

Effect of Ambient-Intrinsic Dimension Gap on Adversarial Vulnerability

Rajdeep Haldar
Purdue University

Yue Xing
Michigan State University

Qifan Song
Purdue University

Abstract

The existence of adversarial attacks on machine learning models imperceptible to a human is still quite a mystery from a theoretical perspective. In this work, we introduce two notions of adversarial attacks: natural or on-manifold attacks, which are perceptible by a human/oracle, and unnatural or off-manifold attacks, which are not. We argue that the existence of the off-manifold attacks is a natural consequence of the dimension gap between the intrinsic and ambient dimensions of the data. For 2-layer ReLU networks, we prove that even though the dimension gap does not affect generalization performance on samples drawn from the observed data space, it makes the clean-trained model more vulnerable to adversarial perturbations in the off-manifold direction. Our main results provide an explicit relationship between the ℓ_2, ℓ_∞ attack strength of the on/off-manifold attack and the dimension gap.

1 Introduction

Neural networks exhibit a universal approximation property, enabling them to learn any continuous function up to arbitrary accuracy (Hornik et al., 1989; Pinkus, 1999; Leshno et al., 1993). This has been used to address complex classification problems by learning their decision boundaries. However, although neural networks are known to generalize well (Liu et al., 2019), they are vulnerable to certain small-magnitude perturbations, hampering classification accuracy, which are known as *adversarial examples* (Szegedy et al., 2013; Goodfellow et al., 2014; Carlini and Wagner, 2017;

Madry et al., 2017).




Figure 1: Mental image: The oracle decision boundary (ODB) (green dashed line) determines the label (blue or red) of any point in the Euclidean space. The observed data space consists of 1-dimensional line segments immersed in the 2-dimensional space. The model learns the estimated decision boundary (EDB) (black dotted line) based on the observed data.

For a classification problem, assume that there exists an *oracle decision boundary* (ODB), which is the human decision boundary that determines the labels for any lawful input (e.g., a consensus digit recognition rule for any arbitrary gray-level images). The goal of training a neural network classifier is to obtain an *estimated decision boundary* (EDB) based on observed data (sample-label pairs observed). To facilitate misclassification by the neural network, a perturbation must push the original sample across the EDB.

Different types of attacks can be defined when the data is only observable in a particular subspace of the lawful input space (the data space). Usually, the EDB can estimate the ODB within the data space reasonably, given a large sample size. Thus, an adversarial perturbation restricted to the data space, which aims to cross over the ODB, tends to cross over the EDB as well. We call these adversarial attacks *natural* or *on-manifold* attacks. Even an oracle will be prone to wrong classifications with these adversarial perturbations.

In contrast, outside the data space, where no observation is made, the EDB curve is merely an extrapolation extended from the data space. It will not necessarily approximate the ODB well. Consequently, an adversarial perturbation outward from the data space, facilitated by crossing the EDB, will not necessarily try to push the sample over the ODB. We call these perturbations *unnatural* or *off-manifold* attacks. Under these attacks, an oracle can accurately distinguish and classify the perturbed sample, but the trained model fails. Figure 1 is provided to illustrate this phenomenon with a simple example. Shamir et al. (2022)’s observations also support the phenomenon we describe, where the EDB learned by neural networks, over the training process, tends to align with the data space first and then form dimples around differently labeled clusters to attain correct classification on the observed data space, leaving them vulnerable to examples outside the data space.



Figure 2: Original image is a digit 5. (a) Natural attack; (b) Unnatural attack (minor changes in pixel values of the background compared to the original). Under both of these attacks, the neural network classifies the image as 3 instead of 5.

Figure 2 shows an example of an *unnatural* and *natural* attack on a digit 5 image in the MNIST dataset (LeCun et al., 2010). The unnatural attacked image can still be classified as 5 by a human/oracle; however, the model classifies it as digit 3. The vulnerability of neural networks to *unnatural* attacks is a pressing issue.

In this work, we obtain distinguished results for off-manifold and on-manifold attacks. Our results imply that unnatural/off-manifold attacks are fundamental consequences of the dimension gap between the local intrinsic dimension of the observed data space and the ambient dimension. That is,

Theorem 1.1 (Informal version of Theorem 4.2). *As the dimension gap increases, the strength of ℓ_2, ℓ_∞ attack required to misclassify the sample decreases, i.e., the model is more vulnerable. Notably, for ℓ_∞ attack, the strength asymptotically goes to 0 w.r.t. the dimension gap.*

We use intrinsic dimension as an umbrella term for well-known concepts of *Manifold dimension*, *Minkowski*

dimension, *Hausdorff dimension* etc. In our theoretical setting (Section 3.1), all of these notions of dimension are equivalent and refer to the signal residing in a lower dimensional set locally.

	Ambient	LPCA	MLE(k=5)	TwoNN
MNIST	784	38	13.36	14.90
CIFAR-10	3072	9	27.66	31.65

Table 1: Intrinsic dimension estimation

To fuel the reader’s motivation, Table 1 is presented to show the estimated local intrinsic dimension for the standard image datasets MNIST and CIFAR-10. Regardless of the estimation method (local PCA(Fukunaga and Olsen, 1971), maximum likelihood estimation(Haro et al., 2008), Two nearest-neighbor(Facco et al., 2017)), the ambient dimension of the images is way larger than the intrinsic dimension of the space they reside in. This paper reveals that this is the root cause of the existence of *unnatural attacks*. Furthermore, as implied by our theorem, a larger ambient-intrinsic dimension gap, e.g., training images with unnecessarily high resolution but no more useful details, leads to higher vulnerability.

2 Related Works

Besides the literature, as mentioned earlier, below is a list of works ordered by relatability to this paper.

Adversarial study motivated by a manifold view (Xiao et al., 2022) assumes a manifold view of the data and focuses on a generative model to learn the data manifold and conceive on-manifold attacks in practice. (Zhang et al., 2022) conducts a theoretical study to decompose the adversarial risk into four geometric components: the standard risk, nearest neighbor risk, in-manifold, and normal (off-manifold in our setting) risk. Our paper also subscribes to the manifold hypothesis; Compared to the above two studies, our focus and mathematical setting are on the relationship between the data dimension gap and adversarial attack strength.

Adversarial examples in random ReLU networks Different from this paper which studies trained neural networks, some other studies, e.g., Daniely and Shacham (2020); Bubeck et al. (2021) also show the existence of ℓ_2 adversarial examples in two-layer ReLU networks with randomized weights and provide a corresponding rate for the strength of the ℓ_2 perturbation. Bartlett et al. (2021) further generalizes to a multi-layer network.

Adversarial examples due to implicit bias in two-layer Relu networks Vardi et al. (2022) and Frei et al. (2023) argue that adversarial examples are a consequence of the implicit bias in clean training despite the existence of robust networks. Extending from their work, we introduce the notion of the intrinsic-ambient dimension gap exhibited by the data, which gives rise to two different types of attacks *natural* or *on-manifold* and *unnatural* or *off-manifold*. The previously mentioned works can be reduced to the *on-manifold setting* in our framework.

Adversarial examples due to low dimensional data A recent work, Melamed et al. (2023), studies the adversarial robustness of clean-trained neural networks with data in a low-dimensional linear subspace and argues that the neural networks are vulnerable to attacks in the orthogonal direction.

Though the motivation is similar, our findings and outcomes are different. Detailed comparisons are provided in Section 3.5 for the model setup, and Section 4 for the final result, respectively. In short, the key differences arise from their setting of linear subspace, lazy training, and weight initialization. Though it is worth consideration in the finite training regime, the adversarial vulnerability considered in their setting does not hold if the neural network training attains convergence. On the other hand, our results are invariant to initialization, and our theoretical setting is more realistic than prior works.

3 Problem Setup

In this section, we introduce the detailed technical settings considered in this paper. We use consistent notations with Frei et al. (2023) in this paper. For any natural number n , the set $\{1, \dots, n\}$ is represented by the shorthand $[n]$. For any vector \mathbf{u} , $\|\mathbf{u}\|$ will denote the euclidean ℓ_2 norm. The notations \mathcal{O} and Θ are the standard asymptotic notations.

3.1 Data

In this paper, we study a binary classification problem with k clusters with respective labels $y^{(r)} \in \{-1, 1\}$, $r \in [k]$ and n observations. The intrinsic signal lies in d dimensions, and we observe some transformation of the signal in D dimensions ($D \geq d$). The dimension gap is denoted by $g = D - d$. Let $\mathbf{x}_i = \boldsymbol{\mu}^{(r)} + \boldsymbol{\xi}_i \in \mathbb{R}^d$ denote the i^{th} , $i \in [n]$ intrinsic signal for the r^{th} cluster. The vectors $\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(k)} \in \mathbb{R}^d$ are the means of the k clusters, and $\boldsymbol{\xi}_i \in \mathbb{R}^d$ is the deviation of the signal around the mean.

In terms of the actual observed data, denote the i^{th}

observation $\tilde{\mathbf{x}}_i = \mathbf{M}\mathbf{x}_i + \boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}_i \in \mathbb{R}^D$, where \mathbf{M} is a $D \times d$ matrix with d - orthonormal columns; the vector $\boldsymbol{\zeta}^{(r)} \in \mathbb{R}^D$ is the additional random cluster effect and $\boldsymbol{\zeta}^{(r)} \sim N(0, \tau^2 \mathbf{I}_D)$ for all $r \in [k]$; and $\boldsymbol{\omega}_i$ represents the ambient noise. We will also use an alternate representation for the i^{th} observation in terms of the transformed mean and deviation term in higher dimensions, $\tilde{\mathbf{x}}_i = \tilde{\boldsymbol{\mu}}^{(r)} + \tilde{\boldsymbol{\xi}}_i$ where $\tilde{\boldsymbol{\mu}}^{(r)} = \mathbf{M}\boldsymbol{\mu}^{(r)}$ and $\tilde{\boldsymbol{\xi}} = \mathbf{M}\boldsymbol{\xi}_i + \boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}_i$. The terms $\boldsymbol{\xi}_i, \boldsymbol{\omega}_i$ are also random, and we do not impose any specific distribution for them; instead, we will restrict their norm for our analysis. The notation $\mathcal{S} = \{(\tilde{\mathbf{x}}_i, y_i)\}_{i=1}^n$ is the collection of all n observation-label pairs used to train our model.

3.2 Data Interpretation

When studying low-dimensional information embedded in a higher-dimensional space, we do not want the original information to be lost or distorted. We restrict our analysis to a class of transformations known as *Isometric Transformations*. *Isometric* transformations are essential as they preserve distances (structure/information) between any pair of points in the origin space.

Any *isometric* transformation can be expressed as an affine transformation with a linear part consisting of orthogonal columns and a geometric translation. Hence, $\mathbf{M}(\cdot) + \boldsymbol{\zeta}^{(r)}$ preserves distance within each cluster. The matrix \mathbf{M} applies a linear transformation on the d -dimensional intrinsic signal (\mathbf{x}_i), immersing it into a D -dimensional space. After the immersion, a translation or shift $\boldsymbol{\zeta}^{(r)}$ is added corresponding to the r^{th} cluster. Then an ambient noise $\boldsymbol{\omega}$ is added to account for the stochasticity in the observation process. Each cluster can be considered a lower d dimensional manifold immersed in a D dimensional space up to ambient noise. Our data space is a union of such manifolds, and our setup can be used as a motivation for real-life datasets that locally reside in a lower dimensional manifold. Examples can be found in Tables 1, 2, 4, 5.

CIFAR class	0	1	...	8	9
LPCA	8	11	...	10	16
MLE (k=5)	18.34	21.20	...	19.80	24.42
TwoNN	21.77	25.98	...	24.77	29.11

Table 2: Intrinsic dimension of CIFAR-10 Dataset. The intrinsic dimension over all classes is similar to that for individual classes. (Complete version: Table 4 in the appendix.)

Within each cluster r , the distances are preserved by the transformation $\mathbf{M}(\cdot) + \boldsymbol{\zeta}^{(r)}$ up to the variance introduced by the ambient noise ($\text{Var}(\boldsymbol{\omega})$). Between clusters, the distances are preserved by $\mathbf{M}(\cdot)$ up to

a slightly larger variance due to cluster translation and the noise term ($2\tau^2 + \text{Var}(\boldsymbol{\omega})$). Intuitively, our setup transforms a low-dimensional signal into a high-dimensional signal via \mathbf{M} , preserving all information (pairwise angle and distances) up to some variance, where the within-cluster variance is lower than the between-cluster variance.

3.3 Neural Network

In this work, we will use a 2-layer ReLU network $\mathcal{N}_\theta : \mathbb{R}^D \rightarrow \mathbb{R}$ of width w as our classification model.

$$\mathcal{N}_\theta(\tilde{\mathbf{x}}) = \sum_{j=1}^w v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) \quad (1)$$

\mathbf{w}, b represent the weights and biases of the first layer; v_j is a scalar weight of the second layer for the j^{th} neuron. The function $\phi(x) = \max(0, x)$ is the ReLU activation function. The vectorized collection of all the parameters $(\mathbf{w}_1^\top, \dots, \mathbf{w}_j^\top, b_1, \dots, b_j, v_1, \dots, v_j)^\top$ for the neural network is represented by $\boldsymbol{\theta}$. Both the two layers will be updated during the training.

3.4 Expression for Network Parameters

Before we proceed, we need to establish the training loss used in our framework. For our neural network model \mathcal{N}_θ parameterized by $\boldsymbol{\theta}$, the *empirical loss* $\mathcal{L}(\boldsymbol{\theta})$ can be defined on the dataset $\mathcal{S} = \{(\tilde{\mathbf{x}}_i, y_i)\}_{i=1}^n$ as follows:

$$\mathcal{L}(\boldsymbol{\theta}) := \frac{1}{n} \sum_{i=1}^n \ell(y_i \cdot \mathcal{N}_\theta(\tilde{\mathbf{x}}_i)). \quad (2)$$

When $\ell(z) = e^{-z}$ or $\ell(z) = \log(1 + e^{-z})$, $\mathcal{L}(\boldsymbol{\theta})$ is the empirical *exponential loss* or *logistic loss* respectively.

