y}{2}$ meet exactly at an equator of \mathbb{B}_d . The phenomenon of waist concentration (see [10], for example) implies that in high dimensions the volume of \mathbb{B}_d is concentrated around its surface and around this equator, implying that this is the threshold value of θ at which the intersection of the two balls contains slightly more than half the volume of \mathbb{B}_d .

What these results suggest is that for any value of $\theta \in [-1, 0]$, the function $f_\theta(d)$ is an invertible function of d , and hence could be used as the basis of a definition of intrinsic dimension. In Definition 1 we use the behaviour at $\theta = 0$ to define our indicative notion of intrinsic dimension simply because it obviates the need to couple the scaling of the support of the distribution and the scaling of θ .

3 Few Shot Learning Is Dependent on Separability

We now consider the scenario of standard binary data classification, and show that the probability of successfully learning to classify data is intrinsically linked to the notion of relative intrinsic dimension. We focus on the case of learning from small data sets, as in this case the link is particularly clear to demonstrate.

Mathematically, we suppose that X and Y are (unknown) probability distributions on an d -dimensional vector space \mathbb{R}^d , and we have a sample $\{y_i\}_{i=1}^k$ of k training points sampled from Y and a sample $\{x_i\}_{i=1}^m$ of m training points sampled from X .

Since the problem setup is symmetric in the roles of X and Y , we only analyse the influence of training data sampled from Y . The role of the data sampled from X (alongside any possible prior knowledge of the data distributions) is incorporated through an arbitrary but fixed point $c \in \mathbb{R}^d$ in the data space.

We consider the following linear classifier to assign the label ℓ_X to data sampled from X and the label ℓ_Y to data sampled from Y :

$$F_\theta(z) = \begin{cases} \ell_Y & \text{if } L(z) \geq \theta, \\ \ell_X & \text{otherwise,} \end{cases} \quad (10)$$

where $L(z) = \frac{1}{k} \sum_{i=1}^k (z - y_i, y_i - c)$. In practice, the value of the threshold θ to be used in the classifier may be determined from the training data $\{y_i\}_{i=1}^k$ and $\{x_i\}_{i=1}^m$, although here we consider it to be a free parameter of the classifier.

Remark 1 (Comparison with similar classifiers). The classifier (10) may be equivalently be expressed in the form of the common Fisher discriminant with a slightly different threshold, viz.

$$F_\theta(z) = \begin{cases} \ell_Y & \text{if } (z - \mu, \mu - c) \geq \theta + \Theta, \\ \ell_X & \text{otherwise,} \end{cases}$$

where $\mu = \frac{1}{k} \sum_{i=1}^k y_i$ and $\Theta = \frac{1}{k} \sum_{i=1}^k \|y_i\|^2 - \|\mu\|^2$. Since the offset Θ to the threshold θ depends only on the same training data as θ , it is clear that the classifier we study is simply a Fisher discriminant. However, we choose to write the classifier in the form (10) because it simplifies some of the forthcoming analysis.

This classifier will successfully learn to classify the training data when both

$$P(F_\theta(y) = \ell_Y) = P(L(y) \geq \theta)$$

is large (where the probability is taken with respect to the evaluation point $y \sim Y$ and the training data $\{y_i \sim Y\}_{i=1}^k$), and

$$P(F_\theta(x) = \ell_X) = P(L(x) < \theta)$$

is also large (where the probability is taken with respect to the evaluation point $x \sim X$ and the training data $\{y_i \sim Y\}_{i=1}^k$). We now show that both of these probabilities can be bounded from above and below by the probability of being able to separate pairs of data points by margin θ . Corollary 1 to this theorem then shows how this simply reduces to upper and lower bounds dependent on the (relative) intrinsic dimension of Y and X when $\theta = 0$.

