LOCAL INTRINSIC DIMENSIONALITY SIGNALS ADVERSARIAL PERTURBATIONS

Sandamal Weerasinghe*
School of Computing and Information Systems
University of Melbourne
Australia

Tansu Alpcan
Department of Electrical and Electronic Engineering
University of Melbourne
Australia

Sarah M. Erfani
University of Melbourne
Australia

Christopher Leckie
University of Melbourne
Australia

Benjamin I. P. Rubinstein
University of Melbourne
Australia

September 27, 2021

ABSTRACT

The vulnerability of machine learning models to adversarial perturbations has motivated a significant amount of research under the broad umbrella of adversarial machine learning. Sophisticated attacks may cause learning algorithms to learn decision functions or make decisions with poor predictive performance. In this context, there is a growing body of literature that uses local intrinsic dimensionality (LID), a local metric that describes the minimum number of latent variables required to describe each data point, for detecting adversarial samples and subsequently mitigating their effects. The research to date has tended to focus on using LID as a practical defence method often without fully explaining why LID can detect adversarial samples. In this paper, we derive a lower-bound and an upper-bound for the LID value of a perturbed data point and demonstrate that the bounds, in particular the lower-bound, has a positive correlation with the magnitude of the perturbation. Hence, we demonstrate that data points that are perturbed by a large amount would have large LID values compared to unperturbed samples, thus justifying its use in the prior literature. Furthermore, our empirical validation demonstrates the validity of the bounds on benchmark datasets.

Keywords local intrinsic dimensionality · adversarial machine learning · poisoning attack

1 Introduction

Machine learning is fast becoming a key instrument in domains such as finance, cybersecurity, and engineering. However, there is a growing body of literature that suggests most machine learning algorithms, including deep neural networks (DNNs), are vulnerable to adversarial attacks. By adding carefully crafted perturbations to data points, adversaries aim to either avoid detection (i.e., test time attack) or alter the decision functions learned by the machine learning models (i.e., training data poisoning). As a consequence, applications that rely on machine learning for high-stakes automated decision making may take incorrect actions with severe consequences. Therefore, investigating defense mechanisms against adversarial attacks is a continuing concern within the adversarial machine learning community.

Adversarial attacks take place in situations where attackers have opportunities to alter existing data points or introduce new malicious data points. For example, training data poisoning can occur through the use of malware or when training data is collected using crowd-sourcing marketplaces, where organizations build data sets with the help of individuals whose authenticity cannot be guaranteed. Due to the size and complexity of datasets, it is infeasible to extensively

---

*email: prameesha.weerasinghe@unimelb.edu.au

validate the quality of all data/labels. Therefore, researchers have explored multiple approaches to address this problem and reduce the vulnerability of machine learning algorithms to adversarial attacks.

Recent research on adversarial machine learning defenses can be categorized as either a certified defense or an empirical defense. A certified defense attempts to learn provably robust prediction models against norm-bounded adversarial perturbations such as ℓ_2 or ℓ_∞. For example, the recent works of Wong and Kolter [2018], Cohen et al. [2019] and Lecuyer et al. [2019] introduce certified defenses particularly for deep neural networks. In contrast, empirical defenses either incorporate adversarial data into the training process [Zhang et al., 2020, Madry et al., 2018, Wang et al., 2019] or incorporate adversarial sample detection components to the machine learning pipeline. More recently, several works have explored empirical defenses that use Local Intrinsic Dimensionality (LID) [Houle, 2017a] as the basis for detecting adversarial samples.

The LID value of a data point identifies the minimum number of latent variables required to represent that data point and recent evidence suggests that LID is an indicator of the degree of being an outlier [Houle et al., 2018]. For example, Ma et al. [2018a] use LID to characterize adversarial subspaces in the context of Deep Neural Networks (DNNs). The authors use a novel learning strategy for DNNs based on LID that monitors the dimensionality of subspaces during training and adapts the loss function to minimize the impact of noisy labels. Weerasinghe et al. [2021a] introduce a black-box defense for nonlinear regression built around LID to identify and reduce the impact of adversarial samples.

Although previous works have established the intuition behind using LID for detecting adversarial samples, they fail to provide a theoretical justification for this behavior of LID. In this paper, we aim to address this by obtaining a lower-bound (and an upper-bound) for the LID value of a perturbed data point and show that there is an association between the size of the perturbation and the (lower-bound of the) LID value of the resulting data point. Specifically, we theoretically analyze the impact of adversarial perturbations using a simple setup comprising three data points, (i) the benign data point being perturbed, $\mathbf{a}$, (ii) its position after being perturbed, $\mathbf{b}$, and (iii) a benign reference data point, $\mathbf{c}$. Using the distribution of distances from the perturbed location (i.e., $\mathbf{b}$), we present a lower bound and an upper bound for its LID value (LID($\mathbf{b}$)). Furthermore, by considering the two possible directions of perturbation (i.e., towards or farther from the reference point $\mathbf{c}$), we identify the least lower and upper bounds when perturbed towards $\mathbf{c}$ and the greatest lower and upper bounds when perturbed farther from $\mathbf{c}$. These theorems indicate a positive association between the LID value and the size of the perturbations, thereby justifying the use of LID as an indicator of adversarial perturbations in previous works.

The main contributions of this paper are summarized as follows. We present a lower bound and an upper bound for the LID value of a perturbed data point based on the distribution of distances to its neighbors. We then identify the least lower and upper bounds when the perturbation is towards the reference point, and the greatest upper and lower bounds when the perturbation is away from it. Finally, we provide empirical evidence on benchmark data sets to support our theoretical contributions.

2 Literature Review

2.1 Theory of Local Intrinsic Dimensionality

The minimum number of latent variables needed to represent the data is considered as the intrinsic dimensionality (ID) of that particular data set. The ID of a dataset is thus an indicator of its complexity. A number of previous works have proposed characterizations of global intrinsic dimensionality (for data sets) and local intrinsic dimensionality (for a single data point) [Houle et al., 2012]. More recently Houle [2017a] proposed a characterization of LID in which the distribution of distances to a query point is modeled in terms of a continuous random variable [Houle, 2017a,b]. Amsaleg et al. [2015] developed several estimators of LID based on extreme value theory, using maximum likelihood estimation (MLE), the method of moments (MoM), probability weighted moments (PWM) and regularly varying functions (RV). Of these, the MLE model of LID estimation has been used for LID estimation in several works in adversarial machine learning including this paper (Definition 2). The accuracy of the MLE estimator is dependent on the number of neighbors considered in the calculation. More recently, Amsaleg et al. [2019] proposed an estimator of LID that achieves accurate results even with small neighborhood sizes.

Intrinsic dimensionality (through the use of estimators) has been utilized in many applications such as dimensionality reduction, outlier detection, similarity search and subspace clustering. In this paper, we focus on its ability to signal the degree of a sample being an outlier in the context of adversarial machine learning.

