Intrinsic Dimensionality and the Compromisability of Data Indexing and Analysis

Dominique Barbe

National Institute of Informatics, Tokyo, Japan
Dept. of Mathematics and Informatics
dominique.barbe@ens-rennes.fr

Abstract. Machine learning systems have been practically shown to misbehave in adversarial environments, including state-of-the-art neural networks. Adversarial examples, objects slightly altered in order to be misclassified with high confidence, while still being close to the original object, can be crafted. Theoretical explanations of this problem are emerging, involving for example the highly linear nature of current classifiers, the low density yet dense repartition of adversarial examples, or the high intrinsic dimensionality of the data. In this report, we study general neighbourhood manipulations, where vicinities of objects are modified by moving them. We show that any neighbourhood manipulation can be achieved with an infinitesimal amount of movement, provided that the database is large and the intrinsic dimensionality is too.

Keywords: Adversarial learning, K-nearest neighbour graph, Intrinsic dimensionality

This reports refers to the work realized at the National Institute of Informatics in Tokyo, under the supervision of Michael E. Houle, from the 06/06/2016 to the 12/08/2016.

1 Introduction

Machine learning techniques have proven their efficiency in many domains: face recognition, fraud detection, online recommendation, speech recognition, etc. All applications are not equally successful, yet these techniques are getting increasingly widespread.

Despite their efficiency, they are highly vulnerable in an adversarial environment. Machine learning models, including neural networks, can be led to misclassify objects with high confidence [8]. To achieve this, a specifically crafted small noise is added to the input. The nature of the noise is not fixed, and the adversarial examples can seem indistinguishable from the original input.

Two others intriguing properties make these attacks very dangerous. First, adversarial examples tend to generalize well across models [9,11], despite their differences in architectures and training sets. Thus, an attacker does not necessarily need an access to the details of the classifier; instead he could train its

own system and use it to craft adversarial examples. Second, these attacks are likely to happen in the physical world: a significant portion of printed adversarial images are still misclassified once photographed back and fed into a classifier [10].

Adversarial examples have to be taken into account. In contexts where there is an adversary, like spam detection, they will occur. In others contexts, they can improve the efficiency of classifiers if used in the training set [11]. There have been several attempts to theoretically explain the adversarial effect [11,13].

In this report, we expand the following work: [13]. In this paper, a first theoretical explanation of the adversarial effect is given. Using the notion of Local Intrinsic Dimensionality as a measure of the vulnerability of machine learning systems, it shows that the adversarial effect is unavoidable as the size of the training set and the intrinsic dimensionality both rise. More particularly: in an Euclidean space, it considers a reference point perturbed toward a target point, so as to have its as its closest neighbour. The amount of perturbation required tends to 0 when both the number of points and the instrincic dimensionality at the reference point's location tends toward infinity.

In this report, this result is extended to more cases. The conclusion can be summarized as such: any desired neighbourhood manipulation can be achieved with an amount of perturbation tending to 0 with a high number of points and a high intrinsic dimensionality.

This report is organized as follows: in Section 2, we present the notion of Local Intrinsic Dimensionality and a theorem that will support our later proofs. In Section 3 we give a proof of our theoretical results. In Section 4 we present the experimental work verifying our results, and exploring some details. Section 5 concludes this report with some possible extensions for this work.

2 Background

2.1 The adversarial effect

Several methods for crafting adversarial objects exist. For example, one of them involves solving the following optimization problem:

$$\min_{\mathbf{d}} \|\mathbf{d}\| + \alpha D_{KL}(\mathbf{p}^A || f(\mathbf{x}+\mathbf{d})), \text{ subject to: } \mathbf{l} \leq \mathbf{x}+\mathbf{d} \leq \mathbf{u} \tag{1}$$

$f : \mathbf{x} \to [p_1, \ldots, p_C]$ is the classifier, a function computing the probability of an object $\mathbf{x} \in \mathbb{R}^d$ to belong to each of the C predefined classes. $\mathbf{p}^A = [\mathbb{1}_{i=a}]$ is the targeted probability. $\mathbf{x}$ and $\mathbf{d}$ are respectively the reference object and the noise added to it. $D_{KL}(\cdot)$ is the Kullback-Leibler divergence, α is a trade-off between the level of noise and the closeness to the targeted output. $\mathbf{l}$ and $\mathbf{u}$ define the lower and upper bounds of the input domain respectively.

With classifiers trained using gradient descent, the above optimization problem can be solved easily, using either gradient descent or box-constrained L-BFGS [9]. An example of an adversarial attack using this strategy is shown in Figure 1.

Fig. 1: Using the optimization formulation in (1), adding adversarial noise (center) to an example of class 'Shark' (left) induces a deep neural network to assign the perturbed result (right) to the class 'Binocular'.

2.2 Local Intrinsic dimensionality (LID)

We introduce here a notion of local dimension, or intrinsic dimensionality (ID).

For the m-dimensional Euclidean space, doubling the radius of a sphere increases its volume by a factor 2^m. More generally, given two balls of different radius, the ratio of the volumes and the distances is linked by the dimension of the space m.

$$\frac{V_2}{V_1} = \left(\frac{r_2}{r_1}\right)^m \tag{2}$$

Thus, the dimension of the space can be recovered from the radius and the volume of two spheres:

$$m = \frac{\ln V_2 - \ln V_1}{\ln r_2 - \ln r_1}\,. \tag{3}$$

For finite data sets, the expansion dimension (ED) is computed by estimating the volume of the spheres by the number of points they enclose [3,2]. This notion of intrinsic dimension can be expanded to a statistical setting, rather than being a characteristic of a given set of points. The distribution of distances to a query point is modelled as a continuous random variable $\mathbf{X}$ [4,6]. The notion of volume then becomes the notion of probability measure. ID can then be modelled as a function of distance $\mathbf{X} = r$, by letting the radii of the two balls be $r_1 = r$ and $r_2 = (1+\epsilon)r$, and letting $\epsilon \to 0^+$.

Definition 1 ([6]). *Let $\mathbf{X}$ be a random distance variable. For any r such that $F_{\mathbf{X}}(r) > 0$, the* local intrinsic dimensionality *of $\mathbf{X}$ at r is given by*

$$\mathrm{ID}_{F_{\mathbf{X}}}(r) \triangleq \lim_{\epsilon \to 0^+} \frac{\ln F_{\mathbf{X}}((1+\epsilon)r) - \ln F_{\mathbf{X}}(r)}{\ln((1+\epsilon)r) - \ln r} = \frac{r \cdot F'_{\mathbf{X}}(r)}{F_{\mathbf{X}}(r)}, \tag{4}$$

wherever the limit exists. The second equality follows by applying l'Hôpital's rule to the limits provided that $F_{\mathbf{X}}$ is positive and differentiable over an open interval containing r.

Under this distributional interpretation, the original data set determines a sample of distances from a given point. The intrinsic dimensionality (here referred to simply as 'local ID', or 'LID') of this distance distribution $F_{\mathbf{X}}$ is estimated. The definition of $\mathrm{ID}_{F_{\mathbf{X}}}$ can be extended to the case where $r = 0$, by

taking the limit of $\mathrm{ID}_{F_\mathbf{X}}(r)$ as $r \to 0^+$, whenever this limit exists:

$$\mathrm{ID}_{F_\mathbf{X}}(0) \triangleq \lim_{r \to 0^+} \mathrm{ID}_{F_\mathbf{X}}(r)\,. \tag{5}$$

For an illustration of the intrinsic dimensionality of distance distributions, see Figure 2.

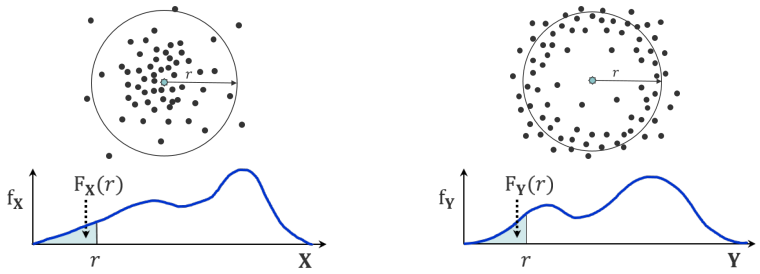

Fig. 2: The random distance variables $\mathbf{X}$ and $\mathbf{Y}$ have different LID values at distance r. Although the total probability measures within distance r are the same (that is, $F_\mathbf{X}(r) = F_\mathbf{Y}(r)$), $\mathrm{ID}_{F_\mathbf{Y}}(r)$ is greater than one would expect for a locally uniform distribution of points in $\mathbb{R}^2$, while $\mathrm{ID}_{F_\mathbf{X}}(r)$ is less.