Some *implicit bias* literature studies the limiting behavior of the training process. Based on Lyu and Li (2020), two-layer ReLU networks have directional convergence to the solution of the *maximum margin problem* (James et al., 2013):

Theorem 3.1 (Lyu and Li (2020)). *Consider $\mathcal{N}_\theta(\tilde{\mathbf{x}})$ trained via gradient descent or gradient flow for the exponential or logistic loss. If for some time t_0 in the training process, the loss reaches a sufficiently small value $\mathcal{L}(\theta(t_0)) < \frac{1}{n}$. Then, under some regularity and smoothness assumptions, $\frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|}$ converges to the KKT solution of the maximum margin problem:*

$$\min_{\boldsymbol{\theta}} \frac{1}{2} \|\boldsymbol{\theta}\|^2 \quad \text{s.t.} \quad \forall i \in [n] \quad y_i \mathcal{N}_\theta(\tilde{\mathbf{x}}_i) \geq 1. \quad (3)$$

Furthermore, $\mathcal{L}(\boldsymbol{\theta}(t)) \rightarrow 0$ and $\|\boldsymbol{\theta}(t)\| \rightarrow \infty$ as $t \rightarrow \infty$.

Theorem 3.1 shows the limiting behavior of the gradient flow in clean training and generally applies to any

homogeneous network defined in Lyu and Li (2020). A consequence of the above theorem is that we can express the parameters in terms of the data, i.e., $\boldsymbol{\theta} = \sum_i^n \lambda_i y_i \nabla_{\boldsymbol{\theta}} \mathcal{N}_\theta(\tilde{\mathbf{x}}_i)$, where $\lambda_i \geq 0$ are the Lagrangian multipliers. Specifically,

$$\mathbf{w}_j = \sum_{i \in [n]} v_j \lambda_i y_i \phi'_{i,j} \tilde{\mathbf{x}}_i; \quad b_j = \sum_{i \in [n]} v_j \lambda_i y_i \phi'_{i,j} \quad (4)$$

$\phi'_{i,j}$ denotes the sub gradient of $\phi(\mathbf{w}_j^\top \tilde{\mathbf{x}}_i + b_j)$.

3.5 Volatile biases

In the previous section, we saw that the network parameters can be explicitly written in terms of the observed dataset. In this section, we will leverage that closed-form expression to justify the existence of vulnerable weight components responsible for *unnatural attacks*.

Let $\mathbf{P}_o = \mathbf{M}(\mathbf{M}^\top \mathbf{M})^{-1} \mathbf{M}^\top = \mathbf{M} \mathbf{M}^\top$ be the projection matrix of the image space where the intrinsic signal is initially immersed. On the contrary, $\mathbf{P} = \mathbf{I}_D - \mathbf{P}_o$ is the projection onto the co-kernel, where no component of the intrinsic signal \mathbf{x} resides. We will associate the terms *on-manifold* and *off-manifold* to the spaces dictated by \mathbf{P}_o and \mathbf{P} , respectively.

In the neural network, each weight \mathbf{w}_j can be decomposed into an *authentic-weights* part $\mathbf{P}_o \mathbf{w}_j$ and a *volatile biases* part $\mathbf{P} \mathbf{w}_j$. Multiplying, \mathbf{P}_o, \mathbf{P} to Eq. (4) we get:

$$\begin{aligned} \mathbf{P}_o \mathbf{w}_j &= \sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} \left[\mathbf{M} \mathbf{x}_i + \mathbf{P}_o \left(\boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}_i \right) \right] \\ \mathbf{P} \mathbf{w}_j &= \sum_{i \in I^{(r)}} \sum_{r \in [k]} \lambda_i y_i v_j \phi'_{i,j} \mathbf{P} \left(\boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}_i \right), \end{aligned} \quad (5)$$

where $I^{(r)}$ is the index set of samples in the r^{th} cluster. The authentic weights are the ones that carry all the information of the intrinsic signal \mathbf{x} from the lower d dimension space; the volatile biases are dummy weights (independent of \mathbf{x}) that act just as additional bias terms during the *clean training*. Any interaction between the immersed intrinsic signal and volatile biases is zero as $\mathbf{P} \mathbf{M} = 0$. However, they can be exploited by *off-manifold* attacks to change the neuron output.

Comparison to Melamed et al. (2023) In Melamed et al. (2023), the setting is equivalent to a linear transformation of a low-dimensional intrinsic signal to a higher dimension, i.e., under our framework, $\tilde{\mathbf{x}} = \mathbf{L} \mathbf{x}$ where \mathbf{L} is any $D \times d$ matrix. Let $\pi^\perp = \mathbf{I}_D - \mathbf{L}(\mathbf{L}^\top \mathbf{L})^{-1} \mathbf{L}^\top$ be the projection matrix. Using Theorem 3.1 and Eq. (4) in their setting, we have $\pi^\perp \mathbf{w}_j = 0, j \in [w]$.

The consequence is that, based on results in Theorem 3.1, at convergence, the orthogonal weights converge

to 0; hence, these weights have no contribution to the neuron output. The adversarial vulnerability described in Melamed et al. (2023) results from finite training and large weight initialization. In real practice, since one can only train finite iterations, the initialization property observed by Melamed et al. (2023) is still a factor to consider. However, it is not the driving force for the existence of *unnatural attacks*. We also validate these arguments with experiments on simulated and real-life datasets (Section 5.1).

3.6 Assumptions

In this section, we introduce some additional assumptions to simplify the analysis.

Assumption 3.2. *We have:*

- (A1) $\|\boldsymbol{\mu}^{(j)}\| = \sqrt{d}$ for all $r \in [k]$.
 (A2) $k(\max_{i \neq j} |\langle \boldsymbol{\mu}^{(i)}, \boldsymbol{\mu}^{(j)} \rangle| + \Delta' + 1) = c'(d - \Delta' + 1)$. ;
 $c' < 1/10$
 (A3) $\|\boldsymbol{\xi}\| \leq \sqrt{2} \ln(d)$; $\|\boldsymbol{\omega}\| \leq 1$
 (A4) $c_3(g, \tau, k) = \frac{g\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13g\tau^2}{8}} > 0$;
 $c_5(d, \tau, k) = \frac{d\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13d\tau^2}{8}} (\sqrt{2 \ln d + 1}) > 0$;

where $\Delta' = \mathcal{O}(\max\{\sqrt{d} \ln d, D\tau^2\})$ (Definition B.2).

Assumptions (A1), (A2) are similar to the assumptions used in Frei et al. (2023). Condition (A1) states that the signal strength must grow with the intrinsic dimension. It can be generalized by replacing \sqrt{d} to $\Theta(\sqrt{d})$ without changing the analysis in any meaningful way. Condition (A2) ensures that the means are near orthogonal to each other; we impose that pairwise interaction between the means is less than $\|\boldsymbol{\mu}^{(r)}\|^2 = d$ for any cluster r . This assumption can be satisfied for some $\tau^2 = \mathcal{O}(d/D)$.

For the other two assumptions, (A3) are bounds on the stochastic deviation and error terms. In general, these bounds can be relaxed to $\mathcal{O}(\ln d)$, $\mathcal{O}(1)$ respectively. These bounds are easy to satisfy for sufficiently large d , for instance, $\boldsymbol{\xi} \sim N(0, \mathbf{I}_d)$; $\boldsymbol{\omega} \sim N(0, D^{-1}\mathbf{I}_D)$ satisfy the above bounds with high probability. In addition, Condition (A4) is a mild assumption. The order of the positive terms is higher than the negative terms for both $c_3(g, \tau, k)$ and $c_5(d, \tau, k)$. A constant can always be incorporated inside τ to satisfy the assumption.

4 Main Results

In this section, we will state our main results. We study the adversarial vulnerability for any reasonable

testing samples that arise from the observable data space (union of manifolds) and are homogeneous with training samples. Formally, we define such examples as *nice examples* 4.1. These examples can be expressed as the transformation of some d dimensional intrinsic signal described in Section 3.1.

Definition 4.1 (Nice example). *An $\tilde{\mathbf{x}} \in \mathbb{R}^D$ is nice if it can be expressed as $\tilde{\mathbf{x}} = \mathbf{M}\mathbf{x} + \boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}$, for some $r \in [k]$ and $\mathbf{x} = \boldsymbol{\mu}^{(r)} + \boldsymbol{\xi}$ such that $\|\boldsymbol{\xi}\| \leq \sqrt{2} \ln(d)$, $\|\boldsymbol{\omega}\| \leq 1$.*

In the following theorem, we show that for any *nice-example*, the clean-trained neural network accurately predicts the labels.

Theorem 4.1. *Suppose \mathcal{S} satisfies the assumption 3.2 and clean training has attained convergence under the setting of Theorem 3.1. Let $\tilde{\mathbf{x}} \in \mathbb{R}^D$ and $r \in [k]$ such that $\tilde{\mathbf{x}}$ is a “nice example” (Definition 4.1). Then with probability $1 - 2 \exp(-\frac{D}{16})$, $\text{sign}(\mathcal{N}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}})) = y^{(r)}$.*

Theorem 4.1 states that the trained neural network can generalize well with a high probability when we have an example from the low-dimensional data space. Asymptotically, as the ambient dimension ($D \rightarrow \infty$) increases, this probability converges to 1.

However, even if the clean training generalizes well, our following theorem states that for any *nice example* $\tilde{\mathbf{x}}$, there exists perturbation \mathbf{z}_{\perp} and \mathbf{z}_{\parallel} in the *on-manifold* and *off-manifold* direction which successfully misclassify the model with high probability.

Theorem 4.2. *Suppose \mathcal{S} satisfies Assumption 3.2 and clean training has attained convergence under the setting of Theorem 3.1. Let $\tilde{\mathbf{x}} \in \mathbb{R}^D$ and $r \in [k]$ such that $\tilde{\mathbf{x}}$ is a “nice example” (Definition 4.1). Then there exist perturbations $\mathbf{z}_{\perp} \propto \mathbf{u}_{\perp}$, $\mathbf{z}_{\parallel} \propto \mathbf{u}_{\parallel}$ with probability at least $1 - \delta_1, 1 - \delta_2$ respectively; Such that, $\mathcal{N}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}} - \mathbf{z}_{\perp})$, $\mathcal{N}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}} - \mathbf{z}_{\parallel}) \leq -1$ for $r \in Q_+$ and $\mathcal{N}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}} + \mathbf{z}_{\perp})$, $\mathcal{N}_{\boldsymbol{\theta}}(\tilde{\mathbf{x}} + \mathbf{z}_{\parallel}) \geq 1$ for $r \in Q_-$. Also,*

$$\begin{aligned} \|\mathbf{z}_{\perp}\| &= \mathcal{O}\left(\frac{d}{c\sqrt{k}g\tau^2}\right), \|\mathbf{z}_{\perp}\|_{\infty} = \mathcal{O}\left(\frac{d\sqrt{2\log(2g)}}{cg\tau}\right), \\ \|\mathbf{z}_{\parallel}\| &= \mathcal{O}\left(\sqrt{\frac{d}{c^2k(2+\tau^2)}}\right), \\ \|\mathbf{z}_{\parallel}\|_{\infty} &= \mathcal{O}\left(\frac{\sqrt{d} + \tau\sqrt{2\log(2d)}}{c(2+\tau^2)}\right), \end{aligned}$$

where $\mathbf{u}_{\parallel} = \sum_{q \in [k]} y^{(q)} (\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)})$ and $\mathbf{u}_{\perp} = \sum_{q \in [k]} y^{(q)} \mathbf{P} \boldsymbol{\zeta}^{(q)}$. Q_+, Q_- are a set of clusters with labels 1, -1 respectively and $c = k^{-1} \min\{|Q_+|, |Q_-|\}$. Probabilities δ_1, δ_2 are defined in Lemmas C.1.5, C.1.7 respectively.

Theorem 4.2 provides an upper bound on the strength

of an *on-manifold* (\mathbf{z}_{\parallel}) and *off-manifold* (\mathbf{z}_{\perp}) adversarial perturbation for any *nice example*. These upper bounds hold with probability $1 - \delta_1, 1 - \delta_2$ for on/off-manifold perturbations respectively, where $\delta_1 \rightarrow 0$ as $D \rightarrow \infty$ and $\delta_2 \rightarrow 0$ as $D, d \rightarrow \infty$.

Note that $\tau^2 = \mathcal{O}(d/D)$ based on our assumptions (Section 3.6), hence we see that ℓ_2 strength of \mathbf{z}_{\perp} decreases as the dimension gap g increases and asymptotically (in g) is equivalent to $\mathcal{O}(d/(c\sqrt{k}))$. For \mathbf{z}_{\parallel} the ℓ_2 strength increases and saturates to $\mathcal{O}(\sqrt{d/(c^2k)})$ as $g \rightarrow \infty$.

On the other hand, the ℓ_{∞} strength of \mathbf{z}_{\perp} and \mathbf{z}_{\parallel} goes to 0 and $\mathcal{O}(\sqrt{d})$, respectively as $g \rightarrow \infty$. That is, asymptotically, the ℓ_2 strength of *on-manifold* is a bit smaller (or comparable¹) than the *off-manifold* perturbation, but in contrast the ℓ_{∞} strength for the *off-manifold* perturbation goes to 0 while the *on-manifold* perturbation is still some constant depending on d . Additionally, the asymptotic upper bound for *on-manifold* perturbation reduces to the rate to Frei et al. (2023); our framework is an extension introducing dimension gap and two distinct attack spaces, Frei et al. (2023) results can be related to our *on-manifold* perturbation results. As these are upper-bound results to an existential result, they also act as an upper bound to the smallest adversarial perturbation. In Section 5, we will validate the tightness of these upper bounds with empirical experiments.