2.2   Applications of LID to Adversarial Machine Learning

To date, several studies have investigated the usability of LID as a detector of adversarial samples. For example, Ma et al. [2018b] use LID for detecting adversarial samples in Deep Neural Networks (DNNs). The authors use the LID values of the deep neural network representation (post-activation values at each layer) of images to train a classifier that identifies flipped labels with high success. For a similar application, Ma et al. [2018a] use a dimensionality-driven learning strategy which dynamically adapts the loss function of the DNN based on the dimensionality of subspaces (measured using LID).

Weerasinghe et al. [2021b] introduce a novel defense mechanism against poisoning and label flipping attacks using Local Intrinsic Dimensionality (LID). The proposed defense employs several label dependent variants of LID (i.e., in-class LID, out-class LID and cross-class LID) against label flipping attacks. Weerasinghe et al. [2021c] proposed a black-box defense for nonlinear regression learning. First, the defense identifies samples that have the greatest influence on the learner's validation error and then uses LID to iteratively identify poisoned samples via a generative probabilistic model, and suppress their influence on the prediction function.

Barua et al. [2019] apply a LID measure called CrossLID to assess the degree to which the distribution of data generated by a Generative Adversarial Network (GAN) conforms to the distribution of real data. The authors demonstrate that by incorporating the proposed LID based evaluation into the learning process, the generation quality of the GAN improves.

The work of Amsaleg et al. [2020] is the most closely related to the work presented in this paper. Amsaleg et al. [2020] introduced an upper-bound for the minimum perturbation required (δ) to change the rank of a data point from 1 to k (or k to 1) w.r.t. a reference point in the context of an information retrieval application. Furthermore, the authors demonstrate that small perturbations are sufficient to change the ranks of data points with relatively larger LID values compared to data points with smaller LID values.

In this work, we provide an upper-bound and a lower-bound for the LID value of a perturbed sample based on the distances (and the cumulative probability distribution of distances) to its starting location and a benign reference point. Furthermore, we identify the greatest (least) lower-bound and upper-bound when the data point is perturbed away from (towards) the reference point.

3   Theoretical Overview of LID

A particular challenge when using high dimensional data (e.g., images, natural language) is identifying lower dimensional representations with minimal loss of information. A considerable amount of literature has been published on characterizing the intrinsic dimensionality (i.e., the minimum number of latent variables required to describe the data) of data. These include fractal methods such as the Hausdorff dimension [Gneiting et al., 2012] and Minkowski–Bouligand dimension [Falconer, 2004], topological approaches [Rozza et al., 2012], and dimensional models such as minimum neighbor distance models [Rozza et al., 2012], expansion dimension [Karger and Ruhl, 2002], generalized expansion dimension [Houle et al., 2012], and LID [Houle, 2013]. Of these, LID [Houle, 2013] has been of particular interest to the machine learning community for detecting adversarial data points (Section 2.2). Therefore, in this paper, we focus on the adversarial detection capabilities of LID.

We start by discussing the theoretical background behind LID. First, we briefly introduce the theory of LID for assessing the dimensionality of data subspaces. Expansion models of dimensionality have previously been successfully employed in a wide range of applications such as manifold learning, dimension reduction, similarity search, and anomaly detection [Houle, 2017a, Amsaleg et al., 2015]. Next, we analyze the effect of perturbations on the LID values of data points.

3.1   Theory of Local Intrinsic Dimensionality.

In the theory of intrinsic dimensionality, classical expansion models measure the rate of growth in the number of data samples encountered as the distance from the sample of interest increases [Houle, 2017a]. As an example, in Euclidean space, the volume of an m-dimensional hyper-sphere grows proportionally to r^m, when its size is scaled by a factor of r. If we consider two hyper-spheres sharing the same center with volumes V_1 and V_2 and radii r_1 and r_2, respectively, the expansion dimension m can be deduced as:

$$\frac{V_2}{V_1} = \left(\frac{r_2}{r_1}\right)^m \Rightarrow m = \frac{\ln\left(V_2/V_1\right)}{\ln\left(r_2/r_1\right)}. \tag{1}$$

Transferring the concept of expansion dimension to the statistical setting of continuous distance distributions leads to the formal definition of LID. By substituting the cumulative probability of distance for volume, LID provides measures

of the intrinsic dimensionality of the underlying data subspace. Refer to the work of Houle [2017a] for more details concerning the theory of LID. The formal definition of LID is given below.

Definition 1 (Local Intrinsic Dimensionality [Houle, 2017a])**.** *Given a data sample* $\mathbf{a} \in \mathcal{S}$*, let* $R > 0$ *be a random variable denoting the distance from* $\mathbf{a}$ *to other data samples. If the cumulative probability distribution function* $F(r)$ *of* R *is positive and continuously differentiable at distance* $r > 0$*, the intrinsic dimensionality of* $\mathbf{a}$ *at distance* r *is given by:*

$$ID_F(r) \triangleq \lim_{\epsilon \to 0^+} \frac{ln\big(F((1+\epsilon) \cdot r)/F(r)\big)}{ln(1+\epsilon)} = \frac{r \cdot F'(r)}{F(r)}, \tag{2}$$

whenever the limit exists.

The last equality of (2) follows by applying L'Hôpital's rule to the limits [Houle, 2017a]. The local intrinsic dimension at $\mathbf{a}$ is in turn defined as the limit when the radius r tends to zero:

$$\text{LID}(\mathbf{a}) = \lim_{r \to 0} ID_F(r). \tag{3}$$

$\text{LID}(\mathbf{a})$ measures the probability density $F'(r)$ normalized by the cumulative probability density $F(r)/r$. In the ideal case where the data in the vicinity of $\mathbf{a}$ is distributed uniformly within a subspace, $\text{LID}(\mathbf{a})$ equals the dimension of the subspace; however, in practice these distributions are not ideal, the manifold model of data does not perfectly apply, and $\text{LID}(\mathbf{a})$ is not an integer [Ma et al., 2018b]. Nevertheless, the local intrinsic dimensionality is an indicator of the dimension of the subspace containing $\mathbf{a}$ that would best fit the data distribution in the vicinity of $\mathbf{a}$.

The smallest distances from point $\mathbf{a}$ can be regarded as extreme events associated with the lower tail of the underlying distance distribution. Using extreme value theory (EVT), the following theorem shows that $F(r)$ completely determines ID_F.

Theorem 1 (Local ID Representation [Houle, 2017c])**.** *Let* $F : \mathbb{R}^{\geq 0} \to \mathbb{R}$ *be a real-valued function, and assume that the* $ID_F(0)$ *exists. Let* r *and* w *be values for which* r/w *and* $F(r)/F(w)$ *are both positive. If* F *is non-zero and continuously differentiable everywhere in the interval containing* $[min\{r, w\}, max\{r, w\}]$*, then*

$$\frac{F(r)}{F(w)} = \left(\frac{r}{w}\right)^{ID_F(0)} \cdot G_{F,w}(r), where$$

$$G_{F,w}(r) \triangleq \exp\left(\int_r^w \frac{ID_F(0) - ID_F(t)}{t} dt\right), \tag{4}$$

whenever the integral exists.