The smallest distances from a given point can be regarded as 'extreme events' associated with the lower tail of the underlying distribution. The modeling of neighborhood distance values can thus be investigated from the viewpoint of extreme value theory (EVT). In [1], it is shown that the EVT representation of the distance distribution $F_\mathbf{X}$ completely determines function $\mathrm{ID}_{F_\mathbf{X}}$, and that the EVT index is in fact identical to $\mathrm{ID}_{F_\mathbf{X}}(0)$.

Theorem 1 ([1]). *Let $F : (0, z) \to \mathbb{R}$ be a function over the range $(0, z)$, for some choice of $z > 0$ (possibly infinite). Let $v \in [0, z)$ be a value for which $\mathrm{ID}_F(v)$ exists. Then for any $r, w \in (0, z)$ such that F is positive and differentiable everywhere over an open interval containing $[\min\{r, w\}, \max\{r, w\}]$,*

$$\frac{F(r)}{F(w)} = \left(\frac{r}{w}\right)^{\mathrm{ID}_F(v)} \cdot G_{F,v,w}(r), \text{ where} \tag{6}$$

$$G_{F,v,w}(r) \triangleq \exp\left(\int_r^w \frac{\mathrm{ID}_F(v) - \mathrm{ID}_F(t)}{t}\,\mathrm{d}t\right)\,. \tag{7}$$

Moreover, let $c > 1$ be a constant, and assume that $\mathrm{ID}_F(0)$ exists. Then

$$\lim_{\substack{w \to 0^+ \\ 0 < w/c \le r \le cw}} G_{F,0,w}(r) \;=\; 1\,. \tag{8}$$

Proof. See [13].

This theorem shows that, asymptotically, a relation between probability measure (volume) and distances (radius) similar to equation 2 holds. The dimension involved is not the representational one anymore, but the LID. The representational dimension disappear because the LID is simply a property of the distribution of distances; the space in which the data live do not matter.

Practical methods that have been developed for the estimation of the index, including expansion-based estimators [4] and the well-known Hill estimator and its variants [7], can all be applied to LID (for a survey, see [5]).

2.3 Previous results

In this section, we refer to the result given in [13]. Our theoretical work, exposed in Section 3, is a generalization and an improvement of this result.

As there is a work with infinitesimal noise, the space the data live in should be continuous Also, angles and distances are to be manipulated. Thus, the work is done in an Euclidean space, assuming that data can be represented in this space.

Also, all the perturbations theorems presented here are asymptotic, either this one or our results in Section 3. We do not work with a fixed database, but with a sequence of such, increasing the number of points. Thus, we also work by expectation. Let $\mathbf{x}$ be a reference point within the Euclidean space $\mathcal{S}$, and $F_{\mathbf{X}}$ the cumulative distribution function of the distribution of distances from $\mathbf{x}$. Let $n \in \mathbb{N}^*$ be the size of the database, and $k < n$. We say that $\mathbf{z}$ is the k-nearest neighbour (k-NN) of x by expectation when $F_{\mathbf{X}}(\mathbf{z}) = {}^{k}/_{n}$.

In [13] a technical lemma is given at first. A reference point $\mathbf{x}$ and a target point $\mathbf{z}$ are defined. Then $\mathbf{x}$ is perturbed toward $\mathbf{z}$, their new distance being reduced by a factor $1 - \delta$, defining the perturbed point $\mathbf{y}$. For any choice of probabilities p and q such that $0 < p < q < 1$, Lemma 1 provides conditions on the proportion δ, which when satisfied guarantee that $F_{\mathbf{Y}}(d(\mathbf{y}, \mathbf{z})) \leq p$ even when $F_{\mathbf{X}}(d(\mathbf{x}, \mathbf{z})) = q > p$.

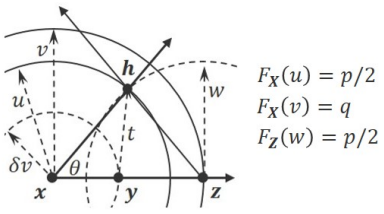

Fig. 3: An illustration of the construction by which the reference point $\mathbf{x}$ is perturbed towards a target point $\mathbf{z}$, so that the rank of $\mathbf{z}$ relative to the resulting perturbed point $\mathbf{y}$ is at most a target value.

Lemma 1. *Let* $\mathbf{x}$ *be a reference point within the Euclidean space* $\mathcal{S}$*. Let* $0 < \delta < 1/2$ *be a fixed real value, and let* p *and* q *be real values such that* $0 < p < q < 1$*. Consider the following construction with points* $\mathbf{y}$ *and* $\mathbf{z}$ *and distance parameters* r*,* u*,* v*, and* w *all depending on* p*,* q *and* δ *(see Figure 3):*

1. *Let* u *and* v *be the infimums of the distances from* $\mathbf{x}$ *at which* $F_{\mathbf{X}}(u) = p/2$ *and* $F_{\mathbf{X}}(v) = q$*, respectively.*
2. *Let* $\mathbf{z} \in \mathcal{S}$ *be any point for which* $d(\mathbf{x}, \mathbf{z}) = v$*.*
3. *Let* $\mathbf{y} \in \mathcal{S}$ *be the point lying on the line segment joining* $\mathbf{x}$ *and* $\mathbf{z}$*, at distance* δv *from* $\mathbf{x}$*.*
4. *Let* r *be the supremum of the distances from* $\mathbf{y}$ *at which* $F_{\mathbf{Y}}(r) = p$*, and let* w *be the infimum of the distances from* $\mathbf{z}$ *at which* $F_{\mathbf{Z}}(w) = p/2$*.*

If $\delta \geq (v^2 - u^2)/(v^2 - u^2 + w^2)$*, then* $d(\mathbf{y}, \mathbf{z}) \leq r$*.*

Proof. See [13]

Under certain assumptions of the smoothness of the underlying data distribution, the construction of Lemma 1 can be used to relate the effect of perturbation on neighbourhoods to the intrinsic dimensionality of the distance distribution from the perturbed point. We say that the local intrinsic dimensionality of the euclidean space $\mathcal{S}$ is continuous at $\mathbf{x} \in \mathcal{S}$ if the following conditions hold:

1. There exists a distance $\rho > 0$ for which all points $\mathbf{z} \in \mathcal{S}$ with $d(\mathbf{x}, \mathbf{z}) \leq \rho$ admit a distance distribution whose cumulative distribution function $F_{\mathbf{Z}}$ is differentiable and positive within some open interval with lower bound 0.
2. $F_{\mathbf{Z}}$ converges in distribution to $F_{\mathbf{X}}$ as $\mathbf{z} \to \mathbf{x}$.
3. For each $\mathbf{z}$ satisfying the condition above, $\mathrm{ID}_{F_{\mathbf{Z}}}(0)$ exists.
4. $\lim_{\mathbf{z} \to \mathbf{x}} \mathrm{ID}_{F_{\mathbf{Z}}}(0) = \mathrm{ID}_{F_{\mathbf{X}}}(0)$.

Note that if the distributions $F_{\mathbf{Z}}$ are assumed to be absolutely continuous, and if uniform convergence is assumed in the second condition, then the third and fourth conditions would follow from the first two conditions.