Rate comparison with Melamed et al. (2023)

Despite the differences in Melamed et al. (2023) and our theoretical setting (Remark 3.5), we can attempt to draw some parallels with their main result. In Melamed et al. (2023), an adversarial perturbation (off-manifold) exists as a function of orthogonal weights when an observation $\tilde{\mathbf{x}}$ strictly resides in a d dimensional linear subspace of a D dimensional Euclidean space (dimension gap $g = D - d$). The ℓ_2 strength of this off-manifold perturbation in our notation can be expressed as $\mathcal{O}\left(N(\tilde{\mathbf{x}})\sqrt{\frac{D}{g}}\right)$ up to some constants. Here, $N(\cdot)$ represents their two-layer ReLU lazy-trained network. For the case when $\tau^2 = \Theta\left(\frac{d}{D}\right)$ in Theorem 4.2, we can obtain a similar rate for $\|\mathbf{z}_{\perp}\|$ in our framework. However, for smaller τ^2 , the rate in our result is smaller.

5 Experiments

In this section, we conduct a simulation study and experiments in the MNIST dataset to exhibit that the data dimension plays an essential role in adversarial robustness. Due to the page limit, we postpone some

¹When $\tau^2 = \Theta\left(\frac{d}{D}\right)$ then $\|\mathbf{z}_{\perp}\|$ and $\|\mathbf{z}_{\parallel}\|$ are both $\mathcal{O}\left(\sqrt{\frac{d}{c^2k}}\right)$ as $g \rightarrow \infty$.

figures to the appendix.

5.1 Simulation Study

To generate the data, we first generate k means $\{\boldsymbol{\mu}^{(r)}\}_{r=1}^k$ corresponding to the k clusters in the lower d dimension from $\boldsymbol{\mu}^{(r)} \sim N(0, \mathbf{I}_d)$. This enforces that $\|\boldsymbol{\mu}^{(r)}\| = \Theta(d)$ with high probability, which is a relaxed version of assumption (A1). Given the k cluster means, we generate $n = 1000$ intrinsic data signals $\{\mathbf{x}_i\}_{i=1}^{1000}$ centered around these means drawn uniformly. For each i , $\mathbf{x}_i = \boldsymbol{\mu}_i + \boldsymbol{\xi}_i$, where $\boldsymbol{\mu}_i \sim \mathcal{U}(\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(k)})$ and $\boldsymbol{\xi}_i \sim N(0, \mathbf{I}_d/d)$, which enforces assumption (A3) with high probability.

After getting the intrinsic data signals $\{\mathbf{x}_i\}_{i=1}^{1000}$, we immerse them into a $D (\gg d)$ dimension space to get our observed data samples $\{\tilde{\mathbf{x}}_i\}_{i=1}^{1000}$. In alignment with our theoretical setting, we apply a matrix \mathbf{M} for the initial linear transformation. We emulate this matrix by generating d, D -dimensional $N(0, D^{-1}\mathbf{I}_D)$ Gaussian vectors m_1, \dots, m_d ; which act as the column of \mathbf{M} . This is justified as when $D \gg d$ and $i \neq j$, $m_i^{\top} m_j = O_p(1/\sqrt{D})$, i.e., the d vectors are almost orthogonal. Lastly, we add translations $\{\boldsymbol{\zeta}^{(r)}\}_{r=1}^k$ corresponding to each cluster. We generate $\boldsymbol{\zeta}^{(r)} \sim N(0, \tau^2 \mathbf{I}_D)$ for each cluster $r = 1, \dots, k$, the choice of $\tau^2 = D^{-1}, dD^{-1}$ corresponds to $\|\boldsymbol{\zeta}^{(r)}\| = \Theta(1), \Theta(\sqrt{d})$ respectively, with high probability. The architecture and training details can be found in Appendix D.1.1.

To validate our theory, we need to find the minimal attack strength required to misclassify an example. This is not computationally feasible; hence as an alternative, we find the minimal attack strength required to misclassify a large proportion of the data. We iterate over attack strengths until reaching 10% robust testing accuracy. The details are in Appendix D.1.2.

In the first simulation, we fix the intrinsic dimension $d = 100$ and vary the ambient dimension D . The results are summarized in Figure 3. The top panel of Figure 3 is the ℓ_2 attack strength for \mathbb{R}^D , on-manifold, and off-manifold attacks. The bottom panel is the proportion of \mathbb{R}^D attack that lies on the manifold. The rows of Figure 3 corresponds to the cases when $\|\boldsymbol{\zeta}^{(r)}\|$ is $\Theta(\sqrt{d})$ and $\Theta(1)$ or $\tau^2 = dD^{-1}, D^{-1}$ respectively.

There are two observations. First, the trends for ℓ_2 attack strength in Figure 3 are as expected as in the discussion in Theorem 4.2.

Second, the bottom panel of the figure shows the ratio of the on-manifold component to the off-manifold component for the \mathbb{R}^D attack. We see that as D increases, this proportion decreases, implying that even when the attack has no constraint with an increase in D or g , the model becomes more vulnerable to off-manifold

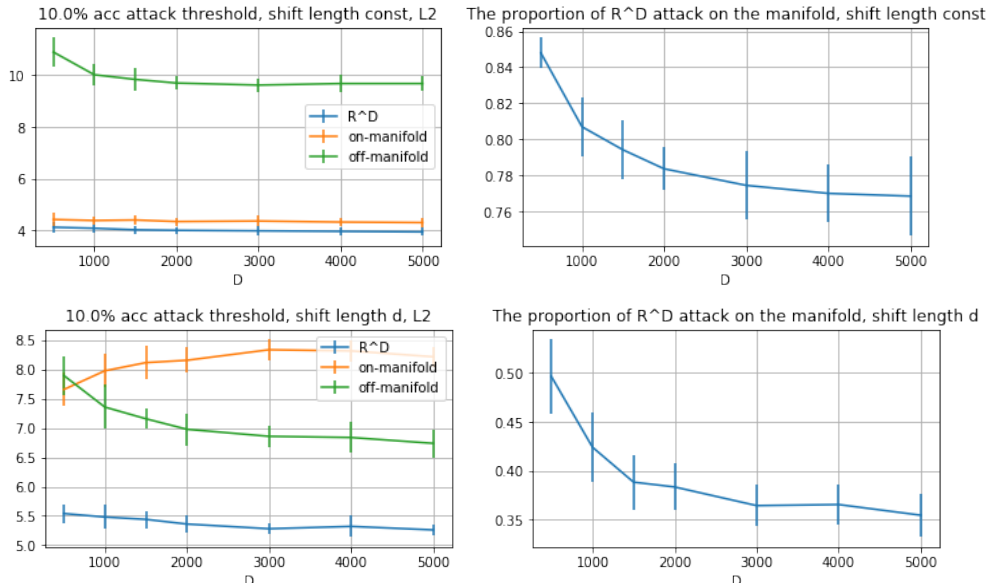


Figure 3: Attack strength threshold associating with 10% robust test accuracy (left) and the on-manifold proportion in \mathbb{R}^D attack. The top and bottom rows correspond to the cases $\|\zeta\| = \Theta(1), \Theta(\sqrt{d})$ respectively.

perturbations.

Figure 9 in the appendix is analogous to Figure 3 but considering the ℓ_∞ norm of the attack. The observations are consistent with the discussion succeeding Theorem 4.2. One difference is that for ℓ_∞ upper bound of the on-manifold attack, it also approaches 0 instead of a constant (in D, g) $\mathcal{O}(\sqrt{d})$. This is an artifact of generating \mathbf{M} from random Gaussian vectors instead of the fixed- \mathbf{M} scenario considered in Theorem 4.2.

In the appendix, we also plot the attack threshold generated by ℓ_∞ attacks. The results are provided in Figure 10. They are identical to the results of Figure 9, implying our bounds provide accurate guidance for the ℓ_∞ case. Intuitively, when D increases, the total attack budget grows when fixing the ℓ_∞ attack strength, so we observe a decreasing trend in the top panel of Figure 10. Besides, if we multiply the attack strength by \sqrt{D} , we can observe a similar trend as the ℓ_2 attack in Figure 3, consistent with our prior discussions.

In contrast to the first simulation, we fix the ambient dimension $D = 2000$ in the second simulation and vary the intrinsic dimension d . We track the attack threshold as a function of d . The results are provided in the appendix, Figure 11. One can see that when d increases, there is more on-manifold information; thus, the model is more robust.

Effect of Initialization We also conduct experiments to examine the effect of neural network initialization and ℓ_2 regularization (weight decay). In the appendix, the results are postponed to Figure 12 and 13. The trained model slightly improves when using a

small initialization or weight decay compared to Figure 3. However, the vulnerability towards off-manifold attack in these two scenarios is similar to Figure 3. That is, dimension-gap hampers robustness regardless of the initialization.

5.2 MNIST Dataset

We also conduct experiments on the MNIST dataset to validate the effect of the gap between intrinsic and ambient dimensions on the adversarial robustness. According to our theoretical setup, the MNIST dataset can be considered as a union of low-dimensional local clusters immersed in a high-dimensional space as bolstered by Table 5 in the appendix.

Different from the simulation, we cannot access the low-dimensional intrinsic signal corresponding to the MNIST images; the original data set already has some dimension gap. Regardless, we can still attempt to track the effect of the additional dimension gap on the adversarial robustness. To make a fair assessment, we implement *isometric transformations* on the original MNIST images that increase the ambient dimension D to track the dependence of attack strength with varying dimensions.

We pad gray (0.5) pixels on the image’s border to increase the ambient dimension (Figure 4). Adding a pad of P pixels along the borders increases the ambient dimension D by $4P^2 + 2P$. Adding gray borders to the images is an isometric transformation, i.e., the distances between every image are preserved, and no information is lost. Padding the images will allow us to emulate the effect of increasing the dimension gap



Figure 4: (a) Original image 28x28, $D = 784$; (b) Padded image 34x34, $P = 3$, $D = 1156$.

without distorting the data. Hence the attack strengths across dimensions are still comparable.

We use a convolution network to train the model. The architecture and training details are in Appendix D.2.1.

We consider two settings: (1) training for 10 epochs and (2) training until the loss is 0.02. While presenting the results for (1) in the main content, we postpone the results for (2) in the appendix. After training the neural network, we examine the robustness against ℓ_2 and ℓ_∞ attacks in the testing data, and the results are summarized in Figure 5 and 6.

From our theory, it is expected that when increasing the ambient dimension D , the model becomes less robust, i.e., we require smaller ℓ_2 attacks until the norm saturates to some constant. In the top panel of Figure 5, one can see that when P increases (in turn increasing D), we only need a smaller ℓ_2 attack to achieve a low accuracy and this small value converges to some constant which aligns with the theory. In the bottom panel of Figure 5, we plot the accuracy w.r.t. different attack strengths. For ℓ_∞ attack, based on the observations for ℓ_2 attack in Figure 5, it is intuitive that the neural network is less robust against ℓ_∞ attack as well when D increases as observed in Figure 6.

5.3 Fashion-MNIST

In the same vein as Section 5.2, we verify the effect of increased dimension gap (by padding grey pixel as in the MNIST experiment) on the vulnerability of the model on FMNIST images (Xiao et al., 2017) in Figure 7. Figure 7 illustrates the same pattern of the bottom plot in Figure 6. We use the same convolution network architecture and training setting used for the MNIST experiment (refer to Section D.2.1) averaging over 30 random seeds.

5.4 Imagenet Binary Classification

We subset two classes (Dog and Church) from the Imagenet dataset (Deng et al., 2009) and scale all the images to a uniform resolution of 320x320 pixels. In contrast to the FMNIST/MNIST experiments, where we emulate the effect of increasing the di-

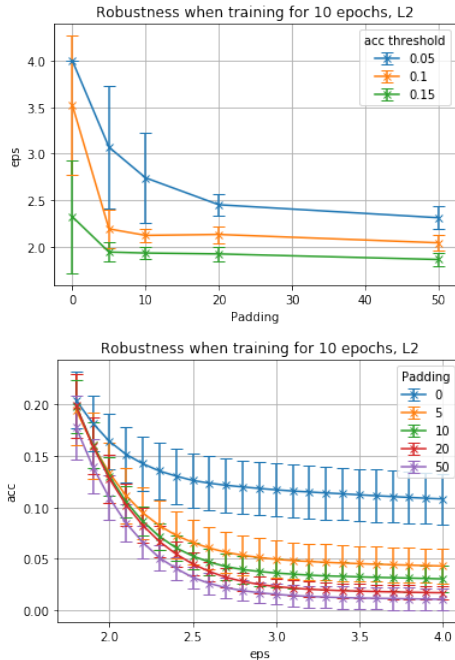


Figure 5: The relationship between attack strength and the padding number P in Table 6 (Top) and the relationship between robust test accuracy and attack strength for different values of P (Bottom). ℓ_2 attack. 10 training epochs.