Furthermore, let $e > 1$ be a constant. Then, as per Theorem 3 of Houle [2017c],

$$\lim_{\substack{w \to 0^+ \\ 0 < 1/e \leq r/w \leq e}} G_{F,w}(r) = 1. \tag{5}$$

The LID formula defined in (3) is the theoretical calculation of LID. We describe below the empirical estimation of LID ($\widehat{\text{LID}}$) as follows.

Definition 2 (Estimation of LID [Amsaleg et al., 2015])**.** *Given a reference sample* $\mathbf{a} \sim \mathcal{P}$*, where* $\mathcal{P}$ *represents the data distribution, the Maximum Likelihood Estimator of the LID at* $\mathbf{a}$ *is defined as follows:*

$$\widehat{LID}(\mathbf{a}) = -\left(\frac{1}{k} \sum_{i=1}^{k} log \frac{r_i(\mathbf{a})}{r_{max}(\mathbf{a})}\right)^{-1}. \tag{6}$$

Here, $r_i(\mathbf{a})$ *denotes the distance between* $\mathbf{a}$ *and its* i*-th nearest neighbor within a sample of* k *points drawn from* $\mathcal{P}$*, and* $r_{max}(\mathbf{a})$ *is the maximum of the neighbor distances.*

The above estimation assumes that samples are drawn from a tight neighborhood, in line with its development from extreme value theory. In practice, the sample set is drawn uniformly from the available training data (omitting x itself), which itself is presumed to have been randomly drawn from $\mathcal{P}$.

4   Adversarial Perturbations and LID

As stated in Section 2, multiple prior works use defenses built around LID for detecting and suppressing the impact of adversarial samples. In the literature, adversarial samples are usually generated by perturbing benign data in directions provided by the attack algorithm [Goodfellow et al., 2015, Moosavi-Dezfooli et al., 2016, Maini et al., 2020]. By perturbing the feature vectors, the adversary moves poisoned samples away from the domain of benign samples. Thus, each poisoned sample would have an irregular distribution of the local distance to its neighbors, which would be reflected by its LID value. For example, the LID estimate of the data point $\mathbf{a} \in \mathcal{S}$, i.e., $\widehat{\mathrm{LID}}(\mathbf{a})$, is an indicator of the dimension of the subspace that contains $\mathbf{a}$, by comparing the LID estimate of a data sample to that of other data points, any data points that have substantially different lower-dimensional subspaces would get highlighted.

Table 1: Summary of the Notation.

|	
$\mathcal{S}$	The domain of data samples.
$\mathbf{a}$	The benign data point being perturbed ($\mathbf{a} \in \mathcal{S}$).
$\mathbf{b}$	The position of $\mathbf{a}$, after the perturbation ($\mathbf{b} \in \mathcal{S}$).
$\mathbf{c}$	A benign reference point ($\mathbf{c} \in \mathcal{S}$).
$d(\mathbf{a}, \mathbf{b})$	A distance measure that measures the distance between any two data points in $\mathcal{S}$.
x	The distance between $\mathbf{a}$ and $\mathbf{c}$ measured using d.
y	The distance between $\mathbf{b}$ and $\mathbf{c}$ measured using d.
δ	A real valued constant that determines the size of perturbations.
δx	The distance between $\mathbf{b}$ and $\mathbf{a}$ measured using d.
$\mathcal{D}$	The univariate distribution of the distances to other data points measured from any data point in $\mathcal{S}$.
$F(r)$	The cumulative probability distribution function of $\mathcal{D}$.
$\mathrm{LID}(\mathbf{a})$	The local intrinsic dimensionality of $\mathbf{a}$.
$\widehat{\mathrm{LID}}(\mathbf{a})$	The numerical estimation of $\mathrm{LID}(\mathbf{a})$.
n	The number of samples (data points) in $\mathcal{S}$ (i.e., $|\mathcal{S}|$).
$p_{\mathbf{c}}$	The cumulative probability of $F_{\mathbf{b}}$ at distance y (i.e., $F_{\mathbf{b}}(y)$).
$k_{\mathbf{c}}$	The expected rank of $\mathbf{c}$ w.r.t the distribution of distances from $\mathbf{b}$ (i.e., $p_{\mathbf{c}} \times n$).
θ	The angle between $\mathbf{b} - \mathbf{a}$ and $\mathbf{c} - \mathbf{a}$.

4.1   Problem Setup and Notation
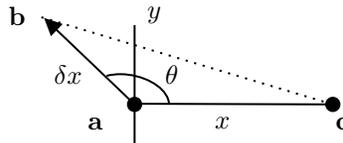
Figure 1: The positions of the benign sample $\mathbf{a}$, its perturbed version $\mathbf{b}$, and reference point $\mathbf{c}$.

We start by describing the problem setup using three data points $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{c}$ as depicted in Figure 1. Take $\mathbf{a} \in \mathcal{S}$ as the initial (benign) point that the attacker aims to perturb and $\mathbf{b} \in \mathcal{S}$ as the resulting perturbed point. Take $\mathbf{c} \in \mathcal{S}$ as some benign reference location. In this section, we analyze how the perturbations affect $\mathrm{LID}(\mathbf{b})$ (i.e., the LID value of the perturbed data point). Define $d(\mathbf{a}, \mathbf{b})$ as the distance measure that measures the distance between any two data points in $\mathcal{S}$. Although we use Euclidean distance as the measure of distance in this paper, we note that any measure of distance for which the definition of LID holds is feasible. Thus, $\mathcal{D}_{\mathbf{b}}$ is defined as a univariate distribution of the distances to other data points measured from $\mathbf{b}$.

Any point, for example $\mathbf{c}$, determines a distance $y = d(\mathbf{c}, \mathbf{b}) = \|\mathbf{b} - \mathbf{c}\|$, which in turn determines a probability $p_{\mathbf{c}} = F_{\mathbf{b}}(y)$. That is, any sample drawn from $\mathcal{S}$ has probability $p_{\mathbf{c}}$ being closer to $\mathbf{b}$. Therefore, with respect to the distribution $\mathcal{D}_{\mathbf{b}}$, the point $\mathbf{c}$ is given a distributional rank $p_{\mathbf{c}}$. We define the expected rank of $\mathbf{c}$ as $k_{\mathbf{c}} := p_{\mathbf{c}} \times n$, where $n = |\mathcal{S}|$.
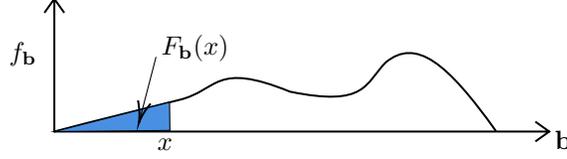
Figure 2: The probability distribution of distances from $\mathbf{b} \in \mathcal{S}$ induced by $\mathcal{D}_{\mathbf{b}}$.