Theorem 2. *Let* $\mathbf{x}$ *be a reference point within the Euclidean space* $\mathcal{S}$*, and let* $F_{\mathbf{X}}$ *be the cumulative distribution function of the distribution of distances from* $\mathbf{x}$*. Let us assume that the local intrinsic dimensionality of* $\mathcal{S}$ *is continuous at* $\mathbf{x}$*. For any real constant* $k > 1$*, consider the real-valued parameter* n *chosen such that* $n > k$*, and let* ρ_n *be the infimum of the distances for which the cumulative distribution function* $F_{\mathbf{X}}$ *achieves* k/n *— that is,* $\rho_n = \inf\{\rho |\, F_{\mathbf{X}}(\rho) = k/n\}$*.*

Let $0 < \delta < 1/2$ *be a fixed real value. With respect to the particular choice of* n*, let* $\mathbf{z}_n \in \mathcal{S}$ *be any point for which* $d(\mathbf{x}, \mathbf{z}_n) = \rho_n$*, and let* $\mathbf{y}_n \in \mathcal{S}$ *be the point lying on the line segment joining* $\mathbf{x}$ *and* $\mathbf{z}_n$ *at distance* $\delta \cdot d(\mathbf{x}, \mathbf{z}_n)$ *from* $\mathbf{x}$*. Then for every real value* $\varepsilon > 0$*, there exists* $n_0 > k$ *such that for all* $n \geq n_0$*, we have that*

$$\delta \geq 1 - (2k)^{\frac{-2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} + \varepsilon \implies F_{\mathbf{Y}_n}\left(d(\mathbf{y}_n, \mathbf{z}_n)\right) \leq 1/n. \tag{9}$$

Proof. See [13]

Thus, the noise required to have the target point become the first nearest neighbour of the perturbed point is small. More precisely, it tends to 0 when the number of points and the intrinsic dimensionality at the reference's location both increase.

3 Theoretical results

In this section, we present a generalization of the result given in Theorem 2, extending it to other neighbourhood manipulations.

We distinguished 4 similar uses of adversarial noise. We define 3 points, $\mathbf{x}$ the reference point, $\mathbf{z}$ the target point and $\mathbf{y}$ the perturbed point. For each case we give an initial state, linking $\mathbf{z}$ and $\mathbf{x}$, and a desired state, linking $\mathbf{z}$ and $\mathbf{y}$:

1. we have $\mathbf{z}$ 1-NN of $\mathbf{x}$, and we want to achieve $\mathbf{z}$ k-NN of $\mathbf{y}$: $\mathbf{x}$ removes $\mathbf{z}$ of its neighbourhood.
2. we have $\mathbf{x}$ 1-NN of $\mathbf{z}$, and we want to achieve $\mathbf{y}$ k-NN of $\mathbf{z}$: $\mathbf{x}$ evades $\mathbf{z}$'s neighbourhood.
3. we have $\mathbf{x}$ k-NN of $\mathbf{z}$, and we want to achieve $\mathbf{y}$ 1-NN of $\mathbf{z}$: $\mathbf{x}$ invades $\mathbf{z}$'s neighbourhood.
4. we have $\mathbf{z}$ k-NN of $\mathbf{x}$, and we want to achieve $\mathbf{z}$ 1-NN of $\mathbf{y}$: $\mathbf{x}$ captures $\mathbf{z}$ in its neighbourhood.

The fourth case identified is the one studied in [13], where an object is made misclassified, by perturbing him toward a new object (and its class). The results proved here are the same for the 4 cases: the amount of noise (relative distance between $\mathbf{z}$ and $\mathbf{y}$) required to achieve the desired effects tends to 0 as both the number of points in the distribution and the LID at $\mathbf{x}$ increases. That is to say: every neighbourhood manipulation is achievable with an infinitesimal noise, provided a condition on LID is met and the database is large.

The rest of this section is organised as such: we expose 4 preliminary lemmas for each case. Then we formalize the 4 situations, and prove that the noise required tends to 0 for each of them.

3.1 Preliminary lemmas

For the following section, we define:

- Let $\mathbf{x}$ be a reference point within the Euclidean space $\mathcal{S}$.
- Let $0 < \delta$ be a fixed real value.
- Let p and q be real values such that $0 < p < q < 1$.

For each of the following lemmas, we will start by choosing a point $\mathbf{z}$ meeting a given condition relatively to $\mathbf{x}$. We then define the perturbed point $\mathbf{y}$, always at distance δv from $\mathbf{x}$. Finally, we give sufficient conditions on δ such that $\mathbf{z}$ meets a new condition relatively to $\mathbf{y}$. Figure 4 illustrates the different positions of $\mathbf{y}$ possible and the notations introduced in the lemmas.

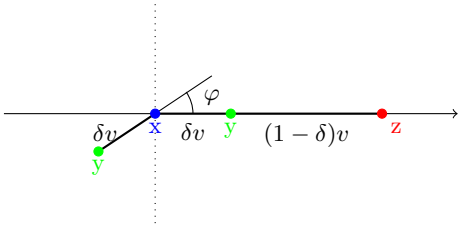

Fig. 4: Relations between $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$

Lemma 2. *1. Let v be the infimum of the distances from $\mathbf{x}$ at which $F_{\mathbf{X}}(v) = p$.*
2. Let $\mathbf{z} \in \mathcal{S}$ be any point for which $d(\mathbf{x}, \mathbf{z}) = v$.
3. Let $\mathbf{y} \in \mathcal{S}$ be a point at distance $\delta \cdot v$ from $\mathbf{x}$ such that $\vec{\mathbf{yx}} \cdot \vec{\mathbf{xz}} \geq 0$. Let then be $\varphi = \angle\vec{\mathbf{yx}}, \vec{\mathbf{xz}}$ (then angle between the vectors $\vec{\mathbf{yx}}$ and $\vec{\mathbf{xz}}$).
4. Let r be the supremum of the distances from $\mathbf{y}$ at which $F_{\mathbf{Y}}(r) = q$.
If $r > v$, then if $\delta > \sqrt{\left(\frac{r^2}{v^2} - 1\right) + \cos^2(\varphi)} - \cos(\varphi)$, then $F_{\mathbf{Y}}(d(\mathbf{y}, \mathbf{z})) > q$.

Lemma 3. *1. Let $\mathbf{z} \in \mathcal{S}$ be any point for which $F_{\mathbf{Z}}(d(\mathbf{x}, \mathbf{z})) = p$, and $v = d(\mathbf{x}, \mathbf{z})$.*
2. Let $\mathbf{y} \in \mathcal{S}$ be a point at distance $\delta \cdot v$ from $\mathbf{x}$ such that $\vec{\mathbf{yx}} \cdot \vec{\mathbf{xz}} \geq 0$. Let then be $\varphi = \angle\vec{\mathbf{yx}}, \vec{\mathbf{xz}}$.
3. Let r be the supremum of the distances from $\mathbf{z}$ at which $F_{\mathbf{Z}}(r) = q$.
If $\delta > \sqrt{\left(\frac{r^2}{v^2} - 1\right) + \cos^2(\varphi)} - \cos(\varphi)$, then $F_{\mathbf{Z}}(d(\mathbf{y}, \mathbf{z})) > q$.

Lemma 4. *1. Let $\mathbf{z} \in \mathcal{S}$ be any point for which $F_{\mathbf{Z}}(d(\mathbf{x}, \mathbf{z})) = q$, and $v = d(\mathbf{x}, \mathbf{z})$.*
2. Let $\mathbf{y} \in \mathcal{S}$ be the point at distance $\delta \cdot v$ from $\mathbf{x}$ lying in the segment joining $\mathbf{x}$ and $\mathbf{z}$.
3. Let r be the infimum of the distances from $\mathbf{z}$ at which $F_{\mathbf{Z}}(r) = p$.
If $\delta > 1 - \frac{r}{v}$, then $F_{\mathbf{Z}}(d(\mathbf{y}, \mathbf{z})) < p$.

Lemma 5. *1. Let v be the infimum of the distances from $\mathbf{x}$ at which $F_{\mathbf{X}}(v) = q$.*
2. Let $\mathbf{z} \in \mathcal{S}$ be any point for which $d(\mathbf{x}, \mathbf{z}) = v$.
3. Let $\mathbf{y} \in \mathcal{S}$ be a point at distance $\delta \cdot v$ from $\mathbf{x}$ lying in the segment joining $\mathbf{x}$ and $\mathbf{z}$.
4. Let r be the infimum of the distances from $\mathbf{y}$ at which $F_{\mathbf{Y}}(r) = p$.
If $\delta > 1 - \frac{r}{v}$, then $F_{\mathbf{Y}}(d(\mathbf{y}, \mathbf{z})) < p$.