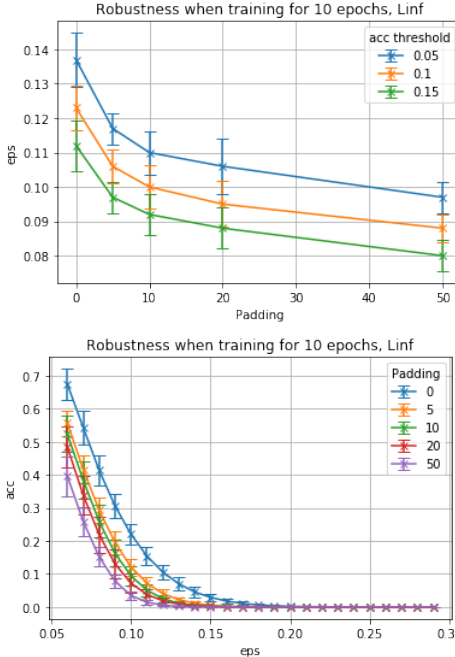


Figure 6: The relationship between attack strength and the padding number P (Top), and the relationship between robust test accuracy and attack strength under different P . (Bottom) ℓ_∞ attack. 10 training epochs.

mension gap via adding padding to the images, for Imagenet, we scale the images to different resolu-




Figure 7: Relationship between robust test accuracy and attack strength under different P . ℓ_∞ attack. 30 training epochs.

tions. The motivation of this experiment is that for higher-resolution images, the ambient dimension increases disproportionately w.r.t the useful features or on-manifold dimensions. From Table 3, even with different resolutions, the intrinsic dimension of the data remains almost constant. Hence, we can use resolution as a surrogate for the intrinsic-ambient dimension gap: higher-resolution images will tend to have a higher dimension gap.

Size(px ²)	Ambient	LPCA	MLE(k=5)	TwoNN
320	307200	12	22.75	31.27
128	49152	12	21.87	30.08
64	12288	12	20.88	28.93

Table 3: Intrinsic dimension is unchanged even with varying resolution. Imagenet (Dogs+Church).

We rescale the images to 64x64, 128x128, and 320x320 (base) pixels and train Resnet18 models (He et al., 2016) with 10^{-3} learning rate and Adam optimizer for 20 epochs in each resolution setting. All models have been trained for 20 epochs for uniform comparison, with ~ 0.02 training loss and 99% accuracy averaging over 30 seeds. After the clean training, the models are subjected to ℓ_∞ PGD attack of varying strengths up to $\epsilon = 16/255$.

Figure 8 illustrates that, as the resolution of an image increases, the vulnerability also increases, bolstering our dimension gap argument further.

6 Discussion and Conclusion

We introduce the notion of natural and unnatural attacks residing in on and off manifold space concerning the low dimensional data space. Despite the generalization capabilities of a cleanly trained model on examples




Figure 8: Relationship between robust test accuracy and attack strength under different resolutions. ℓ_∞ attack.

drawn from the data space, it will exhibit vulnerabilities to perturbations in both off and on manifold directions. In particular, the off-manifold adversarial perturbations can get arbitrarily small (in ℓ_∞ norm) with the increase in dimension gap.

Adversarial Training This work discusses the intrinsic-ambient dimension gap as the root cause for prevalent adversarial attacks in literature. However, we have yet to comment on adversarially trained models that are used to induce robustness to such attacks. We attempt to motivate the success of adversarial training through the lens of our framework. In contrast to clean training, adversarial training aims to minimize the loss over the worst possible attacks the data can exhibit. We can think of adversarial training as clean training on a dataset consisting of adversarial examples of the original data. Even though the original data space is a low dimensional subset of the whole Euclidean space, the dimension of adversarial data space will always be equal to the ambient dimension as we consider a D dimensional ϵ ball around the original data space. Where ϵ is a parameter determining the strength of attacks, adversarial data space makes the dimension gap 0 and, based on our theory, should not induce any *off-manifold* or *unnatural attacks*. Thinking of adversarial training as clean training on a much higher dimensional data space also sheds light on the complexity of adversarial training (Schmidt et al., 2018). A higher dimensional data space requires many more observations for a representative sample.

References

- Adler, R. J. and Taylor, J. E. (2007), “Gaussian inequalities,” *Random Fields and Geometry*, 49–64.
- Bartlett, P., Bubeck, S., and Cherapanamjeri, Y. (2021), “Adversarial examples in multi-layer random relu networks,” *Advances in Neural Information Processing Systems*, 34, 9241–9252.
- Bubeck, S., Cherapanamjeri, Y., Gidel, G., and Tachet des Combes, R. (2021), “A single gradient step finds adversarial examples on random two-layers neural networks,” *Advances in Neural Information Processing Systems*, 34, 10081–10091.
- Carlini, N. and Wagner, D. (2017), “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*, Ieee, pp. 39–57.
- Daniely, A. and Shacham, H. (2020), “Most ReLU Networks Suffer from ℓ^2 Adversarial Perturbations,” in *Advances in Neural Information Processing Systems*, eds. Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., Curran Associates, Inc., vol. 33, pp. 6629–6636.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009), “ImageNet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255.
- Facco, E., d’Errico, M., Rodriguez, A., and Laio, A. (2017), “Estimating the intrinsic dimension of datasets by a minimal neighborhood information,” *Scientific Reports*, 7.
- Frei, S., Vardi, G., Bartlett, P. L., and Srebro, N. (2023), “The Double-Edged Sword of Implicit Bias: Generalization vs. Robustness in ReLU Networks,” *arXiv preprint arXiv:2303.01456*.
- Fukunaga, K. and Olsen, D. (1971), “An Algorithm for Finding Intrinsic Dimensionality of Data,” *IEEE Transactions on Computers*, C-20, 176–183.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014), “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*.
- Haro, G., Randall, G., and Sapiro, G. (2008), “Translated Poisson Mixture Model for Stratification Learning,” *International Journal of Computer Vision*, 80, 358–374.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015), “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034.
- (2016), “Deep Residual Learning for Image Recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hornik, K., Stinchcombe, M., and White, H. (1989), “Multilayer feedforward networks are universal approximators,” *Neural networks*, 2, 359–366.
- James, G., Witten, D., Hastie, T., Tibshirani, R., et al. (2013), *An introduction to statistical learning*, vol. 112, Springer.
- Laurent, B. and Massart, P. (2000), “Adaptive estimation of a quadratic functional by model selection,” *Annals of statistics*, 1302–1338.
- LeCun, Y., Cortes, C., and Burges, C. (2010), “MNIST handwritten digit database,” *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- Ledoux, M. (2006), “Isoperimetry and Gaussian analysis,” *Lectures on Probability Theory and Statistics: Ecole d’Eté de Probabilités de Saint-Flour XXIV—1994*, 165–294.
- Leshno, M., Lin, V. Y., Pinkus, A., and Schocken, S. (1993), “Multilayer feedforward networks with a nonpolynomial activation function can approximate any function,” *Neural networks*, 6, 861–867.
- Liu, J., Bai, Y., Jiang, G., Chen, T., and Wang, H. (2019), “Understanding Why Neural Networks Generalize Well Through GSNR of Parameters,” in *International Conference on Learning Representations*.
- Lyu, K. and Li, J. (2020), “Gradient Descent Maximizes the Margin of Homogeneous Neural Networks,” .
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2017), “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*.
- Melamed, O., Yehudai, G., and Vardi, G. (2023), “Adversarial Examples Exist in Two-Layer ReLU Networks for Low Dimensional Data Manifolds,” *arXiv preprint arXiv:2303.00783*.
- Pinkus, A. (1999), “Approximation theory of the MLP model in neural networks,” *Acta numerica*, 8, 143–195.
- Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. (2018), “Adversarially robust generalization requires more data,” *Advances in neural information processing systems*, 31.

Shamir, A., Melamed, O., and BenShmuel, O. (2022), “The Dimpled Manifold Model of Adversarial Examples in Machine Learning,” .

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013), “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*.

Vardi, G., Yehudai, G., and Shamir, O. (2022), “Gradient methods provably converge to non-robust networks,” *Advances in Neural Information Processing Systems*, 35, 20921–20932.

Xiao, H., Rasul, K., and Vollgraf, R. (2017), “Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms,” .

Xiao, J., Yang, L., Fan, Y., Wang, J., and Luo, Z.-Q. (2022), “Understanding adversarial robustness against on-manifold adversarial examples,” *arXiv preprint arXiv:2210.00430*.

Zhang, W., Zhang, Y., Hu, X., Goswami, M., Chen, C., and Metaxas, D. N. (2022), “A Manifold View of Adversarial Risk,” in *International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 11598–11614.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. **Yes**, Refer to Sections 3.1, 3.6, 5 respectively.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. **Not Applicable**
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. **No**, we only have simulations and simple experiments for which the details provided in Section 5, Appendix D suffice.
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. **Yes**, Refer section 3.6.
 - (b) Complete proofs of all theoretical results. **Yes**, Refer to the proofs in Appendix C.
 - (c) Clear explanations of any assumptions. **Yes**, Refer section 3.6.
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). **No**, We are using simple simulations.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). **Yes**, details provided in Section 5, Appendix D.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). **Yes**, We provide error bars by running multiple runs.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). **No**, Computation is inexpensive and doesn’t need specialized infrastructure.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. **Yes**, Only external asset used is MNIST dataset which has been cited.
 - (b) The license information of the assets, if applicable. **Not Applicable**
 - (c) New assets either in the supplemental material or as a URL, if applicable. **Not Applicable**
 - (d) Information about consent from data providers/curators. **Not Applicable**
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. **Not Applicable**
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. **Not Applicable**
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. **Not Applicable**
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. **Not Applicable**

A Proof Sketch

We adapt the mathematical techniques used in Frei et al. (2023) to our framework.

1. We use the KKT condition in Eq. (4) to connect the neural network weights \mathbf{w}_j, b_j to the data $\tilde{\mathbf{x}}_i, y_i$. Then we can write the outputs at a new sample $\tilde{\mathbf{x}}$ (i.e., $\mathcal{N}_\theta(\tilde{\mathbf{x}}) = \sum_{j=1}^w v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j)$) in terms of interactions ($\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle$) between the observed signal $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}$. For the values of $\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle$, the means $\boldsymbol{\mu}^{(r)}$, $r \in [k]$ are nearly orthogonal [3.2 (A2)] and impose $\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle$ to be larger, if $\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}$ belong to the same cluster compared to the case when they are from different clusters. This ensures that the neuron output is consistent when a sample is from the same cluster. (Theorem 4.1)
2. We further study the adversarial attack. As discussed in Section 3.5, \mathbf{w}_j can be decomposed into $\mathbf{P}_o \mathbf{w}_j$ and $\mathbf{P} \mathbf{w}_j$. We design perturbations \mathbf{z}_\parallel and \mathbf{z}_\perp , such that $\mathbf{P}_o \mathbf{w}_j^\top \mathbf{z}_\parallel$ and $\mathbf{P} \mathbf{w}_j^\top \mathbf{z}_\perp$ can be made arbitrarily large in magnitude compared to $\mathbf{w}_j^\top \tilde{\mathbf{x}}$. For $\mathbf{z} = \mathbf{z}_\perp, \mathbf{z}_\parallel$, we have $\mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}) = \mathbf{w}_j^\top \tilde{\mathbf{x}} + \mathbf{P} \mathbf{w}_j^\top \mathbf{z}_\perp, \mathbf{w}_j^\top \tilde{\mathbf{x}} + \mathbf{P}_o \mathbf{w}_j^\top \mathbf{z}_\parallel$; thus we can manipulate the sign of each neuron after perturbation. (Theorem 4.2)

B Definitions and short hands

$\bar{\zeta}$ is a shorthand notation to bound $\boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}_i$ from above, for any $i \in [n], r \in [k]$ with very high probability. $\bar{\omega}$ will denote the upper bound for $\boldsymbol{\omega}_i$ for $i \in [n]$. If assumption 3.2 hold then $\bar{\omega} \leq 1$ and $\bar{\zeta}$ can be set to a constant in Definition B.2.

Definition B.1 (Semi-inner product). *For any positive semi-definite matrix A we can define a semi-inner product $\langle \cdot, \cdot \rangle_A$ such that for all $a, b \in \mathbb{R}^D$:*

$$\langle a, b \rangle_A = a^T A b$$

If A is a projection matrix, then $\langle a, b \rangle_A = \langle Aa, Ab \rangle$ where $\langle \cdot, \cdot \rangle$ is the standard inner product.

Definition B.2 (Constants).

$$\begin{aligned} c_2 &= 1 + \sqrt{2}; \bar{\zeta}^2 = 1 + 2\tau \sqrt{\frac{13D}{8}} + \frac{13D\tau^2}{8} \\ \Delta &= 2c_2 \sqrt{d} \ln(d) \\ \epsilon_x &= \bar{\zeta}^2 + 2\bar{\zeta} \sqrt{d} c_2 \\ \epsilon_\xi &= \bar{\zeta}^2 + 2\bar{\zeta} \sqrt{2} \ln d \\ \epsilon_{x\xi} &= \bar{\zeta}^2 + \bar{\zeta} (2\sqrt{2} \ln d + \sqrt{d} c_2) \\ \epsilon_\zeta^\parallel &= (\sqrt{2} \ln d + \bar{\omega}) \|\mathbf{P}_o \boldsymbol{\zeta}\|; \epsilon_\zeta^\perp = \bar{\omega} \|\mathbf{P} \boldsymbol{\zeta}\| \\ p &= \max_{i \neq j} |\langle \boldsymbol{\mu}^{(i)}, \boldsymbol{\mu}^{(j)} \rangle| \\ \Delta' &= \Delta + \epsilon_x \end{aligned}$$

C Proofs

Lemma C.0.1 (Properties). *Given \mathbf{x} belongs to the r^{th} cluster, if the assumptions 3.2 hold and $\|\boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}\| \leq \bar{\zeta}$ for any $r \in [k]$, we have the following properties:*

1. $|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi} \rangle| < \Delta$ and $|\langle \tilde{\boldsymbol{\xi}}_i, \tilde{\boldsymbol{\xi}} \rangle| \leq \Delta + \epsilon_\xi \leq \Delta'$.
2. $|\langle \mathbf{x}_i, \boldsymbol{\xi} \rangle| \leq \Delta$ and $|\langle \tilde{\mathbf{x}}_i, \tilde{\boldsymbol{\xi}} \rangle| \leq \Delta + \epsilon_{x\xi} \leq \Delta'$.
3. If $i \notin I^{(r)}$ then $|\langle \mathbf{x}_i, \mathbf{x} \rangle| \leq p + \Delta$ and $|\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle| \leq p + \Delta'$.
4. If $i \in I^{(r)}$ then $|\langle \mathbf{x}_i, \mathbf{x} \rangle - d| \leq \Delta$ and $|\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle - d| \leq \Delta'$.