Lemma 1 considers a point $\mathbf{a}$ at distance x from $\mathbf{c}$ (i.e., $x = \|\mathbf{a} - \mathbf{c}\|$). The point $\mathbf{b}$, which is produced by perturbing $\mathbf{a}$, has distance δx from $\mathbf{a}$ for some proportion $\delta > 0$ (see Fig. 1). Define θ as the angle between $\mathbf{b} - \mathbf{a}$ and $\mathbf{c} - \mathbf{a}$. The lemma gives sufficient conditions on δ and θ for $\mathbf{b}$ to be closer to (resp. farther from) $\mathbf{c}$ compared to $\mathbf{a}$.

Lemma 1. *Define the distance $y = \|\mathbf{b} - \mathbf{c}\|$ and the associated cumulative probability $p_c = F_{\mathbf{b}}(y)$. Similarly, define the distance $\delta x = \|\mathbf{b} - \mathbf{a}\|$ and the associated cumulative probability $p_{\mathbf{a}} = F_{\mathbf{b}}(\delta x)$ where $\delta > 0$ and $x = \|\mathbf{a} - \mathbf{c}\|$. For $y > x$, that is when $\mathbf{b}$ is perturbed farther from $\mathbf{c}$ (relative to $\mathbf{a}$), $\delta > 2\cos\theta$. For $y < x$, (i.e., perturbed towards $\mathbf{c}$) $\delta < 2\cos\theta$ and $-\frac{\pi}{2} < \theta < \frac{\pi}{2}$.*

Proof. Since θ is the angle between $\mathbf{b} - \mathbf{a}$ and $\mathbf{c} - \mathbf{a}$, by construction, we have:

$$
\begin{aligned}
\|\mathbf{b} - \mathbf{c}\|^2 &= \|\mathbf{b} - \mathbf{a} + \mathbf{a} - \mathbf{c}\|^2, \\
&= \|\mathbf{b} - \mathbf{a}\|^2 + \|\mathbf{c} - \mathbf{a}\|^2 - 2(\mathbf{b} - \mathbf{a}).(\mathbf{c} - \mathbf{a}), \\
&= \delta^2 x^2 + x^2 - 2\delta x^2 \cos\theta.
\end{aligned}
\tag{7}
$$

In the first case where $y > x$ ($\mathbf{b}$ is perturbed farther from $\mathbf{c}$) we have,

$$
\begin{aligned}
\|\mathbf{b} - \mathbf{c}\|^2 &> \|\mathbf{a} - \mathbf{c}\|^2, \\
\delta^2 x^2 + x^2 - 2\delta x^2 \cos\theta &> x^2, \\
\delta x^2 (\delta - 2\cos\theta) &> 0.
\end{aligned}
\tag{8}
$$

Since $\delta > 0$ and $x > 0$, we have

$$
\delta > 2\cos\theta.
\tag{9}
$$

Following a similar approach, for the case where $y < x$ ($\mathbf{b}$ is perturbed towards $\mathbf{c}$) we obtain:

$$
\delta < 2\cos\theta \text{ and } -\frac{\pi}{2} < \theta < \frac{\pi}{2}.
\tag{10}
$$

$\square$

5 Theoretical Bounds on LID

We now consider the relationship between local intrinsic dimensionality and the effect of perturbations on neighborhoods under appropriate smoothness conditions of the underlying distance distribution. First, we define the conditions under which the LID is continuous at $\mathbf{b} \in \mathcal{S}$ [Amsaleg et al., 2020]:

1. There exists a distance $y > 0$ for which all points $\mathbf{c} \in \mathbb{R}^d$ with $\|\mathbf{c} - \mathbf{b}\| \leq y$ admit a distance distribution $\mathcal{D}_{\mathbf{b}}$ whose cumulative distribution function $F_{\mathbf{b}}$ is continuously differentiable and positive within some open interval with a lower bound of 0.

2. For any sequence $\mathbf{b} \to \mathbf{c}$ of points satisfying Condition 1, there is convergence in distribution of the sequence of random distance variables defined at $\mathbf{b}$ to the distance variable defined at $\mathbf{c}$; that is, the condition $\lim_{\mathbf{b} \to \mathbf{c}} F_{\mathbf{b}}(\epsilon) = F_{\mathbf{c}}(\epsilon)$ holds for any distance $\epsilon \in (0, y)$.

3. For each $\mathbf{b}$ satisfying Condition 1, $\mathrm{LID}(\mathbf{b})$ exists and is positive.

4. $\lim_{\mathbf{b} \to \mathbf{c}} \mathrm{LID}(\mathbf{b}) = \mathrm{LID}(\mathbf{c})$.

For the remainder of this Section, we assume that the LID is continuous at point $\mathbf{b} \in \mathcal{S}$. In practice, we find that these technical conditions are often satisfied as illustrated by the experimental results in Section 6.

The following theorem derives an upper-bound and a lower bound for the LID of a perturbed data point, $\mathbf{b}$, based on the distances to its original location, $\mathbf{a}$, and a second reference data point $\mathbf{c}$, which are shown in Fig. 1.

Theorem 2. *Define $y = \|\mathbf{b} - \mathbf{c}\|$, $x = \|\mathbf{a} - \mathbf{c}\|$ and $\delta x = \|\mathbf{b} - \mathbf{a}\|$. Let $F_{\mathbf{b}}(\delta x)$ be the cumulative probability of the distance distribution $\mathcal{D}_{\mathbf{b}}$ at distance δx. Similarly, let $F_{\mathbf{b}}(y)$ be the cumulative probability of the distance distribution $\mathcal{D}_{\mathbf{b}}$ at distance y. Thus, $p_{\mathbf{a}} = F_{\mathbf{b}}(\delta x)$ and $p_{\mathbf{c}} = F_{\mathbf{b}}(y)$ are the distributional ranks of $\mathbf{a}$ and $\mathbf{c}$ w.r.t the distribution of distances from $\mathbf{b} \in \mathcal{S}$. Similarly, define $k_{\mathbf{a}} = np_{\mathbf{a}}$ and $k_{\mathbf{c}} = np_{\mathbf{c}}$ as the corresponding expected ranks. Let δ be a real constant that is bounded as $0 < \delta < y/x$. Assume LID($\mathbf{b}$) exists and is positive. Let η be a sufficiently small real value that is bounded as $0 < \eta < \min\left(\frac{y}{\delta x} - 1, \text{LID}(\mathbf{b}) \ln(\varphi)/|\ln(\frac{\delta x}{y})|\right)$, where $\varphi = min\{\frac{y+\delta x \eta}{y}, \frac{y}{y - \delta x \eta}\}$. For feasible values of η, there exists a positive integer $n_0 > \max\{k_{\mathbf{a}}, k_{\mathbf{c}}\}$ for which the following inequalities hold for all choices of integer $n \geq n_0$,*

$$\frac{\ln\left(F_{\mathbf{b}}(y)/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} + \eta)} \leq LID(\mathbf{b}) \leq \frac{\ln\left(F_{\mathbf{b}}(y)/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)}. \tag{11}$$