Remark 1. The proofs are all similar, and can be found in Appendix A).

Remark 2. The first lemma is the only one to include an extra pre-condition ($r > v$). It does not compromise the generality of the result, as it can be shown show that this condition asymptotically holds (*cf.* proof of Theorem 3).

3.2 Main theorems

We define:

- Let $\mathbf{x}$ be a reference point within the Euclidean space $\mathcal{S}$, and let $F_{\mathbf{X}}$ be the cumulative distribution function of the distribution of distances from $\mathbf{x}$. We assume that the LID of $\mathcal{S}$ is continuous at $\mathbf{x}$.
- Let $0 < \delta$ be a fixed real value.
- Let $k > 1$ be any real constant.

Theorem 3. *Consider the real-valued parameter n chosen such that $n > k$, and let ρ_n be the infimum of the distances for which the cumulative distribution function $F_{\mathbf{X}}$ achieves $1/n$ — that is, $\rho_n = \inf\{\rho |\, F_{\mathbf{X}}(\rho) = 1/n\}$.*

With respect to the particular choice of n, let $\mathbf{z}_n \in \mathcal{S}$ be any point for which $d(\mathbf{x}, \mathbf{z}_n) = \rho_n$, and let $\mathbf{y}_n \in \mathcal{S}$ be the point at distance $\delta \cdot d(\mathbf{x}, \mathbf{z}_n)$ from $\mathbf{x}$ such that $\varphi_n = \angle \vec{\mathbf{y}_n\mathbf{x}}, \vec{\mathbf{xz}_n} \in [-\pi/2, \pi/2]$.

Then for every real value $\varepsilon > 0$, there exists $n_0 > k$ such that for all $n \geq n_0$, we have that:

$$\delta > \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}(0)}}} - 1} + \varepsilon \Longrightarrow F_{\mathbf{Y}_n}(d(\mathbf{y}_n, \mathbf{z}_n)) > k/n \tag{10}$$

Theorem 4. *Consider the real-valued parameter n chosen such that $n > k$.*

With respect to the particular choice of n, let $\mathbf{z}_n \in \mathcal{S}$ be any point for which $F_{\mathbf{Z}_n}(d(\mathbf{x}, \mathbf{z}_n)) = 1/n$, and let $\mathbf{y}_n \in \mathcal{S}$ be the point at distance $\delta \cdot d(\mathbf{x}, \mathbf{z}_n)$ from $\mathbf{x}$ such that $\varphi_n = \angle \vec{\mathbf{y}_n\mathbf{x}}, \vec{\mathbf{xz}_n} \in [-\pi/2, \pi/2]$.

Then for every real value $\varepsilon > 0$, there exists $n_0 > k$ such that for all $n \geq n_0$, we have that:

$$\delta > \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}(0)}}} - 1} + \varepsilon \Longrightarrow F_{\mathbf{Z}_n}(d(\mathbf{y}_n, \mathbf{z}_n)) > k/n \tag{11}$$

Theorem 5. *Consider the real-valued parameter n chosen such that $n > k$.*

With respect to the particular choice of n, let $\mathbf{z}_n \in \mathcal{S}$ be any point for which $F_{\mathbf{Z}_n}(d(\mathbf{x}, \mathbf{z}_n)) = k/n$, and let $\mathbf{y}_n \in \mathcal{S}$ be the point at distance $\delta \cdot d(\mathbf{x}, \mathbf{z}_n)$ from $\mathbf{x}$ lying on the segment joining $\mathbf{x}$ and $\mathbf{z}_n$.

Then for every real value $\varepsilon > 0$, there exists $n_0 > k$ such that for all $n \geq n_0$, we have that:

$$\delta > 1 - k^{\frac{-1}{\mathrm{ID}_{F_{\mathbf{X}}(0)}}} + \varepsilon \Longrightarrow F_{\mathbf{Z}_n}(d(\mathbf{y}_n, \mathbf{z}_n)) < 1/n \tag{12}$$

Theorem 6. *Consider the real-valued parameter n chosen such that $n > k$.*

With respect to the particular choice of n, let $\mathbf{z}_n \in \mathcal{S}$ be any point for which $F_{\mathbf{X}}(d(\mathbf{x}, \mathbf{z}_n)) = k/n$, and let $\mathbf{y}_n \in \mathcal{S}$ be the point at distance $\delta \cdot d(\mathbf{x}, \mathbf{z}_n)$ from $\mathbf{x}$ lying on the segment joining $\mathbf{x}$ and $\mathbf{z}_n$.

Then for every real value $\varepsilon > 0$, there exists $n_0 > k$ such that for all $n \geq n_0$, we have that:

$$\delta > 1 - k^{\frac{-1}{\mathrm{ID}_{F_{\mathbf{X}}(0)}}} + \varepsilon \Longrightarrow F_{\mathbf{Y}_n}(d(\mathbf{y}_n, \mathbf{z}_n)) < 1/n \tag{13}$$

Remark 3. Once again, all the proofs are similar, and are given in Appendix A.

Remark 4. Theorem 6 states the same asymptotic effect as the main theorem of the paper this work is based on [13]. However, the bound we exhibit here is tighter, and the proof of both the theorem and the corresponding lemma have been simplified. One can refer to Figure 6 for a comparison of the two theoretical bounds.

Remark 5. For each scenario, instead of manipulating 1-NN and k-NN, one can manipulate respectively k_1-NN and k_2-NN ($k_1 < k_2$). The theoretical bounds would be the same, except that every occurrence of k would be replaced by ${}^{k_2}/_{k_1}$.

3.3 Handling multiple points

The third and fourth cases consider a movement along a segment, whereas the first and second cases consider a movement in a half-space. Therefore, it is possible to handle multiple points for these two scenarios, if the intersection of all the half-spaces is not reduced to one point. More precisely, given a reference point $\mathbf{x}$ and a set of target points $\{\mathbf{z}_1, \cdots, \mathbf{z}_n\}$, can we find a perturbed point $\mathbf{y} \neq \mathbf{x}$ such that $\forall i \in [\![1, n]\!]\ \vec{\mathbf{xy}} \cdot \vec{\mathbf{xz}}_i \leq 0$?

Here, we are only interested in the direction of $\vec{\mathbf{xy}}$, and not in its length. The asymptotic effects examined earlier still apply when handling a fixed number of points: the amount of noise required will tend to zero for each point.

Definition 2. *Let $d \in \mathbb{N}^*$ a dimension and $n \in \mathbb{N}^*$ a number of vectors. Let $(\mathbf{u}_1, \cdots, \mathbf{u}_n) \in \left(\mathbb{R}^d\right)^n$ a family of vectors. We say that $(\mathbf{u}_1, \cdots, \mathbf{u}_n)$* positively spans *$\mathbb{R}^d$ if, and only if, every vector in $\mathbb{R}^d$ can be expressed as a positive linear combination of $\mathbf{u}_1, \cdots, \mathbf{u}_n$. Formally:*

$$\begin{aligned} &(u_i)_{1\leq i\leq n} \text{ positively spans } \mathbb{R}^d \\ &\iff \forall v \in \mathbb{R}^d,\ \exists\, (\lambda_1, \cdots, \lambda_n) \in (\mathbb{R}_+)^n\ v = \textstyle\sum_{i=1}^n \lambda_i \cdot u_i \end{aligned} \tag{14}$$

Having defined the notion of positively spanning the space, we link it to our goal:

Theorem 7. *Let $d \in \mathbb{N}^*$ and $n \in \mathbb{N}^*$. Let $\mathbf{x} \in \mathbb{R}^d$ and $(\mathbf{z}_i)_{1\leq i\leq n} \in (\mathbb{R}^d)^n$.*

$$\begin{aligned} &\exists\, \mathbf{y} \neq \mathbf{x},\ \forall i \in [\![1, n]\!],\ \vec{\mathbf{xy}} \cdot \vec{\mathbf{xz}}_i \leq 0 \\ &\iff (\vec{\mathbf{xz}}_i)_{1\leq i\leq n} \text{ does not positively spans } \mathbb{R}^d \end{aligned} \tag{15}$$

Proof. "If": Let $\mathbf{y}$ be a point such that $\vec{\mathbf{xy}}$ is out of the cone spanned by $(\vec{\mathbf{xz}}_i)_{1\leq i\leq n}$. Using Farkas' lemma, we know there is a hyperplane separating the vector $\vec{\mathbf{xy}}$ from the cone. Thus, (a) can be respected by taking a new point $\mathbf{y}'$ such that $\vec{\mathbf{xy}}'$ is normal to the hyperplane and y' is in the same half-space as y.