Theorem 2 (Pairwise separability and learning). Let $\theta \in \mathbb{R}$ and define

$$p_\theta(Y, X) = P(x \sim X, y \sim Y : (x - y, y - c) \geq \theta),$$

and let $p_\theta(Y) = p_\theta(Y, Y)$. Then, the probability (with respect to the training sample $\{y_i \sim Y\}_{i=1}^k$ and the evaluation point $y \sim Y$) of successfully learning the class Y is bounded by

$$p_\theta^k(Y) \leq P(F_\theta(y) = \ell_Y) \leq 1 - (1 - p_\theta(Y))^k, \tag{11}$$

and the probability (with respect to the training sample $\{y_i \sim Y\}_{i=1}^k$ and the evaluation point $x \sim X$) of successfully learning the class X is bounded by

$$(1 - p_\theta(Y, X))^k \leq P(F_\theta(x) = \ell_X) \leq 1 - p_\theta^k(Y, X). \tag{12}$$

Proof. Let E be the event that $F_\theta(y) = \ell_Y$ for $y \sim Y$. By definition, this occurs when y and $\{y_i\}_{i=1}^k$ are such that $\sum_{i=1}^k (y - y_i, y_i - c) \geq k\theta$. For each $1 \leq i \leq k$, let A_i denote the event that $(y - y_i, y_i - c) \geq \theta$. Then, $\bigwedge_{i=1}^k A_i \Rightarrow E$ and so $P(E) \geq P(\bigwedge_{i=1}^k A_i)$. We may further expand this using the law of total probability as

$$P\left(\bigwedge_{i=1}^k A_i\right) = \int_{\mathbb{R}^d} P\left(\bigwedge_{i=1}^k (y - y_i, y_i - c) \geq \theta \mid y\right) p(y) dy. \tag{13}$$

Since the $\{y_i\}_{i=1}^k$ are independently sampled and identically distributed, it follows that the conditional probability satisfies

$$P\left(\{y_i \sim Y\}_{i=1}^k : \bigwedge_{i=1}^k (y - y_i, y_i - c) \geq \theta \mid y\right) = P(y' \sim Y : (y - y', y' - c) \geq \theta \mid y)^k.$$

Substituting this into (13) shows that $P(\bigwedge_{i=1}^k A_i) = \mathbb{E}_Y[(P(y' \sim Y : (y - y', y' - c) \geq \theta \mid y))^k]$, where the expectation is taken with respect to y . For a random variable X and a convex function g , Jensen's inequality asserts that $\mathbb{E}[g(X)] \geq g(\mathbb{E}[X])$. Applying this here (since the function $g(x) = x^k$ is convex for $k \geq 1$), we find that

$$\begin{aligned} P\left(\bigwedge_{i=1}^k A_i\right) &\geq (\mathbb{E}_Y[P(y' : (y - y', y' - c) \geq \theta \mid y)])^k \\ &= (P(y, y' : (y - y', y' - c) \geq \theta))^k. \end{aligned}$$

Consequently, we deduce the lower bound of (11). The upper bound follows by arguing similarly and using the fact that $\bigwedge_{i=1}^k \text{not } A_i \Rightarrow \text{not } E$, from which it follows that $P(E) \leq 1 - P(\bigwedge_{i=1}^k \text{not } A_i)$. An analogous argument shows the result (12). \square

An immediate consequence of this theorem is that when $\theta = 0$, the probability of successfully learning can be bounded from both above and below using the (relative) intrinsic dimension of the data distributions.

Corollary 1 (Intrinsic dimension and learning). *The probability (with respect to the training sample $\{y_i \sim Y\}_{i=1}^k$ and the evaluation point $y \sim Y$) of successfully learning the class Y is bounded by*

$$\frac{1}{2^{k(n(Y)+1)}} \leq P(F_0(y) = \ell_Y) \leq 1 - \left(1 - \frac{1}{2^{n(Y)+1}}\right)^k, \tag{14}$$

and the probability (with respect to the training sample $\{y_i \sim Y\}_{i=1}^k$ and the evaluation point $x \sim X$) of successfully learning the class X is bounded by

$$1 - \left(1 - \frac{1}{2^{n(Y,X)+1}}\right)^k \leq P(F_0(x) = \ell_X) \leq \frac{1}{2^{k(n(Y,X)+1)}}$$