Proof. For a given choice of n, with $\mathbf{a}, \mathbf{b}, \mathbf{c}, x$ and δ as defined above, we have $p_{\mathbf{a}} = k_{\mathbf{a}}/n$ and $p_{\mathbf{c}} = k_{\mathbf{c}}/n$. Using the local ID characterization formula of the representation theorem (i.e., Theorem 4), we observe that

$$\frac{k_{\mathbf{c}}}{k_{\mathbf{a}}} = \frac{p_{\mathbf{c}}}{p_{\mathbf{a}}} = \frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)} = \left(\frac{y}{\delta x}\right)^{\text{LID}(\mathbf{b})} \cdot G_{F_{\mathbf{b}, \delta x}}(y), \tag{12}$$

where $G_{F_{\mathbf{b}, \delta x}}(r) \triangleq \exp\left(\int_y^{\delta x} \frac{\text{LID}(\mathbf{b}) - \text{ID}_{F_{\mathbf{b}}}(u)}{u} du\right)$. Upon rearranging, we obtain

$$\frac{y}{\delta x} = \left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x) \cdot G_{F_{\mathbf{b}, \delta x}}(y)}\right)^{1/\text{LID}(\mathbf{b})}. \tag{13}$$

Following a logarithmic transformation and substituting for $G_{F_{\mathbf{b}, \delta x}}(y)$, we have

$$\ln\left(\frac{y}{\delta x}\right) = \frac{1}{\text{LID}(\mathbf{b})} \ln\left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)}\right) - \frac{1}{\text{LID}(\mathbf{b})} \int_y^{\delta x} \frac{\text{LID}(\mathbf{b}) - \text{ID}_{F_{\mathbf{b}}}(u)}{u} du. \tag{14}$$

It is assumed that $F_{\mathbf{b}}(y) \neq F_{\delta \mathbf{x}}(y)$, which implies $y \neq \delta x$. Therefore, as per Houle [2017c], for some real value η such that $0 < \eta < \frac{y}{\delta x} - 1$, there must exist a sufficiently small distance value $0 < w$ such that any distance u where $0 < u < w$, implies $|\text{LID}(\mathbf{b}) - \text{ID}_{F_{\mathbf{b}}(u)}| < \eta$. Therefore, the integral in (14) is bounded above by

$$\left|\int_y^{\delta x} \frac{\text{LID}(\mathbf{b}) - \text{ID}_{F_{\mathbf{b}}}(u)}{u} du\right| \leq \eta \cdot \left|\int_y^{\delta x} \frac{1}{u} du\right|. \tag{15}$$

Furthermore, since η is defined as a small positive value, we bound η from above as

$$\eta \leq \frac{\text{LID}(\mathbf{b}) \ln(\varphi)}{|\ln(\frac{\delta x}{y})|}, \tag{16}$$

where φ is the real valued variable

$$\varphi = \min\left\{\frac{y + \delta x \eta}{y}, \frac{y}{y - \delta x \eta}\right\} > 1. \tag{17}$$

Define n_0 as the minimum size of the data set for which F_b is strictly positive and continuously differentiable within some open interval of distances with lower endpoint 0. For choices of η as defined above, whenever the size of the data

set n is beyond the minimum ($n \geq n_0$), we have,

$$\left| \ln\left(\frac{y}{\delta x}\right) - \frac{1}{\text{LID}(\mathbf{b})} \ln\left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)}\right) \right| \leq \frac{1}{\text{LID}(\mathbf{b})} \left| \int_y^{\delta x} \frac{\text{LID}(\mathbf{b}) - \text{ID}_{F_{\mathbf{b}}}(u)}{u} \, du \right|,$$

$$\leq \frac{\eta}{\text{LID}(\mathbf{b})} \left| \int_y^{\delta x} \frac{1}{u} \, du \right|,$$

$$= \frac{\eta}{\text{LID}(\mathbf{b})} \left| \ln\left(\frac{\delta x}{y}\right) \right|,$$

$$\leq \frac{\ln\varphi}{\left| \ln\left(\frac{\delta x}{y}\right) \right|} \left| \ln\left(\frac{\delta x}{y}\right) \right|,$$

$$= \ln\varphi.$$

(18)

By rearranging (18), the we obtain

$$\ln\left(\frac{y}{\delta x \varphi}\right) \leq \frac{1}{\text{LID}(\mathbf{b})} \ln\left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)}\right) \leq \ln\left(\frac{y\varphi}{\delta x}\right)$$

$$\ln\left(\frac{y}{\delta x} \cdot \frac{y - \delta x \eta}{y}\right) \leq \frac{1}{\text{LID}(\mathbf{b})} \ln\left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)}\right) \leq \ln\left(\frac{y}{\delta x} \cdot \frac{y + \delta x \eta}{y}\right)$$

$$\ln\left(\frac{y}{\delta x} - \eta\right) \leq \frac{1}{\text{LID}(\mathbf{b})} \ln\left(\frac{F_{\mathbf{b}}(y)}{F_{\mathbf{b}}(\delta x)}\right) \leq \ln\left(\frac{y}{\delta x} + \eta\right).$$

(19)

Upon rearranging, we obtain

$$\frac{\ln\left(F_{\mathbf{b}}(y)/F_{\mathbf{b}}(\delta x)\right)}{\ln\left(\frac{y}{\delta x} + \eta\right)} \leq \text{LID}(\mathbf{b}) \leq \frac{\ln\left(F_{\mathbf{b}}(y)/F_{\mathbf{b}}(\delta x)\right)}{\ln\left(\frac{y}{\delta x} - \eta\right)}.$$

(20)

□

Theorem 2 derives a lower-bound and an upper-bound for $\text{LID}(\mathbf{b})$ using the cumulative probability values (of the distribution of distances from $\mathbf{b}$) of the remaining two data points and the distances to them. Since the maximum value $F_{\mathbf{b}}(\delta x)$ can have is one, we observe a strong association between the size of perturbations δ and the LID value of $\mathbf{b}$. This theoretical result justifies the use of LID as an indicator to detect adversarial perturbations in prior works (Section 2).