Proof. Using cauchy-schwartz, (A3) and $\ln x < \sqrt{x}$; we have:

1. $|\langle \xi_i, \xi \rangle| \leq |2 \ln^2 d| < 2\sqrt{d} \ln d < \Delta$. Consequently,

$$\begin{aligned} |\langle \tilde{\xi}_i, \tilde{\xi} \rangle| &= |\langle \mathbf{M}\xi_i + \zeta^{(r')} + \omega_i, \mathbf{M}\xi + \zeta^{(r)} + \omega \rangle| \\ &= |\langle \xi_i, \xi \rangle + \langle \zeta^{(r')} + \omega_i, \zeta^{(r)} + \omega \rangle + \langle \mathbf{M}\xi_i, \zeta^{(r)} + \omega \rangle + \langle \mathbf{M}\xi, \zeta^{(r')} + \omega_i \rangle| \\ &\leq |\langle \xi_i, \xi \rangle| + \bar{\zeta}^2 + 2\|\mathbf{M}\|_{op}\sqrt{2} \ln d \cdot \bar{\zeta} \leq \Delta + \epsilon_\xi \leq \Delta + \epsilon_x = \Delta' \end{aligned}$$

$\|\mathbf{M}\|_{op} = 1$ as \mathbf{M} has orthogonal columns.

2. $|\langle \mathbf{x}_i, \xi \rangle| = |\langle \mu^{(r')} + \xi_i, \xi \rangle| \leq |\langle \mu^{(r')}, \xi \rangle| + |\langle \xi_i, \xi \rangle| \leq \sqrt{2}\sqrt{d} \ln d + 2 \ln^2 d < 2c_2\sqrt{d} \ln d = \Delta$.

$$\begin{aligned} \langle \tilde{\mathbf{x}}_i, \tilde{\xi} \rangle &= |\langle \mathbf{M}\mathbf{x}_i + \zeta^{(r')} + \omega_i, \mathbf{M}\xi + \zeta^{(r)} + \omega \rangle| \\ &\leq |\langle \mathbf{x}_i, \xi \rangle + \langle \zeta^{(r')} + \omega_i, \zeta^{(r)} + \omega \rangle + \langle \mathbf{M}\xi_i, \zeta^{(r)} + \omega \rangle + \langle \underbrace{\mathbf{M}\mathbf{x}_i}_{\mu^{(r')} + \xi_i}, \zeta^{(r')} + \omega_i \rangle| \\ &\leq \Delta + \bar{\zeta}^2 + \|\mathbf{M}\|_{op}\sqrt{2} \ln d \cdot \bar{\zeta} + \|\mathbf{M}\|_{op}(\sqrt{d} + \sqrt{2} \ln d) \cdot \bar{\zeta} \\ &\leq \Delta + \epsilon_{x\xi} \leq \Delta + \epsilon_x = \Delta' \end{aligned}$$

3. If $i \notin I^{(r)}$ then, $|\langle \mathbf{x}_i, \mathbf{x} \rangle| = |\langle \mu^{(r')} + \xi_i, \mu^{(r)} + \xi \rangle| \leq |\langle \mu^{(r')}, \mu^{(r)} \rangle| + |\langle \xi_i, \xi \rangle| \leq p + \Delta$.

Similarly, $|\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle| = |\langle \tilde{\mu}^{(r')} + \tilde{\xi}_i, \tilde{\mu}^{(r)} + \tilde{\xi} \rangle| \leq |\langle \tilde{\mu}^{(r')}, \tilde{\mu}^{(r)} \rangle| + |\langle \tilde{\xi}_i, \tilde{\xi} \rangle| = |\langle \mu^{(r')}, \mu^{(r)} \rangle| + |\langle \xi_i, \xi \rangle| \leq p + \Delta'$.

4. If $i \in I^{(r)}$ then, $|\langle \mathbf{x}_i, \mathbf{x} \rangle - d| = |\langle \mu^{(r)} + \xi_i, \mu^{(r)} + \xi \rangle - d| \leq \|\mu^{(r)}\| - d + \langle \xi_i, \xi \rangle \leq 0 + \Delta$.

Similarly, $|\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}} \rangle - d| = |\langle \tilde{\mu}^{(r)} + \tilde{\xi}_i, \tilde{\mu}^{(r)} + \tilde{\xi} \rangle - d| \leq \|\mu^{(r)}\| - d + \langle \tilde{\xi}_i, \tilde{\xi} \rangle \leq 0 + \Delta'$.

□

C.1 High probability statements

In this section, we show that any of the following events hold with a very high probability for any $q \in [k]$:

$$(E1) \quad \|\zeta^{(q)}\| \asymp \tau\sqrt{D}, \text{ in particular } \tau\sqrt{\frac{D}{2}} < \|\zeta^{(q)}\| < \tau\sqrt{\frac{13D}{8}}$$

$$(E2) \quad \zeta^{(q)\top} \mathbf{P}\zeta^{(q)} \asymp g\tau^2, \text{ in particular } \frac{g\tau^2}{2} < \zeta^{(q)\top} \mathbf{P}\zeta^{(q)} < \frac{13g\tau^2}{8}$$

$$(E3) \quad \zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)} \asymp d\tau^2, \text{ in particular } \frac{d\tau^2}{2} < \zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)} < \frac{13d\tau^2}{8}$$

$$(E4) \quad \|\mathbf{P}\zeta^{(q)}\|_\infty = \mathcal{O}\left(\tau\sqrt{2\log(2g)}\right), \text{ in particular } \|\mathbf{P}\zeta^{(q)}\|_\infty < 3\tau\sqrt{2\log(2g)}$$

$$(E5) \quad \|\mathbf{P}_o\zeta^{(q)}\|_\infty = \mathcal{O}\left(\tau\sqrt{2\log(2d)}\right), \text{ in particular } \|\mathbf{P}_o\zeta^{(q)}\|_\infty < 3\tau\sqrt{2\log(2d)}$$

$$(E6) \quad k \cdot \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}}| + \epsilon_\zeta^\perp\right) < \frac{1}{10} \cdot \left(\zeta^{(q)\top} \mathbf{P}\zeta^{(q)} - \epsilon_\zeta^\perp\right)$$

$$(E7) \quad k \cdot \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| + \epsilon_\zeta^\parallel\right) < \frac{1}{10} \cdot \left(\zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)} - \epsilon_\zeta^\parallel\right)$$

Lemma C.1.1. $\|\zeta^{(q)}\| \asymp \tau\sqrt{D} = \sqrt{h}$, $\zeta^{(q)\top} \mathbf{P}\zeta^{(q)} \asymp g\tau^2$ and $\zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)} \asymp d\tau^2$ w.h.p. for any $q \in [k]$. Specifically,

$$\begin{aligned} P\left(\frac{\zeta^{(q)\top} \mathbf{P}\zeta^{(q)}}{\tau^2} > \frac{13g}{8}\right) &\leq e^{-\frac{g}{16}}; P\left(\frac{\zeta^{(q)\top} \mathbf{P}\zeta^{(q)}}{\tau^2} < \frac{g}{2}\right) \leq e^{-\frac{g}{16}} \\ P\left(\frac{\zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)}}{\tau^2} > \frac{13d}{8}\right) &\leq e^{-\frac{d}{16}}; P\left(\frac{\zeta^{(q)\top} \mathbf{P}_o\zeta^{(q)}}{\tau^2} < \frac{d}{2}\right) \leq e^{-\frac{d}{16}} \\ P\left(\|\zeta^{(q)}\| > \tau\sqrt{\frac{13D}{8}}\right) &\leq e^{-\frac{D}{16}}; P\left(\|\zeta^{(q)}\| < \tau\sqrt{\frac{D}{2}}\right) \leq e^{-\frac{D}{16}} \end{aligned}$$

Proof. As \mathbf{P}, \mathbf{P}_o are projection matrices of rank g, d respectively, they are idempotent and hence $\zeta^{(q)\top} \mathbf{P} \zeta^{(q)} \tau^{-2} \sim \chi^2(g), \zeta^{(q)\top} \mathbf{P}_o \zeta^{(q)} \tau^{-2} \sim \chi^2(d)$. Also, $\|\zeta\|^2 \tau^{-2} \sim \chi^2(D)$.

Using Laurent and Massart (2000) lemma 1:

$$P\left(\frac{\zeta^{(q)\top} \mathbf{P} \zeta^{(q)}}{\tau^2} > g + 2\sqrt{g}\sqrt{x} + 2x\right) \leq e^{-x}; P\left(\frac{\zeta^{(q)\top} \mathbf{P} \zeta^{(q)}}{\tau^2} < g - 2\sqrt{g}\sqrt{x}\right) \leq e^{-x}$$

Setting $x = \frac{g}{16}$ we get:

$$P\left(\frac{\zeta^{(q)\top} \mathbf{P} \zeta^{(q)}}{\tau^2} > \frac{13g}{8}\right) \leq e^{-\frac{g}{16}}; P\left(\frac{\zeta^{(q)\top} \mathbf{P} \zeta^{(q)}}{\tau^2} < \frac{g}{2}\right) \leq e^{-\frac{g}{16}}$$

Using union bound we have $P\left(\frac{g}{2} < \frac{\zeta^{(q)\top} \mathbf{P} \zeta^{(q)}}{\tau^2} < \frac{13g}{8}\right) \geq 1 - 2e^{-\frac{g}{16}}$. One can get an identical statement for $\zeta^{(q)\top} \mathbf{P}_o \zeta^{(q)}$ replacing g with d .

Similarly, one can get

$$P\left(\frac{\|\zeta^{(q)}\|^2}{\tau^2} > \frac{13D}{8}\right) \leq e^{-\frac{D}{16}}; P\left(\frac{\|\zeta^{(q)}\|^2}{\tau^2} < \frac{D}{2}\right) \leq e^{-\frac{D}{16}}$$

Consequently, $P\left(\frac{D}{2} < \frac{\|\zeta^{(q)}\|^2}{\tau^2} < \frac{13D}{8}\right) \geq 1 - 2e^{-\frac{D}{16}}$ \square

Lemma C.1.2. For the projection matrix \mathbf{P} for every $i \in [D]$ we have $0 < p_{ii} < 1$. Where p_{ij} is the entry in i^{th} row and j^{th} column of \mathbf{P} . Also, there exists an indexing set L of length d such that for all $i \in L, p_{ii} = 0$.

Proof. As $\mathbf{P}^\top = \mathbf{P}$ and $\mathbf{P}^2 = \mathbf{P}$, we have $p_{ii} = p_{ii}^2 + \sum_{i \neq j} p_{ij}^2$.

This means $p_{ii}^2 \leq p_{ii}$, hence $0 < p_{ii} < 1$.

Considering, the eigenvalue decomposition of $\mathbf{P} = O^\top \Lambda O$ WLOG we can assume that Λ is ordered such that the first g diagonal elements are 1 and the rest d diagonal elements are 0. One can check that this enforces the first g diagonal elements of \mathbf{P} to non-zero and the rest d diagonal elements to be 0.

In general Λ won't be ordered hence an indexing set L can be used corresponding to the 0 eigenvalues; and $p_{ii} = 0$ for $i \in L$. \square

Corollary C.1.1. For the projection matrix \mathbf{P}_o for every $i \in [D]$ we have $0 < r_{ii} < 1$. Where r_{ij} is the entry in i^{th} row and j^{th} column of \mathbf{P}_o . Also, there exists an indexing set L_o of length g such that for all $i \in L_o, r_{ii} = 0$.

Proof. Notice that, $\mathbf{P}_o = \mathbf{I} - \mathbf{P} = O^\top (\mathbf{I} - \Lambda) O$. Mimicking the arguments of Lemma C.1.2 the proof follows. \square

Lemma C.1.3. $\|\mathbf{P} \zeta^{(q)}\|_\infty = \mathcal{O}\left(\tau \sqrt{2 \log(2g)}\right), \|\mathbf{P}_o \zeta^{(q)}\|_\infty = \mathcal{O}\left(\tau \sqrt{2 \log(2d)}\right)$ w.h.p. for any $q \in [k]$.

Specifically,

$$P\left[\|\mathbf{P} \zeta^{(q)}\|_\infty > 3\tau \sqrt{2 \log(2g)}\right] \leq e^{-4 \log(2g)}$$

$$P\left[\|\mathbf{P}_o \zeta^{(q)}\|_\infty > 3\tau \sqrt{2 \log(2d)}\right] \leq e^{-4 \log(2d)}$$

Proof. First we will bound $\mathbb{E}(\|\mathbf{P} \zeta\|_\infty)$ for any $\zeta \sim N(0, \tau^2 \mathbf{I}_D)$.

Let, $Z = \|\mathbf{P} \zeta\|_\infty$ equivalently $Z = \max_{i \in [D]} \max\{(\mathbf{P} \zeta)_i, -(\mathbf{P} \zeta)_i\}$.

We can define $X_i = (\mathbf{P} \zeta)_i \sim N(0, \tau^2 p_{ii})$ and $X_{2i} = -(\mathbf{P} \zeta)_i \sim N(0, \tau^2 p_{ii})$. Where p_{ii} is the i^{th} diagonal element of \mathbf{P} . From Lemma C.1.2 there exists an indexing set L such that $p_{ii} = 0 \forall i \in L$. This implies $X_i = 0, X_{2i} = 0 \forall i \in L$. Hence, $Z = \max_{i \in [D] \setminus L} \max\{(\mathbf{P} \zeta)_i, -(\mathbf{P} \zeta)_i\}$.