We note that the best lower bound which can be shown by (14) is $\frac{1}{2}$, due to the fact that the classifier with $\theta = 0$ will pass through the centre of the Y distribution. Despite this, Corollary 1 shows that the intrinsic dimension of Y is sufficient to know whether the probability of correctly learning the class Y is less than $\frac{1}{2}$. Arguing symmetrically, a more refined analysis taking more account of the training set $\{x_i\}_{i=1}^m$ could instead show a version of the bound (14) which depends on the relative intrinsic dimension $n(X, Y)$.

These bounds are tuned to the case when the size k of the training set sampled from Y is small, and the upper and lower bounds separate from each other as

k grows, and alternative arguments would be required to get sharp bounds in the case of large k . However, even for large values of k , if the (relative) intrinsic dimension of the data distributions is sufficiently large or small, the bounds above will provide tight guarantees on the success of learning.

4 Learning with Polynomial Kernels

As an application of our proposed notion of intrinsic dimension, we use it to find the optimal polynomial kernel for a classification problem — i.e. the degree of the polynomial feature map in which two data sets become easiest to separate.

For fixed bias $b > 1$ and polynomial degree $k \geq 0$, let the polynomial kernel $\kappa : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be given by

$$\kappa(x, y) = (b^2 + x \cdot y)^k. \tag{15}$$

There exists a polynomial feature map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^N$, where $N = \binom{d+k}{k}$, such that $\kappa(x, y) = (\phi(x), \phi(y))$ (see [12], for example, for details).

Consider

$$P(x, y, \sim \mathcal{U}(\mathbb{B}_d) : (\phi(x) - \phi(y), \phi(y) - c) \geq \theta),$$

where $c = \frac{1}{V_d^{\text{ball}}(1)} \int_{\mathbb{B}_d} \phi(z) dz$ is the empirical mean of the data in feature space. Then, expanding the inner product,

$$\begin{aligned} (\phi(x) - \phi(y), \phi(y) - c) &= k(x, y) - k(y, y) + \int_{\mathbb{B}_d} \frac{k(y, z) - k(x, z)}{V_d^{\text{ball}}(1)} dz \\ &= (b^2 + x \cdot y)^k - (b^2 + \|y\|^2)^k + \int_{\mathbb{B}_d} \frac{(b^2 + y \cdot z)^k - (b^2 + x \cdot z)^k}{V_d^{\text{ball}}(1)} dz. \end{aligned}$$

Exploiting the spherical symmetry of $\mathcal{U}(\mathbb{B}_d)$, we have

$$\frac{1}{V_d^{\text{ball}}(1)} \int_{\mathbb{B}_d} (b^2 + x \cdot z)^k dz = \int_{-1}^1 \frac{V_{d-1}^{\text{ball}}((1-t^2)^{1/2})}{V_d^{\text{ball}}(1)} (b^2 + t\|x\|)^k dt = q(\|x\|),$$

for $b \geq 1$, where $q : [0, 1] \rightarrow \mathbb{R}$ is given by $q(\|x\|) := b^{2k} {}_2F_1\left(\frac{1-k}{2}, -\frac{k}{2}; \frac{d}{2}+1; \frac{\|x\|^2}{b^4}\right)$, with ${}_2F_1$ denoting the hypergeometric function. Therefore $(\phi(x) - \phi(y), \phi(y) - c) \geq \theta$ if and only if

$$\cos(\beta(x, y)) \geq Q(\|x\|, \|y\|)$$

where $\beta(x, y) = \arccos\left(\frac{(x,y)}{\|x\|\|y\|}\right)$ denotes the angle between x and y , and

$$Q(s, t) := (st)^{-1} \left((\theta + (b^2 + t^2)^k + q(s) - q(t))^{1/k} - b^2 \right).$$