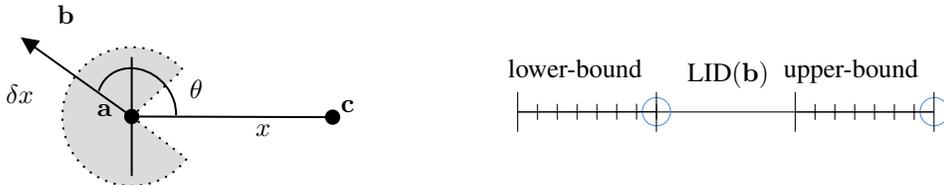## 5.1 Additional Bounds based on the Perturbation Directions

Figure 3: The figure on the left shows the valid range of θ (grey) when $\mathbf{b}$ is perturbed away from $\mathbf{c}$. The figure on the right shows the least lower and upper bounds when $\mathbf{b}$ is perturbed away from $\mathbf{c}$.

We now consider the relationship between LID and the direction of perturbations. The bounds in Theorem 2 hold for $\mathbf{b}$ irrespective of the perturbation directions. In Lemma 1, we present the conditions δ and θ must satisfy given the two possible directions of perturbation. In the following theorem, we fix δ, and consider the possible angles θ can take to obtain the greatest lower-bound and greatest lower-bound of (11) when $\mathbf{b}$ is perturbed away from $\mathbf{c}$ (relative to $\mathbf{a}$) (refer Figure 3). Conversely, we obtain the least lower-bound and least upper-bound of (11) when $\mathbf{b}$ is perturbed towards $\mathbf{c}$ (refer Figure 4).
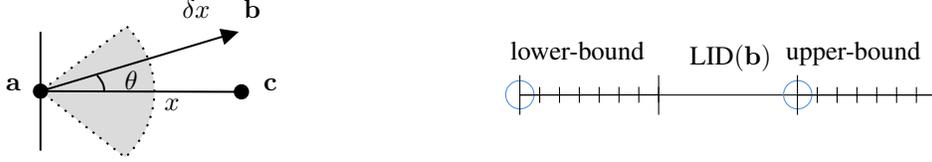
Figure 4: The figure on the left shows the valid range of θ (grey) when $\mathbf{b}$ is perturbed towards $\mathbf{c}$. The figure on the right shows the least lower and upper bounds when $\mathbf{b}$ is perturbed towards $\mathbf{c}$.

Theorem 3. *Denote the distances $x = \|\mathbf{a} - \mathbf{c}\|$, $\|\mathbf{b} - \mathbf{c}\|$ as y and $\delta x = \|\mathbf{a} - \mathbf{b}\|$ for an appropriate choice of $0 < \delta < y/x$. We denote the distributional ranks of $\mathbf{a}$ and $\mathbf{c}$ as $p_\mathbf{a} = F_\mathbf{b}(\delta x)$ and $p_\mathbf{c} = F_\mathbf{b}(y)$, and the distribution of distances from $\mathbf{b} \in \mathcal{S}$ as $\mathcal{D}_\mathbf{b}$. Assume LID($\mathbf{b}$) exists and is positive. Assume η is a sufficiently small real value that satisfies the constraints outlined in Theorem 2. Additionally, it satisfies the constraint $1 - \frac{1}{\delta} < \eta < \frac{1}{\delta} - 1$.*

By considering the boundary values of $\cos\theta$ as per Lemma 1, when $\mathbf{b}$ is perturbed farther from $\mathbf{c}$ (i.e., $y > x$), the greatest lower and upper bounds of (20) *are as follows:*

$$\frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{y}{\delta x} + \eta)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < LID(\mathbf{b}) \le$$

$$\frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)}. \tag{21}$$

Conversely, when $\mathbf{b}$ is perturbed towards $\mathbf{c}$ (i.e., $y < x$), the least lower and upper bounds of (20) *are as follows:*

$$\frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{y}{\delta x} + \eta)} \le LID(\mathbf{b}) <$$

$$\frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)}. \tag{22}$$

Proof. We consider the two possible directions of perturbation separately.
Case 1: $\mathbf{b}$ perturbed farther from $\mathbf{c}$
First, we consider the case where $\mathbf{b}$ is perturbed farther from $\mathbf{c}$ (i.e., $y > x$). Using (7) on the denominator of the left inequality of (20), we obtain

$$\ln\left(\frac{y}{\delta x} + \eta\right) = \ln\left(\frac{\sqrt{\delta^2 x^2 + x^2 - 2\delta x^2 \cos\theta}}{\delta x} + \eta\right),$$

$$= \ln\left(\sqrt{1 + \frac{1}{\delta^2} - \frac{2\cos\theta}{\delta}} + \eta\right). \tag{23}$$

We then consider the upper boundary value of $\cos\theta$ as shown in (9) (i.e., $\cos\theta < \delta/2$), where we get

$$\ln\left(\sqrt{1 + \frac{1}{\delta^2} - \frac{2\cos\theta}{\delta}} + \eta\right) > \ln\left(\sqrt{1 + \frac{1}{\delta^2} - 1} + \eta\right)$$

$$= \ln\left(\frac{1}{\delta} + \eta\right). \tag{24}$$

Thus, for all feasible values of θ, we observe that $\ln(\frac{y}{\delta x} + \eta) > \ln(\frac{1}{\delta} + \eta)$. Therefore, by replacing $\ln(\frac{y}{\delta x} + \eta)$ in (20) with $\ln(\frac{1}{\delta} + \eta)$, we obtain the following greatest lower bound for LID($\mathbf{b}$)

$$\frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln\left(\frac{y}{\delta x} + \eta\right)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < LID(\mathbf{b}). \tag{25}$$

The greatest upper bound of (21) can be obtained by replacing $\ln\left(\frac{y}{\delta x} - \eta\right)$ with $\ln(\frac{1}{\delta} - \eta)$ following an identical approach.

$$LID(\mathbf{b}) \le \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)} < \frac{\ln\left(F_\mathbf{b}(y)/F_\mathbf{b}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)}. \tag{26}$$

9

Thus, from (25) and (26), we see that when $\mathbf{b}$ is perturbed farther from the reference sample $\mathbf{c}$, the greatest lower and upper bounds of (20) are as follows:

$$\frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} + \eta)} < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < \text{LID}(\mathbf{b}) \leq$$
$$\frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)} < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)}. \tag{27}$$

Case 2: b perturbed closer to c

We now consider the case where $\mathbf{b}$ is perturbed towards $\mathbf{c}$ (i.e., $y < x$). Take the lower boundary value of $\cos\theta$ as shown in (10) (i.e., $\cos\theta > \delta/2$) in order to obtain the greatest possible denominator value as follows:

$$\begin{aligned}
\ln\left(\frac{y}{\delta x} + \eta\right) &= \ln\left(\sqrt{1 + \frac{1}{\delta^2} - \frac{2\cos\theta}{\delta}} + \eta\right) \\
&< \ln\left(\sqrt{1 + \frac{1}{\delta^2} - 1} + \eta\right) \\
&= \ln\left(\frac{1}{\delta} + \eta\right).
\end{aligned} \tag{28}$$