"Only if": Let $\mathbf{y}$ be a point different from $\mathbf{x}$ satisfying $\forall i \in [\![1, n]\!],\ \vec{\mathbf{xy}} \cdot \vec{\mathbf{xz}}_i \leq 0$. Again, using Farkas' lemma:

- either $\vec{\mathbf{xy}}$ lies within the cone spanned by $(\vec{\mathbf{xz}}_i)_{1\leq i\leq n}$. Then there exists $(\lambda)_{1\leq i\leq n} \in \mathbb{R}_+^d$ such that $\vec{\mathbf{xy}} = \sum_{i=1}^{n} \lambda_i \vec{\mathbf{xz}}_i$. Thus $\|\vec{\mathbf{xy}}\|^2 = \sum_{i=1}^{n} \lambda_i \vec{\mathbf{xy}} \cdot \vec{\mathbf{xz}}_i \leq 0$, hence $\mathbf{y} = \mathbf{x}$, which is impossible.
- either there exists a hyperplane separating $\vec{\mathbf{xy}}$ from the cone spanned by $(\vec{\mathbf{xz}}_i)_{1\leq i\leq n}$. In this case, $(\vec{\mathbf{xz}}_i)_{1\leq i\leq n}$ does not positively spans $\mathbb{R}^d$.

Having established this theorem, we give a new characterization of the property "positively spanning the space"; it allows us to deduce an algorithm to check this property.

Theorem 8. *Let $d \in \mathbb{N}^*$ and $n \in \mathbb{N}^*$. Let $\mathbf{x} \in \mathbb{R}^d$ and $(\mathbf{z}_i)_{1\leq i\leq n} \in (\mathbb{R}^d)^n$.*

$$(\vec{\mathbf{xz}}_i)_{1\leq i\leq n} \text{ positively spans } \mathbb{R}^d \iff \begin{cases} \text{(a) } (\vec{\mathbf{xz}}_i)_{1\leq i\leq n} \text{ linearly spans } \mathbb{R}^d \\ \text{(b) } \forall i \in [\![1,n]\!] - \vec{\mathbf{xz}}_i \text{ is in the convex} \\ \quad \text{cone spanned by the remaining } \vec{\mathbf{xz}}_j \end{cases} \tag{16}$$

Proof. See theorems 3.6 and 3.7 from [12].

Remark 6. (a) can be checked easily (by rank computation for example). For (b), each of the n sub-conditions can be expressed by the following optimisation problem:

- Minimize: $f : (\lambda_j)_{1\leq j\leq n, j\neq i} \longrightarrow \| \sum_{1\leq j\leq n, j\neq i} (\lambda_j \vec{\mathbf{xz}}_j) + \vec{\mathbf{xz}}_i \|$
- Respecting: $\forall j \in [\![1,n]\!],\ j \neq i,\ \lambda_j \geq 0$

We define $\Lambda_i = (\lambda_j)_{1\leq j\leq n, j\neq i}$. Achieving $f(\Lambda_i) = 0$ for some i means that $-\vec{\mathbf{xz}}_i$ is a positive linear combination of the remaining $\vec{\mathbf{xz}}_j$. Achieving $f(\Lambda_i) = 0$ for all $i \in [\![1,n]\!]$ means that condition (b) is respected.

4 Experimental results

4.1 The adversarial effect

We conducted 4 series of experiments to confirm the trends announced by the theoretical results. For each experiment, we sample uniformly points in a d-dimensional hypercube, and compute the median amount of noise, across all points in the distribution, required to achieve a desired neighbourhood perturbation. Uniformly sampled points are expected to have the same LID as the representational dimension; thus, we expect the amount of noise computed to tend to 0 as the dimension increases.

For each scenario we use artificial data sets of 32000 points, representational dimension taking values in $(2, 4, 8, 16, 32, 64)$ and rank in $(2, 4, 8, 16, 32, 64, 128)$ (the rank is used to choose an index k when talking about the k-NN of a point). The results are shown in figures 5a to 5d, referring respectively theorems 3 to 6.

Figure 6 is borrowed from [13]. It corresponds of the fourth scenario we identified (getting the target point $\mathbf{z}$ to be the 1-NN of the reference point). It uses the BIGANN_SIFT1B dataset; the experiment involves $n = 10^9$ points in

dimension $d = 128$. For $n_q = 10000$ query points, the minimum value for δ to have the 100-NN become the 1-NN is plotted against the estimated LID at the query point position. The points are plotted in blue. the mean minimum noise level, along with the standard deviation, are plotted in green (LID values are grouped in integer bins). The theoretical bound from [13] is plotted in red; the bound from this report is plotted in yellow.

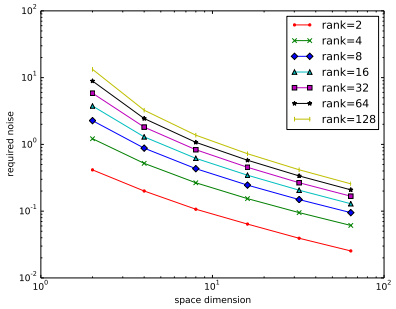

(a) Median amount of noise required for $\mathbf{x}$ to remove $\mathbf{z}$ from its neighbourhood

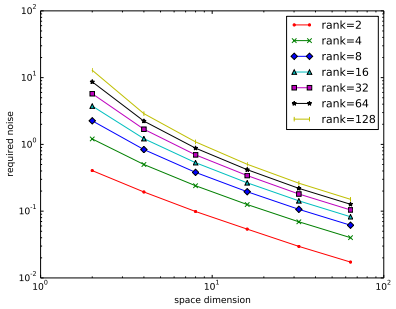

(b) Median amount of noise required for $\mathbf{x}$ to evade $\mathbf{z}$'s neighbourhood

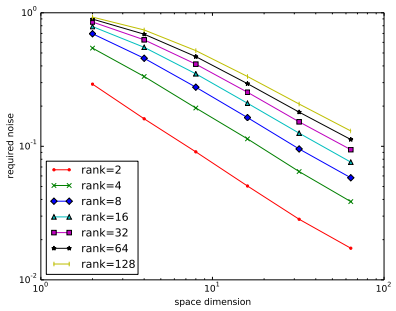

(c) Median amount of noise required for $\mathbf{x}$ to invade $\mathbf{z}$'s neighbourhood

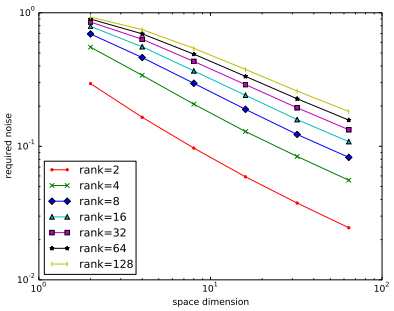

(d) Median amount of noise required for $\mathbf{x}$ to capture $\mathbf{z}$ in its neighbourhood

Fig. 5: Amount of noise required to achieve a neighbourhood perturbation

4.2 The feasibility of escaping multiple points

In Section 3.3, we studied how to check to feasibility of moving away from multiple points. However we do not know how many points we can expect to evade from at the same time. The problem is the following: given a set of points $\mathcal{X}$ in the Euclidean space $\mathcal{S}$ and a point $\mathbf{x} \in \mathcal{X}$, what is the maximum index i such that the family $(\mathbf{x}\vec{\mathbf{z}}_k)_{1 \leq k \leq i}$ does not positively spans $\mathcal{S}$, $(\mathbf{z}_k)_{1 \leq k \leq i}$ being the i nearest neighbours of $\mathbf{x}$ in $\mathcal{X}$.