By Jensen's inequality for any $t > 0$, we have

$$\exp\{t\mathbb{E}[Z]\} \leq \mathbb{E} \exp\{tZ\} = \mathbb{E} \max_{i \in [2D] \setminus (L \cup 2L)} \exp\{tX_i\} \leq \sum_{i \in [2D] \setminus (L \cup 2L)} \mathbb{E}[\exp\{tX_i\}]$$

$$\stackrel{(i)}{=} \sum_{i \in [D] \setminus L} 2 \exp\{t^2 p_{ii} \tau^2 / 2\} \stackrel{(ii)}{\leq} 2g \exp\{t^2 \tau^2 / 2\}$$

(i) uses the moment-generating function for Gaussian. (ii) uses the fact that $p_{ii} < 1$ and $|[D] \setminus L| = g$ from Lemma C.1.2.

$$\mathbb{E}[Z] \leq \frac{\log 2g}{t} + \frac{t\tau^2}{2}$$

Setting $t = \frac{\sqrt{2\log 2g}}{\tau}$ we get $\mathbb{E}[Z] \leq \tau\sqrt{2\log(2g)}$. For any $u > 0$ we have

$$\begin{aligned} P[Z > \tau\sqrt{2\log(2g)} + u] &\leq P(Z > \mathbb{E}Z + u) \\ &\stackrel{(iii)}{\leq} e^{-\frac{u^2}{2\tau^2 \max_i p_{ii}}} \leq e^{-\frac{u^2}{2\tau^2}} \end{aligned}$$

Where, (iii) uses Borell-TIS inequality Adler and Taylor (2007).

Setting, $u = 2\tau\sqrt{2\log(2g)}$ we get $P(Z > 3\tau\sqrt{2\log(2g)}) \leq e^{-4\log(2g)}$. Finally $P(\|\mathbf{P}\zeta\|_\infty < 3\tau\sqrt{2\log(2g)}) = 1 - e^{-4\log(2g)}$. One can get an identical result for $\mathbf{P}_o\zeta^{(g)}$ replacing $g \rightarrow d$ and using Corollary C.1.1 instead of C.1.2. \square

Lemma C.1.4. For any $\delta > 0$ and $i \neq j \in [k]$; $|\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}}|, |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| \leq \delta$ w.h.p. Specifically,

$$\begin{aligned} P\left[\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}} \geq \delta\right] &\leq 4e^{-\frac{2\delta}{\sqrt{8g}\tau^2}}; P\left[\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}} \leq -\delta\right] \leq 4e^{-\frac{2\delta}{\sqrt{8g}\tau^2}} \\ P\left[\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o} \geq \delta\right] &\leq 4e^{-\frac{2\delta}{\sqrt{8d}\tau^2}}; P\left[\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o} \leq -\delta\right] \leq 4e^{-\frac{2\delta}{\sqrt{8d}\tau^2}} \end{aligned}$$

Proof. Let, $Z_1 = \mathbf{P}\zeta^{(i)} \sim N(0, \tau^2\mathbf{P})$ and $Z_2 = \zeta^{(j)} \sim N(0, \tau^2\mathbf{I}_D)$ for $i \neq j$.

$$\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}} = \zeta^{(i)\top} \mathbf{P} \zeta^{(j)} = (\mathbf{P}\zeta^{(i)})^\top \zeta^{(j)} = Z_1^\top Z_2$$

The inner product can be expressed as: $Z_1^\top Z_2 = \|Z_1\| \cdot \frac{Z_1^\top}{\|Z_1\|} Z_2$.

Notice that, $\frac{Z_1^\top}{\|Z_1\|} Z_2 \sim N(0, \tau^2)$ and is independent of Z_1 . Hence, using the standard Gaussian- tail bounds for $\delta > 0, C \geq 1$.

$$P\left(\frac{Z_1^\top}{\|Z_1\|} Z_2 > \frac{\delta}{C}\right) \leq e^{-\frac{\delta^2}{2C^2\tau^2}} \quad (6)$$

We can also get a concentration inequality for $\|Z_1\|$, using Ledoux (2006) formula (3.5).

$$\begin{aligned} P(\|Z_1\| > C) &\leq 4e^{-\frac{C^2}{8\mathbb{E}(\|Z_1\|^2)}} \\ &\stackrel{(i)}{=} 4e^{-\frac{C^2}{8g\tau^2}} \end{aligned} \quad (7)$$

Where (i) follows from the fact that $\|Z_1\|^2 \tau^{-2} \sim \chi^2(g)$, hence $\mathbb{E}(\|Z_1\|^2) = g\tau^2$.

Combining, Eq. (7) and Eq. (6), we have:

$$P(Z_1^\top Z_2 > \delta) \leq 4 \exp\left(-\frac{\delta^2}{2C^2\tau^2} - \frac{C^2}{8g\tau^2}\right)$$

Setting $C^2 = \sqrt{4g}\delta$ we get :

$$P(Z_1^\top Z_2 > \delta) \leq 4 \exp\left(-\frac{\delta}{\sqrt{4g}\tau^2}\right)$$

One can convince themselves, that $Z_1^\top Z_2$ has a symmetric distribution it is a product of Gaussians. Hence, we also have:

$$P(Z_1^\top Z_2 < -\delta) \leq 4 \exp\left(-\frac{\delta}{\sqrt{4g}\tau^2}\right)$$

Using the union bound,

$$P(|Z_1^\top Z_2| < \delta) \geq 1 - 8 \exp\left(-\frac{\delta}{\sqrt{4g}\tau^2}\right)$$

The proof for $\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}$ is identical. \square

Lemma C.1.5. For all $r \in [k]$ and $j \neq r$ consider the following events hold with high probability:

1. $\frac{g\tau^2}{2} < \zeta^{(r)\top} \mathbf{P} \zeta^{(r)} = \|\mathbf{P} \zeta^{(r)}\|^2 < \frac{13g\tau^2}{8}$
2. $\frac{D\tau^2}{2} < \|\zeta^{(r)}\|^2 < \frac{13D\tau^2}{8}$
3. $\|\mathbf{P} \zeta^{(r)}\|_\infty < 3\tau\sqrt{2\log(2g)}$
4. $\max_{r \neq j} |\langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}}| \leq \frac{g\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13g\tau^2}{8}} = c_3(g, \tau, k) \asymp \left(\frac{g\tau^2}{k}\right)$

Specifically,

$$P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}}| + \epsilon_\zeta^\perp\right) < \frac{1}{10} \cdot [\zeta^{(r)\top} \mathbf{P} \zeta^{(r)} - \epsilon_\zeta^\perp]\right) \geq 1 - \delta_1$$

$$\delta_1 = \left(4k(k-1) \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right) + 2k \exp\left(-\frac{g}{16}\right) + 2k \exp\left(-\frac{D}{16}\right) + k \exp(-4 \log 2g)\right)$$

With all events (1-4) being satisfied.

Proof. If the above events hold true then $\epsilon_\zeta^\perp = \bar{\omega} \|\mathbf{P} \zeta\| \leq \sqrt{\frac{13g\tau^2}{8}}$.

Using Lemma C.1.5 with $\delta = c_3(g, \tau, k)$ and due to the assumption of $c_3(g, \tau, k) > 0$ we have:

$$P\left[\langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}} \geq c_3(g, \tau, k)\right] \leq 4 \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right) \quad (8)$$

$$P\left[\langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}} \leq -c_3(g, \tau, k)\right] \leq 4 \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right)$$

Let,

$$P(A) = P\left\{\bigcup_{r \neq j} \left[\langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}} \geq c_3 \cup \langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}} \leq -c_3\right]\right\} \stackrel{(i)}{\leq} 4k(k-1) \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right)$$

$$P(B) = P\left\{\bigcup_{r \in [k]} \left[\zeta^{(r)\top} \mathbf{P} \zeta^{(r)} \geq \frac{13g\tau^2}{8} \cup \zeta^{(r)\top} \mathbf{P} \zeta^{(r)} \leq \frac{g\tau^2}{2}\right]\right\} \stackrel{(ii)}{\leq} 2k \exp\left(-\frac{g}{16}\right)$$

$$P(C) = P\left\{\bigcup_{r \in [k]} \left[\|\zeta^{(r)}\|^2 \geq \frac{13D\tau^2}{8} \cup \|\zeta^{(r)}\|^2 \leq \frac{D\tau^2}{2}\right]\right\} \stackrel{(iii)}{\leq} 2k \exp\left(-\frac{D}{16}\right)$$

$$P(D) = P\left\{\bigcup_{r \in [k]} \left[\|\mathbf{P} \zeta^{(r)}\|_\infty > 3\tau\sqrt{2\log(2g)}\right]\right\} \stackrel{(iv)}{\leq} k \exp(-4 \log 2g)$$

Where (i) uses Eq. (8); (ii), (iii) uses Lemma C.1.1 and (iv) uses Lemma C.1.3, along with the Union bound. The probability that all of the high-probability events happen can be computed by the union bound: $1 - P(A \cup B \cup C \cup D) \geq 1 - \left(4k(k-1) \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right) + 2k \exp\left(-\frac{g}{16}\right) + 2k \exp\left(-\frac{D}{16}\right) + k \exp(-4 \log 2g)\right)$.

Hence,

$$P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}}| + \epsilon_\zeta^\perp\right) < \frac{1}{10} \cdot [\zeta^{(r)\top} \mathbf{P} \zeta^{(r)} - \epsilon_\zeta^\perp]\right) \geq P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}}|\right) < c_3(g, \tau, k)\right)$$

$$\geq 1 - \left(4k(k-1) \exp\left(-\frac{c_3}{\sqrt{4g\tau^2}}\right) + 2k \exp\left(-\frac{g}{16}\right) + 2k \exp\left(-\frac{D}{16}\right) + k \exp(-4 \log 2g)\right)$$

□

Lemma C.1.6. For some $q \in [k]$, let $V_q = \sum_{r \in [k]} y^{(r)} y^{(q)} \langle \tilde{\mu}^{(q)}, \zeta^{(r)} \rangle_{\mathbf{P}_o}$. Then w.h.p. $|V_q| < \frac{1}{10} \cdot c_4(d, \tau, k, p)$, where $c_4(d, \tau, k, p) = d - \Delta' - k(p + \Delta') + \frac{9}{10} \cdot \left(\frac{d\tau^2}{2} - \sqrt{\frac{13d\tau^2}{8}}(\sigma\sqrt{2\ln d} + 1)\right)$

Proof. As each $\zeta^{(r)}$ $\overset{i.i.d.}{\sim} N(0, \tau^2 \mathbf{I}_D)$ and $\|\tilde{\boldsymbol{\mu}}\|^2 = d$, we have $y^{(r)} y^{(q)} \langle \tilde{\boldsymbol{\mu}}^{(q)}, \zeta^{(r)} \rangle_{\mathbf{P}_o} \overset{i.i.d.}{\sim} N(0, \tau^2 d)$. Consequently, $V_q \sim N(0, \tau^2 dk)$.

Using gaussian tail bound we have:

$$P\left(|V_q| \geq \frac{1}{10} \cdot c_4(d, \tau, k, p)\right) \leq 2 \exp\left(-\frac{c_4^2}{20\tau^2 dk}\right)$$

□

Lemma C.1.7. For all $r \in [k]$ and $j \neq r$ consider the following events hold with high probability:

1. $\frac{d\tau^2}{2} < \zeta^{(r)\top} \mathbf{P}_o \zeta^{(r)} = \|\mathbf{P}_o \zeta^{(r)}\|^2 < \frac{13d\tau^2}{8}$
2. $\frac{D\tau^2}{2} < \|\zeta^{(r)}\|^2 < \frac{13D\tau^2}{8}$
3. $\|\mathbf{P}_o \zeta^{(r)}\|_\infty < 3\tau\sqrt{2\log(2d)}$
4. $|V_r| < \frac{1}{10} \cdot c_4(d, \tau, k, p)$
5. $\max_{r \neq j} |\langle \zeta^{(r)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| \leq \frac{d\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2\ln d + 1}) = c_5(d, \tau, k) \asymp \left(\frac{d\tau^2}{k}\right)$

Specifically,

$$P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| + \epsilon_\zeta^{\text{II}}\right) < \frac{1}{10} \cdot [\zeta^{(r)\top} \mathbf{P}_o \zeta^{(r)} - \epsilon_\zeta^{\text{II}}]\right) \geq 1 - \delta_2$$

$$\delta_2 = \left(4k(k-1) \exp\left(-\frac{c_5}{\sqrt{4d\tau^2}}\right) + 2k \exp\left(-\frac{d}{16}\right) + 2k \exp\left(-\frac{D}{16}\right) + k \exp(-4 \log 2d) + 2k \exp\left(-\frac{c_4^2}{20\tau^2 dk}\right)\right)$$

With all events (1-5) being satisfied.