Thus, for all feasible values of θ, we have $\ln(\frac{y}{\delta x} + \eta) < \ln(\frac{1}{\delta} + \eta)$. By replacing $\ln(\frac{y}{\delta x} + \eta)$ in (22) with $\ln(\frac{1}{\delta} + \eta)$, we obtain the following least lower bound for $\text{LID}(\mathbf{b})$

$$\frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln\left(\frac{y}{\delta x} + \eta\right)} \leq \text{LID}(\mathbf{b}). \tag{29}$$

The least upper bound of (21) can be obtained by replacing $\ln\left(\frac{y}{\delta x} - \eta\right)$ with $\ln(\frac{1}{\delta} - \eta)$ following an identical approach.

$$\text{LID}(\mathbf{b}) < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)} < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)}. \tag{30}$$

Therefore, from (29) and (30), we observe the least lower and upper bounds of (20) when $\mathbf{b}$ is perturbed towards from the reference sample $\mathbf{c}$.

$$\frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} + \eta)} < \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} + \eta)} \leq \text{LID}(\mathbf{b}) <$$
$$\frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{1}{\delta} - \eta)} \leq \frac{\ln\left(F_{\mathbf{b}}(y)\big/F_{\mathbf{b}}(\delta x)\right)}{\ln(\frac{y}{\delta x} - \eta)}. \tag{31}$$

$\square$

We note that in practice, as evidenced by the experimental results in the section that follows, attack algorithms tend to perturb data away from the domain of the benign data. Therefore, the most important result of Theorem 3 is the greatest lower-bound in (27). This greatest-lower bound further supports the strong association seen between the size of perturbations δ and the LID value of $\mathbf{b}$.

6 Experimental Validation and Discussion

In this section, we empirically evaluate the lower-bound and upper-bound derived in Theorem 2 for the LID value of the perturbed data point $\mathbf{b}$. Our code is available at `https://github.com/sandamal/adversarial_lid_bounds`.

To approximate the cumulative distributions of the distance variable (i.e., F_b), we first use kernel density estimation with a Gaussian kernel to obtain the probability density function of $\mathcal{D}_b$ [Silverman, 2018]. Subsequently, the integral is calculated to obtain the cumulative distribution. We empirically evaluate the theoretical bounds on two real-world datasets: MNIST [LeCun et al., 2010] and CIFAR-10 [Krizhevsky, 2009]. The datasets contain 10,000 samples with 782 and 3,072 features respectively. It should be noted that, similar to the work of Amsaleg et al. [2020], the bounds hold asymptotically as the number of data samples n tends to infinity. However, due to the high computational complexity of kernel density estimation, we perform our experiments on a randomly sampled subset of data points with $n \in \{1000, 5000\}$.

(a) MNIST: $k = 100$, $n_q = 50$, $n = 1000$.

(b) CIFAR-10: $k = 100$, $n_q = 50$, $n = 1000$.

(c) MNIST: $k = 1000$, $n_q = 50$, $n = 5000$.
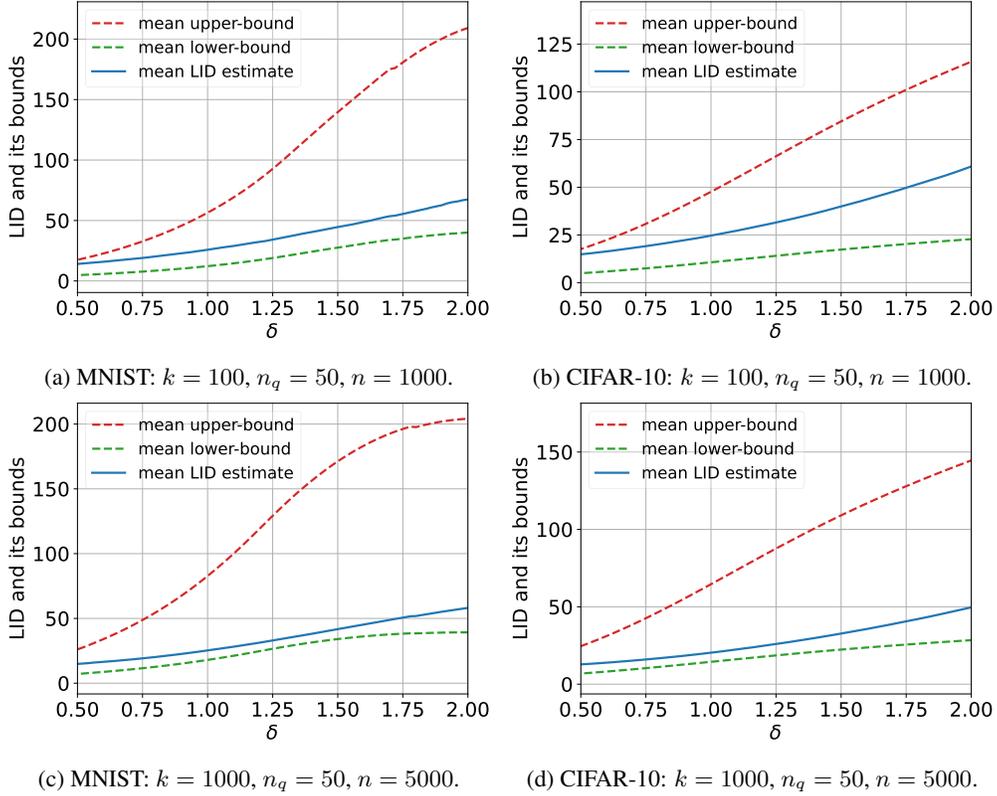
(d) CIFAR-10: $k = 1000$, $n_q = 50$, $n = 5000$.

Figure 5: The LID estimates and theoretical bounds on MNIST and CIFAR-10 when data points are poisoned using the gradient-based attack.

We randomly select $n_q = 50$ query points from each dataset (i.e., $\mathbf{a}$) and select the perturbation direction using a gradient based attack [Goodfellow et al., 2015]. For the LID estimations (6) we select the neighborhood size k from $\{100, 1000\}$. For each query point $\mathbf{a}$, we select the data point with the rank $k/2$ from it as the reference point (i.e., $\mathbf{c}$). We then obtain the perturbed points (i.e., $\mathbf{b}$) by perturbing the query points in the direction that maximizes the training error (according to the attack algorithm).

Figure 5 shows the empirical evaluations of the theoretical bounds in Theorem 2. It shows the averaged bounds for all feasible η values as well as the estimated LID values against the perturbation δ. As depicted in the figure, when δ increases, the lower-bound, LID estimate and the upper-bound increase. The most important indicator here is the lower-bound that increases with δ as it supports the intuition that LID values of adversarial samples tend to be larger.