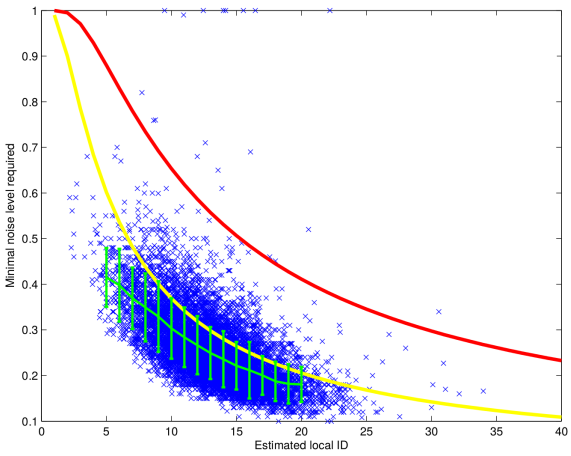

Fig. 6: Comparison of the theoretical bounds

We did not find a mention of a similar problem in the literature. Given the time, we only studied experimentally this question. Our experiments involve artificial data sets of uniformly sampled points in the d-dimensional hypercube.

In Figure 7a, we plotted the histogram of the maximum size of the neighbourhood one point of the data set can escape, for a data set of 204800 points in dimension 8. The bar on the left represent the proportion of points able to escape more than their 200 nearest neighbours. In Figure 7b, we plotted the proportion of points able to escape from their 100 nearest neighbours for data sets of size 1600, 3200, 6400, 12800, 25600 (and 51200) and dimension 2, 4, 8 and 16.

There is a lot to say about this problem, but it starts leaving the scope of this study. The shape of the the histogram in Figure 7a is very regular. Others experiments with different parameters yield the same kind of figures. It is unknown at the moment if is follows a known distribution. The functions plotted in Figure 7b are tending toward 0, which leads to an interesting conclusion: augmenting the number of points diminish the perturbation required for the adversarial effect, but also seems to diminish the chances of finding a direction of escape. Finally, it should be noted that this problem should be viewed with the representational dimension in mind, and not the LID. When using different distributions, the feasibility proportions do not move significantly.

5 Conclusion

The subject of this internship was the one of adversarial learning, where one tries to fool a machine learning system, here by applying small modification to objects in order to achieve a desired effect.

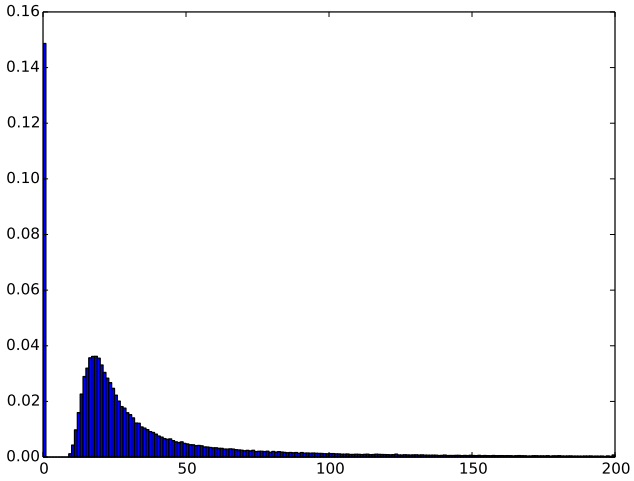

(a) Histogram of the size of the neighbourhood one point can escape

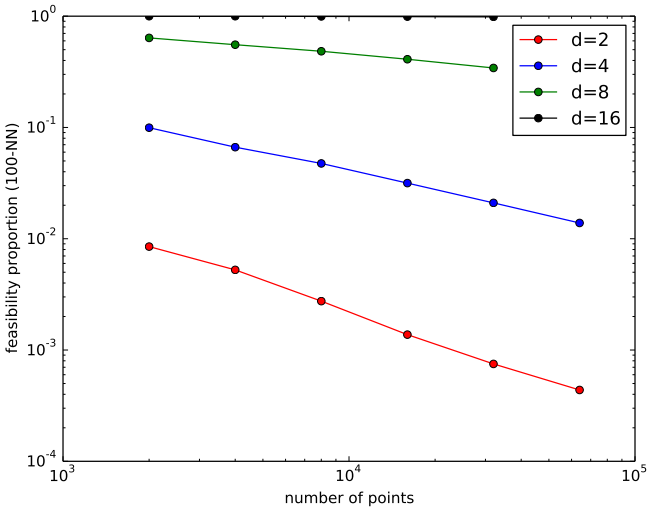

(b) Proportion of points able to evade their 100 nearest neighbours

Fig. 7: Experimental study of the feasibility of escaping multiple points

The results from [13] were extended, showing that any neighbourhood manipulations is achievable with a small perturbation, provided that both the size of the data set the LID are high. These results apply to any classifier of continuously-distributed data. The proofs given here are simpler — shorter, yet using the same arguments — than the ones from the the document this work is based on, and the theoretical bounds are also tighter. Then, an experimental observation of the trends announced in the theoretical part is provided.

Future works on this topic could include the construction of machine learning systems resistant to adversarial perturbations. Sufficient conditions are given for the adversarial effect to apply. Studying these conditions could lead to more robust systems. Also, it could be possible to generalize these results to more general spaces than the euclidean ones. As only distance and angles are manipulated, an extension to Hilbert spaces seems manageable.

References

1. M. E. Houle. Inlierness, outlierness, hubness and discriminability: an extreme-value-theoretic foundation. Technical Report NII-2015-002E, NII, Mar 2015. http://www.nii.ac.jp/TechReports/public_html/15-002E.pdf
2. D. R. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *STOC*, pages 741–750, 2002.
3. M. E. Houle, H. Kashima, and M. Nett. Generalized expansion dimension. In *ICDMW*, pages 587–594, 2012.
4. L. Amsaleg, O. Chelly, T. Furon, S. Girard, M. E. Houle, K. Kawarabayashi, and M. Nett. Estimating local intrinsic dimensionality. In *KDD*, pages 29–38, 2015. http://mistis.inrialpes.fr/~girard/Fichiers/p29-amsaleg.pdf
5. M. I. Gomes *et al.*, Statistics of extremes for IID data and breakthroughs in the estimation of the extreme value index: Laurens de Haan leading contributions. *Extremes*, 11:3–34, 2008.
6. Michael E. Houle. Dimensionality, Discriminability, Density & Distance Distributions. In ICDMW, pages 468–473, 2013.
7. Huisman, R. and Koedijk, K. G. and Kool, C. J. M. and Palm, F. Tail-index estimates in small samples Journal of Business and Economic Statistics, 19(2):208–216, 2001.
8. Lowd, Daniel and Meek, Christopher. Adversarial Learning In Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining, pages 641–647, 2005. http://doi.acm.org/10.1145/1081870.1081950,
9. C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013. http://arxiv.org/abs/1312.6199
10. Alexey Kurakin, Ian Goodfellow, Samy Bengio. Adversarial examples in the physical world http://arxiv.org/abs/1607.02533
11. Ian J. Goodfellow, Jonathon Shlens, Christian Szegedy. Explaining and Harnessing Adversarial Examples http://arxiv.org/abs/1412.6572
12. Chandler Davis, Theory of Positive Linear Dependence, American Journal of Mathematics, Vol. 76, 1954, pp. 733-746. http://dx.doi.org/10.2307/2372648
13. L. Amsaleg et al., The Vulnerability of Learning to Adversarial Perturbation Increases with Intrinsic Dimensionality, NII Technical Report (NII-2016-005E). http://www.nii.ac.jp/TechReports/public_html/16-005E.html

A Remaining proofs

We give here the omitted proofs of Section 3. As told before, they all follow the same layout.