Proof. If the above events hold true then, $\epsilon_\zeta^{\text{II}} = (\sqrt{2}\sigma \ln d + \bar{\omega}) \|\mathbf{P}_o \zeta\| \leq \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2\ln d + 1})$. We can mimic the proof of Lemma C.1.5 replacing $\mathbf{P} \rightarrow \mathbf{P}_o, g \rightarrow d$ and using the corresponding probability bounds. Use Lemma C.1.4 with $\delta = c_5(d, \tau, k)$ to get bound for event A. In addition to the events A, B, C, D defined in Lemma C.1.5 we need to subtract the probability of an event E from our union bound so that the event 4 of this lemma is satisfied w.h.p.

$$P(E) = P\left\{\bigcup_{r \in [k]} |V_r| > \frac{1}{10} \cdot c_4(d, \tau, k, p)\right\} \stackrel{(i)}{\leq} 2k \exp\left(-\frac{c_4^2}{20\tau^2 dk}\right)$$

Where (i) is a consequence of Lemma C.1.6. Using the union bound $1 - P(A) - P(B) \dots - P(E)$, we have:

$$\begin{aligned} & P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| + \epsilon_\zeta^{\text{II}}\right) < \frac{1}{10} \cdot [\zeta^{(r)\top} \mathbf{P}_o \zeta^{(r)} - \epsilon_\zeta^{\text{II}}]\right) \geq P\left(k \left(\max_{i \neq j} |\langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o}| < c_5(g, \tau, k)\right)\right) \\ & \geq 1 - \left(4k(k-1) \exp\left(-\frac{c_5}{\sqrt{4d\tau^2}}\right) + 2k \exp\left(-\frac{d}{16}\right) + 2k \exp\left(-\frac{D}{16}\right) + k \exp(-4 \log 2d) + 2k \exp\left(-\frac{c_4^2}{20\tau^2 dk}\right)\right) \end{aligned}$$

□

C.2 Generalisation

Now onwards, define the index set of positive/negative second layer weights as $J+ = \{j : v_j > 0\}, J- = \{j : v_j < 0\}$.

Lemma C.2.1. We have $\Delta' \leq \frac{d}{21}$.

Proof. Recall that $k(p + \Delta' + 1) \leq \frac{d - \Delta' + 1}{10}$. Since $k \geq 2$ and $p \geq 0$ it implies that $2(\Delta' + 1) \leq \frac{d - \Delta' + 1}{10}$. Hence, $\Delta' \leq \frac{d-19}{21} \leq \frac{d}{21}$. □

The proofs of the following four lemmas C.2.2 to C.3.2 can be directly translated from some of the lemmas proved in Frei et al. (2023) to our setting. The reader needs to just replace the $\mathbf{x}, \boldsymbol{\xi}, \Delta$ terms in their proofs to $\tilde{\mathbf{x}}, \tilde{\boldsymbol{\xi}}, \Delta'$ in our setting. Under the hood, one can mimic the proofs by just replacing such terms because of the properties established in Lemma C.0.1. For the properties to hold, we need $\|\boldsymbol{\zeta}^{(r)} + \boldsymbol{\omega}\| \leq \bar{\zeta}$ for any $r \in [k]$. Note that this gets satisfied if $\|\boldsymbol{\zeta}^{(r)}\| < \frac{13D\tau^2}{8}$ which is true with probability $1 - \exp(-D/16)$ (Lemma C.1.1). So, our previously mentioned lemmas will hold w.h.p. for brevity's sake, we omit their proofs. However, if one wants, they can refer to the analogous lemmas in Frei et al. (2023) cited accordingly.

Lemma C.2.2 (Translated from Lemma A.10 Frei et al. (2023)). *If \mathcal{S} satisfies assumptions 3.2, then for all $q \in Q$ we have the following with probability $1 - \exp(-D/16)$.*

$$\max \left\{ \sum_{i \in I^{(q)}} \sum_{j \in J_+} v_j^2 \lambda_i \phi'_{i,j}, \sum_{i \in I^{(q)}} \sum_{j \in J_-} v_j^2 \lambda_i \phi'_{i,j} \right\} \leq \frac{1}{(1 - 2c')(d - \Delta' + 1)}.$$

Lemma C.2.3 (Translated from Lemma A.11 Frei et al. (2023)). *If \mathcal{S} satisfies assumptions 3.2, then for all $q \in Q_+$ we have the following with probability $1 - \exp(-D/16)$.*

$$\sum_{i \in I^{(q)}} \sum_{j \in J_+} v_j^2 \lambda_i \phi'_{i,j} \geq \left(1 - \frac{c'}{1 - 2c'}\right) \frac{1}{3d + \Delta' + 1},$$

and for all $q \in Q_-$ we have

$$\sum_{i \in I^{(q)}} \sum_{j \in J_-} v_j^2 \lambda_i \phi'_{i,j} \geq \left(1 - \frac{c'}{1 - 2c'}\right) \frac{1}{3d + \Delta' + 1}.$$

Proof of Theorem 4.1. Working is shown for the case when $r \in Q_+$. The proof for $r \in Q_-$ is similar. By Eq. (4), for every $j \in [w]$ we have:

$$\begin{aligned} \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j &= \left(\sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} \right) + \sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} \\ &= \sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} (\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \\ &= \left(\sum_{i \in I^{(r)}} \lambda_i v_j \phi'_{i,j} (\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \right) + \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i y_i v_j \phi'_{i,j} (\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1). \end{aligned} \quad (9)$$

Now,

$$\mathcal{N}_\theta(\tilde{\mathbf{x}}) = \sum_{j \in [w]} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) \geq \sum_{j \in J_+} v_j (\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) + \sum_{j \in J_-} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j). \quad (10)$$

By Eq. (9) we have

$$\begin{aligned} \sum_{j \in J_+} v_j (\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) &= \sum_{j \in J_+} \left[\left(\sum_{i \in I^{(r)}} \lambda_i v_j^2 \phi'_{i,j} (\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \right) + \sum_{q \in Q \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i y_i v_j^2 \phi'_{i,j} (\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \right] \\ &\geq \sum_{j \in J_+} \left[\left(\sum_{i \in I^{(r)}} \lambda_i v_j^2 \phi'_{i,j} (d - \Delta' + 1) \right) - \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i v_j^2 \phi'_{i,j} (p + \Delta' + 1) \right] \\ &= \left((d - \Delta' + 1) \sum_{i \in I^{(r)}} \sum_{j \in J_+} \lambda_i v_j^2 \phi'_{i,j} \right) \\ &\quad - \left((p + \Delta' + 1) \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \sum_{j \in J_+} \lambda_i v_j^2 \phi'_{i,j} \right). \end{aligned}$$

By Lemma C.2.2 and Lemma C.2.3 the above is at least

$$\begin{aligned}
 & (d - \Delta' + 1) \left(1 - \frac{c'}{1 - 2c'}\right) \frac{1}{3d + \Delta' + 1} - (p + \Delta' + 1)k \cdot \frac{1}{(1 - 2c')(d - \Delta' + 1)} \\
 &= \left(1 - \frac{c'}{1 - 2c'}\right) \frac{d - \Delta' + 1}{3d + \Delta' + 1} - (p + \Delta' + 1)c' \cdot \frac{d - \Delta' + 1}{p + \Delta' + 1} \cdot \frac{1}{(1 - 2c')(d - \Delta' + 1)} \\
 &= \left(1 - \frac{c'}{1 - 2c'}\right) \frac{d - \Delta' + 1}{3d + \Delta' + 1} - \frac{c'}{1 - 2c'}. \tag{11}
 \end{aligned}$$

Likewise, we have

$$\begin{aligned}
 \sum_{j \in J_-} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) &= \sum_{j \in J_-} v_j \phi \left(\left(\sum_{i \in I^{(r)}} \lambda_i v_j \phi'_{i,j}(\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \right) + \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i y_i v_j \phi'_{i,j}(\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}} + 1) \right) \\
 &\geq \sum_{j \in J_-} v_j \phi \left(\left(\sum_{i \in I^{(r)}} \lambda_i v_j \phi'_{i,j}(d - \Delta' + 1) \right) + \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i |v_j| \phi'_{i,j}(p + \Delta' + 1) \right) \\
 &\geq \sum_{j \in J_-} v_j \phi \left(\sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i |v_j| \phi'_{i,j}(p + \Delta' + 1) \right) \\
 &= - \sum_{j \in J_-} \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \lambda_i v_j^2 \phi'_{i,j}(p + \Delta' + 1) \\
 &= -(p + \Delta' + 1) \sum_{q \in [k] \setminus \{r\}} \sum_{i \in I^{(q)}} \sum_{j \in J_-} \lambda_i v_j^2 \phi'_{i,j}.
 \end{aligned}$$

By Lemma C.2.2 the above is at least

$$\begin{aligned}
 -(p + \Delta' + 1)k \cdot \frac{1}{(1 - 2c')(d - \Delta' + 1)} &= -(p + \Delta' + 1)c' \cdot \frac{d - \Delta' + 1}{p + \Delta' + 1} \cdot \frac{1}{(1 - 2c')(d - \Delta' + 1)} \\
 &= -\frac{c'}{1 - 2c'}. \tag{12}
 \end{aligned}$$

Combining Eq. (10), (11), and (12), we get

$$\mathcal{N}_\theta(\tilde{\mathbf{x}}) \geq \left(1 - \frac{c'}{1 - 2c'}\right) \frac{d - \Delta' + 1}{3d + \Delta' + 1} - \frac{c'}{1 - 2c'} - \frac{c'}{1 - 2c'}.$$

Using $c' \leq \frac{1}{10}$ and $\Delta' \leq d/21$ (which holds by Lemma C.2.1), the above is at least

$$\frac{7}{8} \cdot \frac{d - \Delta' + 1}{3d + \Delta' + 1} - \frac{2}{8} \geq \frac{7}{8} \cdot \frac{d - \Delta'}{3d + \Delta'} - \frac{2}{8}.$$

By Lemma C.2.1, the displayed equation is at least

$$\frac{7}{8} \cdot \frac{d - d/21}{3d + d/21} - \frac{2}{8} = \frac{7}{8} \cdot \frac{5}{16} - \frac{2}{8} > 0.$$

□

C.3 Existence of Adversarial perturbations

Lemma C.3.1 (Translated from Lemma D.2 Frei et al. (2023)). *Suppose \mathcal{S} satisfies assumption 3.2. Let $\tilde{\mathbf{x}} \in \mathbb{R}^D$ and $r \in Q$ such that $\tilde{\mathbf{x}} = \tilde{\boldsymbol{\mu}}^{(r)} + \tilde{\boldsymbol{\xi}}$ is a 'nice example'. Then, $|\mathcal{N}_\theta(\tilde{\mathbf{x}})| \leq 2$ with probability $1 - \exp(-D/16)$.*

Lemma C.3.2 (Translated from Lemma D.3 Frei et al. (2023)). *Suppose \mathcal{S} satisfies assumption 3.2. Let $\tilde{\mathbf{x}} \in \mathbb{R}^D$ and $r \in Q$ such that $\tilde{\mathbf{x}} = \tilde{\boldsymbol{\mu}}^{(r)} + \tilde{\boldsymbol{\xi}}$ is a 'nice example'. Then, for all $j \in [w]$ we have the following with probability $1 - \exp(-D/16)$:*

$$\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j \geq - \sum_{i \in [n]} \lambda_i |v_j| \phi'_{i,j}(2\Delta' + p + 1).$$

Lemma C.3.3. Suppose \mathcal{S} satisfies assumptions 3.2.

Let $\mathbf{u}_\perp = \sum_{q \in [k]} y^{(q)} \mathbf{P}_o \boldsymbol{\zeta}^{(q)}$, $\mathbf{u}_\parallel = \sum_{q \in [k]} y^{(q)} (\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)})$. Then $\mathbf{u}_\perp, \mathbf{u}_\parallel$ satisfies the following with probability at least $1 - \delta_1, 1 - \delta_2$ respectively.

For every $j \in J_+$ we have :

$$\mathbf{w}_j^\top \mathbf{u}_\perp \geq \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) > 0; \quad \mathbf{w}_j^\top \mathbf{u}_\parallel \geq \frac{9}{10} \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_4(d, \tau, k, p) > 0,$$

For every $j \in J_-$ we have :

$$\mathbf{w}_j^\top \mathbf{u}_\perp \leq \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) < 0; \quad \mathbf{w}_j^\top \mathbf{u}_\parallel \leq \frac{9}{10} \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_4(d, \tau, k, p) < 0,$$

Where, $c_4(d, \tau, k, p) = d - \Delta' + \frac{9}{10} \cdot \left(\frac{d\tau^2}{2} - \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2} \ln d + 1) \right) - k(p + \Delta')$ and $c_6(g, \tau) = \frac{9}{10} \cdot \left(\frac{g\tau^2}{2} - \sqrt{\frac{13g\tau^2}{8}} \right)$. δ_1, δ_2 are defined in Lemma C.1.5 and Lemma C.1.7 respectively.

Proof. Note that from Lemma C.3.4 the constants $c_4(d, \tau, k, p), c_6(g, \tau)$ are positive.