It should be noted that when δ is close to zero, the bounds in (11) exhibit some numerical instability as $F_{\mathbf{b}}(y)/F_{\mathbf{b}}(\delta x)$, $\frac{y}{\delta x} + \eta$ and $\frac{y}{\delta x} - \eta$ tend to infinity. However, in practice, it is expected for adversarial perturbations to be sufficiently large in order to influence the learner's prediction capabilities. Therefore it is reasonable to expect δ to not be close to zero.

We observe similar numerical instabilities for (27) and (31) of Theorem 3 as well. Furthermore, when δ tends to one (i.e., the theoretical maximum value for which Theorem 3 theoretically holds), since η is a small positive value, we observe that both lower-bounds and upper-bounds in (27) and (31) tend to infinity. Although we have considered δ values up to 2.5 (which makes the maximum perturbation $2.5x$), in practice, we observe that adversarial attacks tend to have bounded perturbations (to avoid detection).

As for the directions of perturbation, we observed that, on average, only 5.56% and 0.65% of the total perturbations were made in the direction towards the benign reference point $\mathbf{c}$ for MNIST and CIFAR-10 respectively. Intuitively, we expect a gradient-based attack that perturbs data in the direction that maximizes the training error to perturb data points away from the benign data domain. This finding, while preliminary, supports this intuition.

7 Conclusions

The purpose of this paper was to theoretically determine the relationship between the magnitude of the perturbations and the LID value of the resulting perturbed data point. To that end, we derived a lower-bound and an upper-bound for the LID value of a perturbed data point based on the cumulative probability (of its distance distribution) values to its original position and a third benign reference point. The bounds, in particular the lower-bound, showed that there is a positive association between the magnitude of the perturbation δ and the LID value of the perturbed point. We then empirically evaluated the effectiveness of these bounds on two reference data sets using a gradient-based poisoning attack. Furthermore, by considering the two possible directions of perturbation, we obtained the least lower-bound and upper-bound when the direction is towards the benign reference point and the greatest lower-bound and upper-bound when the direction is away from it. These findings support the use of LID as a tool to detect adversarial samples in prior state of the art works.

References

Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pages 5286–5295. PMLR, 2018.

Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310–1320. PMLR, 2019.

Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 656–672. IEEE, 2019.

Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *37th International Conference on Machine Learning*, pages 11278–11287, 2020.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019.

Michael E. Houle. Local intrinsic dimensionality I: an extreme-value-theoretic foundation for similarity applications. In Christian Beecks, Felix Borutta, Peer Kröger, and Thomas Seidl, editors, *Similarity Search and Applications*, pages 64–79, 2017a.

Michael E. Houle, Erich Schubert, and Arthur Zimek. On the correlation between local intrinsic dimensionality and outlierness. In *International Conference on Similarity Search and Applications*, pages 177–191. Springer, 2018.

Xingjun Ma, Yisen Wang, Michael E. Houle, Shuo Zhou, Sarah Erfani, Shutao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3355–3364, 2018a.

Sandamal Weerasinghe, Tamas Abraham, Tansu Alpcan, Sarah M. Erfani, Christopher Leckie, and Benjamin I. P. Rubinstein. Closing the big-lid: An effective local intrinsic dimensionality defense for nonlinear regression poisoning. In *30th International Joint Conferences on Artificial Intelligence*, 2021a.

Michael E. Houle, Hisashi Kashima, and Michael Nett. Generalized expansion dimension. In *2012 IEEE 12th International Conference on Data Mining Workshops*, pages 587–594, 2012.

Michael E. Houle. Local intrinsic dimensionality II: multivariate analysis and distributional support. In *International Conference on Similarity Search and Applications*, pages 80–95, 2017b.

Laurent Amsaleg, Oussama Chelly, Teddy Furon, Stephane Girard, Michael E. Houle, Ken-Ichi Kawarabayashi, and Michael Nett. Estimating local intrinsic dimensionality. In *21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 29–38. ACM, 2015.

Laurent Amsaleg, Oussama Chelly, Michael E Houle, Ken-Ichi Kawarabayashi, Miloš Radovanović, and Weeris Treeratanajaru. Intrinsic dimensionality estimation within tight localities. In *SIAM International Conference on Data Mining*, pages 181–189, 2019.

Xingjun Ma, Bo Li, Yisen Wang, Sarah M. Erfani, Sudanthi N. R. Wijewickrema, Grant Schoenebeck, Dawn Song, Michael E. Houle, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. In *6th International Conference on Learning Representations, ICLR 2018*, 2018b.

Sandamal Weerasinghe, Tansu Alpcan, Sarah M. Erfani, and Christopher Leckie. Defending support vector machines against data poisoning attacks. *IEEE Transactions on Information Forensics and Security*, 16:2566–2578, 2021b.

Sandamal Weerasinghe, Tamas Abraham, Tansu Alpcan, Sarah M. Erfani, Christopher Leckie, and Benjamin I. P. Rubinstein. Closing the big-lid: An effective local intrinsic dimensionality defense for nonlinear regression poisoning. In *Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3176–3184, 2021c.

Sukarna Barua, Xingjun Ma, Sarah M. Erfani, Michael E. Houle, and James Bailey. Quality evaluation of gans using cross local intrinsic dimensionality. *CoRR*, abs/1905.00643:1–31, 5 2019.

Laurent Amsaleg, James Bailey, Amelie Barbe, Sarah M. Erfani, Teddy Furon, Michael E Houle, Miloš Radovanović, and Xuan Vinh Nguyen. High intrinsic dimensionality facilitates adversarial attack: Theoretical evidence. *IEEE Transactions on Information Forensics and Security*, 16:854–865, 2020.

Tilmann Gneiting, Hana Ševčíková, and Donald B. Percival. Estimators of fractal dimension: Assessing the roughness of time series and spatial data. *Statistical Science*, pages 247–277, 2012.

Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2004.

Alessandro Rozza, Gabriele Lombardi, Claudio Ceruti, Elena Casiraghi, and Paola Campadelli. Novel high intrinsic dimensionality estimators. *Machine learning*, 89:37–65, 2012.

David R. Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Thiry-fourth annual ACM symposium on Theory of computing*, pages 741–750, 2002.

Michael E. Houle. Dimensionality, discriminability, density and distance distributions. In *IEEE 13th International Conference on Data Mining Workshops*, pages 468–473, 2013.

Michael E. Houle. Local intrinsic dimensionality i: an extreme-value-theoretic foundation for similarity applications. In *International Conference on Similarity Search and Applications*, pages 64–79. Springer, 2017c.

Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016.

Pratyush Maini, Eric Wong, and Zico Kolter. Adversarial robustness against the union of multiple perturbation models. In *International Conference on Machine Learning*, pages 6640–6650, 2020.

Bernard W Silverman. *Density estimation for statistics and data analysis*. Routledge, 2018.

Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2, 2010.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.