Geometric arguments show that for any $\alpha \in [-1, 1]$,

$$P(x, y \sim \mathcal{U}(\mathbb{B}_d) : \cos(\beta(x, y)) \geq \alpha \mid \|x\|, \|y\|) = T_d^{\text{cap}}(\alpha)$$

where $T_d^{\text{cap}}(\alpha)$ denotes the proportion of the surface area of a unit sphere which falls within a spherical cap with opening angle $\arccos(\alpha)$, given for $d > 1$ by

$$T_d^{\text{cap}}(\alpha) = \begin{cases} 0, & \alpha > 1, \\ \frac{1}{2} I_{(\sin(\arccos(\alpha)))^2} \left(\frac{d-1}{2}, \frac{1}{2} \right), & \alpha \in [0, 1], \\ 1 - T_d^{\text{cap}}(-\alpha), & \alpha \in (-1, 0), \\ 1, & \alpha \leq -1, \end{cases}$$

where $I_x(a, b)$ is the regularised incomplete beta function, and for $d = 1$ by

$$T_1^{\text{cap}}(\alpha) = \begin{cases} 0 & \text{for } \alpha > 1; \\ \frac{1}{2} & \text{for } \alpha \in (-1, 1]; \\ 1 & \text{for } \alpha \leq -1 \end{cases}$$

Let E be the event that $x, y \sim \mathcal{U}(\mathbb{B}_d)$ are such that $\cos(\beta) \geq Q(\|x\|, \|y\|)$. Then, by the law of total probability,

$$P(E) = \int_0^1 \int_0^1 P(E \mid \|x\| = s, \|y\| = t) \hat{p}(s) \hat{p}(t) ds dt,$$

where $\hat{p}(t) = \frac{S_d^{\text{ball}}(t)}{V_d^{\text{ball}}(1)} = dt^{d-1}$ denotes the density associated with $\|z\|$ for $z \sim \mathcal{U}(\mathbb{B}_d)$.

The arguments above therefore prove the following theorem, from which Theorem 1 arises as a simplified special case when $k = 1$

Theorem 3 (Separability in polynomial feature space). *Let $k > 0$, let d be a fixed positive integer, and let ϕ denote the feature map associated with the polynomial kernel (15) with degree k in dimension d . Then, for $\theta \in \mathbb{R}$,*

$$\begin{aligned} P(x, y \sim \mathcal{U}(\mathbb{B}_d) : (\phi(x) - \phi(y), \phi(y) - c) \geq \theta) \\ = d^2 \int_0^1 \int_0^1 T_d^{\text{cap}}(Q(s, t)) s^{d-1} t^{d-1} ds dt. \end{aligned}$$

Figure 5 shows how the intrinsic dimension of the unit ball in various dimensions is affected by applying a polynomial feature mapping. Since the degree k polynomial feature map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^N$, where $N = \binom{d+k}{k}$, increases the apparent dimension of the space as k increases, the rule of thumb encapsulated by the blessing of dimensionality would lead us to expect that high order polynomial kernels should make the data more separable. However, this is not what we observe. Instead, the intrinsic dimension reveals that there is an ‘optimal’ polynomial degree, for which the data is most separable, and increasing the polynomial degree further beyond the point can actually have the detrimental effect of making the data less separable.




Fig. 5. The intrinsic dimension of the image of $\mathcal{U}(\mathbb{B}_d)$ under a polynomial feature map, for different polynomial degrees and data space dimensions d .

5 Conclusion

We have introduced a new notion of the intrinsic dimension of a data distribution, based on the pairwise separability properties of data points sampled from this distribution. Alongside this, we have also introduced a notion of the relative intrinsic dimension of a data distribution relative to another distribution. Theorem 2 shows how these notions of intrinsic dimension occupy a fundamental position in the theory of learning, as they directly provide upper and lower bounds on the probability of successfully learning in a generalisable fashion.