Lemma 2

Proof.

$$\begin{aligned} F_{\mathbf{Y}}(d(\mathbf{y},\mathbf{z})) > q &\iff d(\mathbf{y},\mathbf{z}) > r && (17)\\ &\iff \|\vec{\mathbf{yz}}\|^2 > r^2 && (18)\\ &\iff \|\vec{\mathbf{yx}}\|^2 + \|\vec{\mathbf{xz}}\|^2 + 2\,\vec{\mathbf{yx}}\cdot\vec{\mathbf{xz}} > r^2 && (19)\\ &\iff \delta^2 v^2 + v^2 + 2\delta v^2\cos(\varphi) > r^2 && (20)\\ &\iff \delta^2 + 2\delta\cos(\varphi) + \left(1 - \frac{r^2}{v^2}\right) > 0 && (21) \end{aligned}$$

We assume $r > v$. Then, the polynomial $\delta^2 + 2\delta\cos(\varphi) + \left(1 - \frac{r^2}{v^2}\right)$ admits two real roots, and thus is strictly positive if $\delta > \sqrt{\left(\frac{r^2}{v^2} - 1\right) + \cos^2(\varphi)} - \cos(\varphi)$ (using its highest root).

Lemma 3

Proof.

$$\begin{aligned} F_{\mathbf{Z}}(d(\mathbf{y},\mathbf{z})) > q &\iff d(\mathbf{y},\mathbf{z}) > r && (22)\\ &\iff \|\vec{\mathbf{yz}}\|^2 > r^2 && (23)\\ &\iff \|\vec{\mathbf{yx}}\|^2 + \|\vec{\mathbf{xz}}\|^2 + 2\,\vec{\mathbf{yx}}\cdot\vec{\mathbf{xz}} > r^2 && (24)\\ &\iff \delta^2 v^2 + v^2 + 2\delta v^2\cos(\varphi) > r^2 && (25)\\ &\iff \delta^2 + 2\delta\cos(\varphi) + \left(1 - \frac{r^2}{v^2}\right) > 0 && (26) \end{aligned}$$

We have $q > p$, so $F_{\mathbf{Z}}(r) > F_{\mathbf{Z}}(v)$. $F_{\mathbf{Z}}$ being a monotone function, we have $r > v$. Thus, the polynomial $\delta^2 + 2\delta\cos(\varphi) + \left(1 - \frac{r^2}{v^2}\right)$ admits two real roots, and is strictly positive if $\delta > \sqrt{\left(\frac{r^2}{v^2} - 1\right) + \cos^2(\varphi)} - \cos(\varphi)$ (using its highest root).

Lemma 4

Proof.

$$\begin{aligned} F_{\mathbf{Z}}(d(\mathbf{y},\mathbf{z})) < p &\iff d(\mathbf{y},\mathbf{z}) < r && (27)\\ &\iff (1-\delta)\cdot v < r && (28)\\ &\iff 1 - \frac{r}{v} < \delta && (29) \end{aligned}$$

Lemma 5

Proof.

$$F_{\mathbf{Y}}(d(\mathbf{y},\mathbf{z})) < p \iff d(\mathbf{y},\mathbf{z}) < r \tag{30}$$

$$\iff (1-\delta)\cdot v < r \tag{31}$$

$$\iff 1 - \frac{r}{v} < \delta \tag{32}$$

Proof. (Theorem 3)

Theorem 3 For a given choice of n, consider the construction in the statement of Lemma 2, with $p = 1/n$ and $q = k/n$, where $\mathbf{z}_n = \mathbf{z}$, $\mathbf{y}_n = \mathbf{y}$, $\varphi = \varphi\ v_n = d(\mathbf{x},\mathbf{z}_n) = \rho_n$, and $r_n = r$. Next, when $r_n > v_n$, we define:

$$\delta_n \triangleq \sqrt{\left(\frac{r_n^2}{v_n^2} - 1\right) + \cos^2(\varphi_n)} - \cos(\varphi_n) \tag{33}$$

$$\delta_n' \triangleq \sqrt{\frac{r_n^2}{v_n^2} - 1} \tag{34}$$

$$\alpha_n \triangleq n \cdot F_{\mathbf{Y}_n}(v_n) \tag{35}$$

We note that $\delta_n' \geq \delta_n$. Thus, it suffices to take $\delta \geq \delta_n'$ to apply Lemma 2 to have $F_{\mathbf{Y}_n}(d(\mathbf{y}_n,\mathbf{z}_n)) < 1/n$.

Using the local ID characterization formula of Theorem 1, we observe that:

$$\frac{k}{\alpha_n} = \frac{F_{\mathbf{Y}_n}(r_n)}{F_{\mathbf{Y}_n}(v_n)} = \left(\frac{r_n}{v_n}\right)^{\mathrm{ID}_{F_{\mathbf{Y}_n}}(0)} \cdot G_{F_{\mathbf{Y}_n},0,v_n}(r_n) \tag{36}$$

Thus:

$$\frac{r_n}{v_n} = \left(\frac{k}{\alpha_n \cdot G_{F_{\mathbf{Y}_n},0,v_n}(r_n)}\right)^{1/\mathrm{ID}_{F_{\mathbf{Y}_n}}(0)} \tag{37}$$

Note that since $F_{\mathbf{Y}_n}$ is assumed to converge in distribution to $F_{\mathbf{X}}$ as $n \to \infty$, we have:

$$\lim_{n\to\infty} \alpha_n = \lim_{n\to\infty} n \cdot F_{\mathbf{Y}_n}(v_n) \tag{38}$$

$$= \lim_{n\to\infty} \lim_{m\to\infty} \left(\frac{F_{\mathbf{Y}_m}(v_n)}{F_{\mathbf{X}}(v_n)} \cdot \frac{F_{\mathbf{X}}(v_n)}{1/n}\right) \tag{39}$$

$$= \lim_{n\to\infty} \left(\frac{F_{\mathbf{X}}(v_n)}{F_{\mathbf{X}}(v_n)} \cdot \frac{1/n}{1/n}\right) \tag{40}$$

$$= 1 \tag{41}$$

Furthermore, Theorem 1 and the continuity of the intrinsic dimension of $\mathcal{S}$ imply that:

$$\lim_{n\to\infty} G_{F_{\mathbf{Y}_n},0,v_n}(r_n) = 1 \tag{42}$$

$$\lim_{n\to\infty} \mathrm{ID}_{F_{\mathbf{Y}_n}}(0) = \mathrm{ID}_{F_{\mathbf{X}}}(0) \tag{43}$$

Together, these two statements establish that:

$$\lim_{n\to\infty} \frac{r_n}{v_n} = k^{1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{44}$$

Since $k > 1$, r_n/v_n tends to a value strictly greater than 1. The first consequence of this limit is that the hypothesis $r_n \geq v_n$ asymptotically holds:

$$\exists n_1 \in \mathbb{N},\ \forall n \in \mathbb{N},\ n > n_1 \Rightarrow r_n > v_n \tag{45}$$

δ'_n is well defined for any $n > n_1$, and its limit is:

$$\lim_{n\to\infty} \delta_n = \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1} \tag{46}$$

For any real value $\varepsilon > 0$, the limit of δ_n ensures the existence of a constant $n_0 > n_1$ such that for all $n \geq n_0$, we have that:

$$\left| \delta_n - \left(\sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1} \right) \right| \leq \varepsilon \tag{47}$$

Any choice of δ satisfying:

$$\delta > \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1} + \varepsilon \tag{48}$$

thus ensures that $\delta > \delta_n$; from this, Lemma 2 can be applied with $p = 1/n$ and $q = k/n$ to yield:

$$F_{\mathbf{Y}_n}(d(y_n, z_n)) > k/n \tag{49}$$

as required.

Theorem 4

Proof. For a given choice of n, consider the construction in the statement of Lemma 3, with $p = 1/n$ and $q = k/n$, where $\mathbf{z}_n = \mathbf{z}$, $\mathbf{y}_n = \mathbf{y}$, $\varphi = \varphi$, $v_n = d(\mathbf{x}, \mathbf{z}_n)$, and $r_n = r$. Next we define:

$$\delta_n \triangleq \sqrt{\left(\frac{r_n^2}{v_n^2} - 1\right) + \cos^2(\varphi_n)} - \cos(\varphi_n) \tag{50}$$

$$\delta'_n \triangleq \sqrt{\frac{r_n^2}{v_n^2} - 1} \tag{51}$$

$$\tag{52}$$

We note that $\delta'_n \geq \delta_n$. Thus, it suffices to take $\delta \geq \delta'_n$ to apply Lemma 2 to have $F_{\mathbf{Z}_n}(d(\mathbf{y}_n, \mathbf{z}_n)) > k/n$.

Using the local ID characterization formula of Theorem 1, we observe that:

$$k = \frac{F_{\mathbf{Z}_n}(r_n)}{F_{\mathbf{Z}_n}(v_n)} = \left(\frac{r_n}{v_n}\right)^{\mathrm{ID}_{F_{\mathbf{Z}_n}}(0)} \cdot G_{F_{\mathbf{Z}_n},0,v_n}(r_n) \tag{53}$$

Thus:

$$\frac{r_n}{v_n} = \left(\frac{k}{G_{F_{\mathbf{Z}_n},0,v_n}(r_n)}\right)^{1/\mathrm{ID}_{F_{\mathbf{Z}_n}}(0)} \tag{54}$$

Furthermore, Theorem 1 and the continuity of the intrinsic dimension of $\mathcal{S}$ imply that:

$$\lim_{n\to\infty} G_{F_{\mathbf{Z}_n},0,v_n}(r_n) = 1 \tag{55}$$

$$\lim_{n\to\infty} \mathrm{ID}_{F_{\mathbf{Z}_n}}(0) = \mathrm{ID}_{F_{\mathbf{X}}}(0) \tag{56}$$

Together, these two statements establish that:

$$\lim_{n\to\infty} \frac{r_n}{v_n} = k^{1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{57}$$

Thus, the limit of δ_n is:

$$\lim_{n\to\infty} \delta'_n = \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1} \tag{58}$$

For any real value $\varepsilon > 0$, the limit of δ_n ensures the existence of a constant n_0 such that for all $n \geq n_0$, we have that:

$$\left|\delta'_n - \left(\sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1}\right)\right| \leq \varepsilon \tag{59}$$

Any choice of δ satisfying:

$$\delta > \sqrt{k^{\frac{2}{\mathrm{ID}_{F_{\mathbf{X}}}(0)}} - 1} + \varepsilon \tag{60}$$

thus ensures that $\delta > \delta_n$; from this, Lemma 3 can be applied with $p = 1/n$ and $q = k/n$ to yield:

$$F_{\mathbf{Z}_n}(d(y_n, z_n)) > k/n \tag{61}$$

as required.

Theorem 5

Proof. For a given choice of n, consider the construction in the statement of Lemma 4, with $p = 1/n$ and $q = k/n$, where $\mathbf{z}_n = \mathbf{z}$, $\mathbf{y}_n = \mathbf{y}$, $v_n = d(\mathbf{x}, \mathbf{z}_n)$ and $r_n = r$. Next we define:

$$\delta_n \triangleq 1 - \frac{r_n}{v_n} \tag{62}$$

Using the local ID characterization formula of Theorem 1, we observe that:

$$\frac{1}{k} = \frac{F_{\mathbf{Z}_n}(r_n)}{F_{\mathbf{Z}_n}(v_n)} = \left(\frac{r_n}{v_n}\right)^{\mathrm{ID}_{F_{\mathbf{Z}_n}}(0)} \cdot G_{F_{\mathbf{Z}_n},0,v_n}(r_n) \tag{63}$$

Thus:

$$\frac{r_n}{v_n} = \left(\frac{1}{k \cdot G_{F_{\mathbf{Z}_n},0,v_n}(r_n)}\right)^{1/\mathrm{ID}_{F_{\mathbf{Z}_n}}(0)} \tag{64}$$

Furthermore, Theorem 1 and the continuity of the intrinsic dimension of $\mathcal{S}$ imply that:

$$\lim_{n\to\infty} G_{F_{\mathbf{Z}_n},0,v_n}(r_n) = 1 \tag{65}$$

$$\lim_{n\to\infty} \mathrm{ID}_{F_{\mathbf{Z}_n}}(0) = \mathrm{ID}_{F_{\mathbf{X}}}(0) \tag{66}$$

Together, these two statements establish that:

$$\lim_{n\to\infty} \frac{r_n}{v_n} = k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{67}$$

Thus, the limit of δ_n is:

$$\lim_{n\to\infty} \delta_n = 1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{68}$$

For any real value $\varepsilon > 0$, the limit of δ_n ensures the existence of a constant n_0 such that for all $n \geq n_0$, we have that:

$$\left|\delta_n - \left(1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)}\right)\right| \leq \varepsilon \tag{69}$$

Any choice of δ satisfying:

$$\delta > 1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} + \varepsilon \tag{70}$$

thus ensures that $\delta > \delta_n$; from this, Lemma 4 can be applied with $p = 1/n$ and $q = k/n$ to yield:

$$F_{\mathbf{Z}_n}(d(y_n, z_n)) < 1/n \tag{71}$$

as required.

Theorem 6

Proof. For a given choice of n, consider the construction in the statement of Lemma 5, with $p = 1/n$ and $q = k/n$, where $\mathbf{z}_n = \mathbf{z}$, $\mathbf{y}_n = \mathbf{y}$, $v_n = d(\mathbf{x}, \mathbf{z}_n)$, and $r_n = r$. Next we define:

$$\delta_n \triangleq 1 - \frac{r_n}{v_n} \tag{72}$$

We also define $k_n \triangleq n \cdot F_{\mathbf{Y}_n}(v_n)$.

Using the local ID characterization formula of Theorem 1, we observe that:

$$\frac{1}{k_n} = \frac{F_{\mathbf{Y}_n}(r_n)}{F_{\mathbf{Y}_n}(v_n)} = \left(\frac{r_n}{v_n}\right)^{\mathrm{ID}_{F_{\mathbf{Y}_n}}(0)} \cdot G_{F_{\mathbf{Y}_n},0,v_n}(r_n) \tag{73}$$

by defining

Thus:

$$\frac{r_n}{v_n} = \left(\frac{1}{k_n \cdot G_{F_{\mathbf{Y}_n},0,v_n}(r_n)}\right)^{1/\mathrm{ID}_{F_{\mathbf{Y}_n}}(0)} \tag{74}$$

Note that since $F_{\mathbf{Y}_n}$ is assumed to converge in distribution to $F_{\mathbf{X}}$ as $n \to \infty$, we have:

$$\lim_{n\to\infty} k_n = \lim_{n\to\infty} n \cdot F_{\mathbf{Y}_n}(v_n) \tag{75}$$

$$= \lim_{n\to\infty} \lim_{m\to\infty} \left(\frac{F_{\mathbf{Y}_m}(v_n)}{F_{\mathbf{X}}(v_n)} \cdot \frac{F_{\mathbf{X}}(v_n)}{1/n}\right) \tag{76}$$

$$= \lim_{n\to\infty} \left(\frac{F_{\mathbf{X}}(v_n)}{F_{\mathbf{X}}(v_n)} \cdot \frac{k/n}{1/n}\right) \tag{77}$$

$$= k \tag{78}$$

Furthermore, Theorem 1 and the continuity of the intrinsic dimension of $\mathcal{S}$ imply that:

$$\lim_{n\to\infty} G_{F_{\mathbf{Y}_n},0,v_n}(r_n) = 1 \tag{79}$$

$$\lim_{n\to\infty} \mathrm{ID}_{F_{\mathbf{Z}_n}}(0) = \mathrm{ID}_{F_{\mathbf{X}}}(0) \tag{80}$$

Together, these two statements establish that:

$$\lim_{n\to\infty} \frac{r_n}{v_n} = k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{81}$$

Thus, the limit of δ_n is:

$$\lim_{n\to\infty} \delta_n = 1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} \tag{82}$$

For any real value $\varepsilon > 0$, the limit of δ_n ensures the existence of a constant n_0 such that for all $n \geq n_0$, we have that:

$$\left|\delta_n - \left(1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)}\right)\right| \leq \varepsilon \tag{83}$$

Any choice of δ satisfying:

$$\delta > 1 - k^{-1/\mathrm{ID}_{F_{\mathbf{X}}}(0)} + \varepsilon \tag{84}$$

thus ensures that $\delta > \delta_n$; from this, Lemma 5 can be applied with $p = 1/n$ and $q = k/n$ to yield:

$$F_{\mathbf{Y}_n}(d(y_n, z_n)) < 1/n \tag{85}$$

as required.