Many open questions remain, however, such as how to accurately determine the intrinsic dimension of a data distribution using just sampled data from that distribution, and how best to utilise these insights to improve neural network learning. This work also opens the door to generalising the concept beyond just simple linear functionals of the data distribution to notions of intrinsic dimensionality based around other more interesting models. The idea also generalises beyond examining individual points sampled from distributions, to studying the collective behaviour of groups, or ‘granules’ of sampled data.

Acknowledgements. The authors are grateful for financial support by the UKRI and EPSRC (UKRI Turing AI Fellowship ARaISE EP/V025295/1). I.Y.T. is also grateful for support from the UKRI Trustworthy Autonomous Systems Node in Verifiability EP/V026801/1.

References

1. Albergante, L., Bac, J., Zinovyev, A.: Estimating the effective dimension of large biological datasets using fisher separability analysis. In: 2019 International Joint Conference on Neural Networks (IJCNN) (2019)
2. Anderson, J., Belkin, M., Goyal, N., Rademacher, L., Voss, J.: The more, the merrier: the blessing of dimensionality for learning large Gaussian mixtures. In: Conference on Learning Theory, pp. 1135–1164. PMLR (2014)

3. Bac, J., Mirkes, E.M., Gorban, A.N., Tyukin, I., Zinovyev, A.: Scikit-Dimension: a python package for intrinsic dimension estimation. *Entropy* **23**(10), 1368 (2021). <https://doi.org/10.3390/e23101368>
4. Donoho, D., Tanner, J.: Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philos. Trans. Royal Soc. A Math. Phys. Eng. Sci.* **367**(1906), 4273–4293 (2009)
5. Gorban, A.N., Tyukin, I.Y.: Stochastic separation theorems. *Neural Netw.* **94**, 255–259 (2017)
6. Gorban, A.N., Tyukin, I.Y.: Blessing of dimensionality: mathematical foundations of the statistical physics of data. *Philos. Trans. Royal Soc. A Math. Phys. Eng. Sci.* **376**(2118), 20170237 (2018)
7. Gorban, A.N., Tyukin, I.Y., Romanenko, I.: The blessing of dimensionality: separation theorems in the thermodynamic limit. *IFAC-PapersOnLine* **49**(24), 64–69 (2016)
8. Gorban, A., Golubkov, A., Grechuk, B., Mirkes, E., Tyukin, I.: Correction of AI systems by linear discriminants: probabilistic foundations. *Inf. Sci.* **466**, 303–322 (2018). <https://doi.org/10.1016/j.ins.2018.07.040>
9. Kainen, P.C., Kůrková, V.: Quasiorthogonal dimension. In: Kosheleva, O., Shary, S.P., Xiang, G., Zapatin, R. (eds.) *Beyond Traditional Probabilistic Data Processing Techniques: Interval, Fuzzy etc. Methods and Their Applications*. SCI, vol. 835, pp. 615–629. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-31041-7_35
10. Ledoux, M.: *The Concentration of Measure Phenomenon*, vol. 89, American Mathematical Society (2001)
11. Smola, A., Gretton, A., Song, L., Schölkopf, B.: A hilbert space embedding for distributions. In: Hutter, M., Servedio, R.A., Takimoto, E. (eds.) *Algorithmic Learning Theory*, pp. 13–31. Springer, Berlin Heidelberg, Berlin, Heidelberg (2007)
12. Sutton, O.J., Gorban, A.N., Tyukin, I.Y.: Towards a mathematical understanding of learning from few examples with nonlinear feature maps. <https://doi.org/10.48550/ARXIV.2211.03607><https://arxiv.org/abs/2211.03607> (2022)
13. Tyukin, I.Y., Gorban, A.N., Grechuk, B., Green, S.: Kernel stochastic separation theorems and separability characterizations of kernel classifiers. In: 2019 International Joint Conference on Neural Networks (IJCNN), pp. 1–6 (2019). <https://doi.org/10.1109/IJCNN.2019.8852278>