On manifold perturbation

For $\mathbf{w}_j^\top \mathbf{u}_\parallel$ we have:

$$\begin{aligned} \mathbf{w}_j^\top \sum_{r \in [k]} y^{(r)} (\tilde{\boldsymbol{\mu}}^{(r)} + \mathbf{P}_o \boldsymbol{\zeta}^{(r)}) &= \sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} \tilde{\mathbf{x}}_i^\top \sum_{r \in [k]} y^{(r)} (\tilde{\boldsymbol{\mu}}^{(r)} + \mathbf{P}_o \boldsymbol{\zeta}^{(r)}) \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i y^{(q)} v_j \phi'_{i,j} \tilde{\mathbf{x}}_i^\top \left(y^{(q)} (\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)}) + \sum_{r \in [k] \setminus \{q\}} y^{(r)} (\tilde{\boldsymbol{\mu}}^{(r)} + \mathbf{P}_o \boldsymbol{\zeta}^{(r)}) \right) \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left((y^{(q)})^2 \tilde{\mathbf{x}}_i^\top (\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)}) + \sum_{r \in [k] \setminus \{q\}} y^{(q)} y^{(r)} \tilde{\mathbf{x}}_i^\top (\tilde{\boldsymbol{\mu}}^{(r)} + \mathbf{P}_o \boldsymbol{\zeta}^{(r)}) \right) \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left[(\tilde{\boldsymbol{\mu}}^{(q)})^\top \tilde{\boldsymbol{\mu}}^{(q)} + \tilde{\boldsymbol{\xi}}_i^\top \boldsymbol{\mu}^{(q)} + (\tilde{\boldsymbol{\mu}}^{(q)})^\top \mathbf{P}_o \boldsymbol{\zeta}^{(q)} + (\mathbf{M} \boldsymbol{\xi}_i + \boldsymbol{\omega}_i)^\top \mathbf{P}_o \boldsymbol{\zeta}^{(q)} + \|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|^2 \right. \\ &\quad \left. \sum_{r \in [k] \setminus \{q\}} y^{(q)} y^{(r)} \left((\tilde{\boldsymbol{\mu}}^{(q)})^\top \tilde{\boldsymbol{\mu}}^{(r)} + \tilde{\boldsymbol{\xi}}_i^\top \tilde{\boldsymbol{\mu}}^{(r)} + \langle \tilde{\boldsymbol{\mu}}^{(q)}, \boldsymbol{\zeta}^{(r)} \rangle_{\mathbf{P}_o} + \langle \boldsymbol{\zeta}^{(q)}, \boldsymbol{\zeta}^{(r)} \rangle_{\mathbf{P}_o} + (\mathbf{M} \boldsymbol{\xi}_i + \boldsymbol{\omega}_i)^\top \mathbf{P}_o \boldsymbol{\zeta}^{(r)} \right) \right]. \quad (13) \end{aligned}$$

For $j \in J_+$ the above is at least

$$\begin{aligned} &\sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(V_q + d - \Delta' - \epsilon_\zeta^\parallel + \|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|^2 - \sum_{r \in [k] \setminus \{q\}} (p + \Delta' + \langle \boldsymbol{\zeta}^{(q)}, \boldsymbol{\zeta}^{(r)} \rangle_{\mathbf{P}_o} + \epsilon_\zeta^\parallel) \right) \\ &\geq \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(V_q + d - \Delta' - \epsilon_\zeta^\parallel + \|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|^2 - k(p + \Delta' + \max_{i \neq j} \langle \boldsymbol{\zeta}^{(i)}, \boldsymbol{\zeta}^{(j)} \rangle_{\mathbf{P}_o} + \epsilon_\zeta^\parallel) \right) \\ &\stackrel{(i)}{\geq} \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(V_q + d - \Delta' + \frac{9}{10} \cdot (\|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|^2 - \epsilon_\zeta^\parallel) - k(p + \Delta') \right) \\ &\stackrel{(ii)}{\geq} \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(V_q + d - \Delta' + \frac{9}{10} \cdot \left(\frac{d\tau^2}{2} - \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2} \ln d + 1) \right) - k(p + \Delta') \right) \\ &\stackrel{(iii)}{\geq} \frac{9}{10} c_4(d, \tau, k, p) \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} > 0 \end{aligned}$$

We use the Lemma C.1.7 for the following:

(i) we used that $\max_{i \neq j} \langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}_o} < \frac{1}{10} \cdot \|\mathbf{P}_o \zeta^{(q)}\|$ w.h.p.

(ii) we used that $\|\mathbf{P}_o \zeta^{(q)}\| > \frac{d\tau^2}{2}$ and $\epsilon_\zeta^u \leq \sqrt{\frac{13d\tau^2}{8}}(\sigma\sqrt{2} \ln d + 1)$ w.h.p.

(iii) we used that $V_q > -\frac{1}{10}c_4(d, \tau, k, p)$ w.h.p.

Similarly, for $j \in J_-$, Eq. (13) is at most

$$\frac{9}{10}c_4(d, \tau, k, p) \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} < 0$$

Off manifold perturbation

For $\mathbf{w}_j^\top \mathbf{u}_\perp$ we have:

$$\begin{aligned} \mathbf{w}_j^\top \sum_{r \in [k]} y^{(r)} \mathbf{P} \zeta^{(r)} &= \sum_{i \in [n]} \lambda_i y_i v_j \phi'_{i,j} \tilde{\mathbf{x}}_i^\top \sum_{r \in [k]} y^{(r)} \mathbf{P} \zeta^{(r)} \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i y^{(q)} v_j \phi'_{i,j} \tilde{\mathbf{x}}_i^\top \left(y^{(q)} \mathbf{P} \zeta^{(q)} + \sum_{r \in [k] \setminus \{q\}} y^{(r)} \mathbf{P} \zeta^{(r)} \right) \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left((y^{(q)})^2 \tilde{\mathbf{x}}_i^\top \mathbf{P} \zeta^{(q)} + \sum_{r \in [k] \setminus \{q\}} y^{(q)} y^{(r)} \tilde{\mathbf{x}}_i^\top \mathbf{P} \zeta^{(r)} \right) \\ &= \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left[\omega_i^\top \mathbf{P} \zeta^{(q)} + \|\mathbf{P} \zeta^{(q)}\|^2 + \sum_{r \in [k] \setminus \{q\}} y^{(q)} y^{(r)} \left(\langle \zeta^{(q)}, \zeta^{(r)} \rangle_{\mathbf{P}} + \omega_i^\top \mathbf{P} \zeta^{(r)} \right) \right]. \end{aligned} \quad (14)$$

For $j \in J_+$ the above is at least

$$\begin{aligned} &\sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(\|\mathbf{P} \zeta^{(q)}\|^2 - \epsilon_\zeta^\perp - \sum_{r \in [k] \setminus \{q\}} (\langle \zeta^{(q)}, \zeta^{(r)} \rangle_{\mathbf{P}} + \epsilon_\zeta^\perp) \right) \\ &\geq \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(\|\mathbf{P} \zeta^{(q)}\|^2 - \epsilon_\zeta^u - k(\max_{i \neq j} \langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}} + \epsilon_\zeta^\perp) \right) \\ &\stackrel{(iv)}{\geq} \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(\frac{9}{10} \cdot (\|\mathbf{P} \zeta^{(q)}\|^2 - \epsilon_\zeta^\perp) \right) \\ &\stackrel{(v)}{\geq} \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} \left(\frac{9}{10} \cdot \left(\frac{g\tau^2}{2} - \sqrt{\frac{13g\tau^2}{8}} \right) \right) = \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) > 0 \end{aligned}$$

We use the Lemma C.1.5 for the following:

(iv) we used that $\max_{i \neq j} \langle \zeta^{(i)}, \zeta^{(j)} \rangle_{\mathbf{P}} < \frac{1}{10} \cdot \|\mathbf{P} \zeta^{(q)}\|$ w.h.p.

(v) we used that $\|\mathbf{P} \zeta^{(q)}\| > \frac{g\tau^2}{2}$ and $\epsilon_\zeta^\perp \leq \sqrt{\frac{13g\tau^2}{8}}$ w.h.p.

Similarly, for $j \in J_-$, Eq. (14) is at most

$$c_6(g, \tau) \sum_{q \in [k]} \sum_{i \in I^{(q)}} \lambda_i v_j \phi'_{i,j} < 0$$

□

Lemma C.3.4.

$$\frac{18}{21} \cdot \frac{g\tau^2}{2} \leq c_6(g, \tau) \leq \frac{9}{10} \cdot \frac{g\tau^2}{2}$$

and,

$$\frac{18}{21} \cdot d \left(1 + \frac{\tau^2}{2} \right) \leq c_4(d, \tau, k, p) \leq d \left(1 + \frac{9\tau^2}{20} \right)$$

Proof. From assumption 3.2, $c_3(g, \tau, k) > 0$:

$$\frac{g\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13g\tau^2}{8}} > 0 \text{ but } k \geq 2 \text{ implying } \frac{g\tau^2}{42} \geq \sqrt{\frac{13g\tau^2}{8}}.$$

$$\frac{9}{10} \cdot \frac{g\tau^2}{2} \geq c_6(g, \tau) = \frac{9}{10} \cdot \left(\frac{g\tau^2}{2} - \sqrt{\frac{13g\tau^2}{8}} \right) \geq \frac{18}{21} \cdot \frac{g\tau^2}{2}$$

From assumption 3.2, $c_5(d, \tau, k) > 0$:

$$\frac{d\tau^2}{20k} - \frac{10k+1}{10k} \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2\ln d + 1}) > 0, \text{ and } k \geq 2 \text{ implying}$$

$$\frac{d\tau^2}{42} \geq \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2\ln d + 1}) \quad (15)$$

$$\begin{aligned} c_4(d, \tau, k, p) &= d - \Delta' + \frac{9}{10} \cdot \left(\frac{d\tau^2}{2} - \sqrt{\frac{13d\tau^2}{8}} (\sigma\sqrt{2\ln d + 1}) \right) - k(p + \Delta') \\ &\stackrel{(i)}{\geq} d - \Delta' + \frac{18}{21} \cdot \frac{d\tau^2}{2} - k(p + \Delta' + 1) \\ &\stackrel{(ii)}{\geq} \frac{9}{10} \cdot (d - \Delta') + \frac{1}{10} + \frac{18}{21} \cdot \frac{d\tau^2}{2} \\ &\stackrel{(iii)}{\geq} \frac{18}{21} \cdot d + \frac{18}{21} \cdot \frac{d\tau^2}{2} = \frac{18}{21} \cdot d \left(1 + \frac{\tau^2}{2} \right) \end{aligned}$$

For (i) use Eq. (15); for (ii) use 3.2 $\frac{k(p+\Delta'+1)}{d-\Delta'+1} \leq \frac{1}{10}$; for (iii) use Lemma C.2.1. The upper-bound for $c_4(d, \tau, k, p)$ is straight forward and can be obtained by omitting the negative terms. \square

Lemma C.3.5. *Suppose \mathcal{S} satisfies assumption 3.2. Let $\tilde{\mathbf{x}} \in \mathbb{R}^D$ and $r \in [k]$ such that $\tilde{\mathbf{x}} = \tilde{\boldsymbol{\mu}}^{(r)} + \tilde{\boldsymbol{\xi}}$ is a 'nice example'. Let $\mathbf{z}_\perp = \eta^\perp \sum_{q \in [k]} y^{(q)} \mathbf{P}\boldsymbol{\zeta}^{(q)}$, $\mathbf{z}_\parallel = \eta^\parallel \sum_{q \in [k]} y^{(q)} \tilde{\boldsymbol{\mu}}^{(q)}$, for $\eta^\perp \geq \frac{2\Delta' + p + 1}{c_6(g, \tau)}$, $\eta^\parallel \geq \frac{10}{9} \cdot \frac{2\Delta' + p + 1}{c_4(d, \tau, k, p)}$. Then, $\mathbf{z}_\perp, \mathbf{z}_\parallel$ satisfies the following with probability at least $1 - \delta_1, 1 - \delta_2$ respectively.*

For all $j \in J_-$ we have:

$$\mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}_\perp) + b_j \geq 0; \quad \mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}_\parallel) + b_j \geq 0$$

and for all $j \in J_+$:

$$\mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}_\perp) + b_j \geq 0; \quad \mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}_\parallel) + b_j \geq 0,$$

δ_1, δ_2 are defined in Lemma C.1.5 and Lemma C.1.7 respectively.

Proof. Let $j \in J_-$. By Lemma C.3.2 w.h.p , we have

$$\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j \geq \sum_{i \in I} \lambda_i v_j \phi'_{i,j} (2\Delta' + p + 1) .$$

By Lemma C.3.3, we have

$$-\mathbf{w}_j^\top \mathbf{z}_\perp \geq -\eta^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau); \quad -\mathbf{w}_j^\top \mathbf{z}_\parallel \geq -\eta^\parallel \frac{9}{10} \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_4(d, \tau, k, p)$$

Combining the last three displayed equations, we get

$$\begin{aligned} \mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}_\perp) + b_j &= \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j - \mathbf{w}_j^\top \mathbf{z}_\perp \\ &\geq \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} (2\Delta' + p + 1) - \eta^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) \\ &= \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} (2\Delta' + p + 1 - c_6(g, \tau) \eta^\perp) . \end{aligned}$$

Similarly for \mathbf{z}'' :

$$\mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}_n) + b_j \geq \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} \left(2\Delta' + p + 1 - \frac{9}{10} c_4(d, \tau, k, p) \eta'' \right).$$

Note that by Lemma C.3.4 we have $c_6(g, \tau), c_4(d, \tau, k, p) > 0$. Hence, for $\eta^\perp \geq \frac{2\Delta' + p + 1}{c_6(g, \tau)}$, $\eta'' \geq \frac{10}{9} \cdot \frac{2\Delta' + p + 1}{c_4(d, \tau, k, p)}$ we have $\mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}^\perp) + b_j, \mathbf{w}_j^\top (\tilde{\mathbf{x}} - \mathbf{z}'') + b_j \geq 0$ respectively.

The proof for $j \in J_+$ is similar. Namely, by Lemmas C.3.2 and C.3.3, we have

$$\begin{aligned} \mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}_\perp) + b_j &= \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j + \mathbf{w}_j^\top \mathbf{z}_\perp \\ &\geq \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} (2\Delta' + p + 1) + \eta^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) \\ &= \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} (-2\Delta' - p - 1 + c_6(g, \tau) \eta^\perp). \end{aligned}$$

and hence for $\eta^\perp \geq \frac{2\Delta' + p + 1}{c_6(g, \tau)}$ we get $\mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}^\perp) + b_j \geq 0$. Finally using identical arguments, for $\eta'' \geq \frac{10}{9} \cdot \frac{2\Delta' + p + 1}{c_4(d, \tau, k, p)}$ we get $\mathbf{w}_j^\top (\tilde{\mathbf{x}} + \mathbf{z}'') + b_j \geq 0$ \square

Proof of Theorem 4.2. We showcase the explicit steps for off-manifold $\mathbf{z}_\perp = (\eta_1^\perp + \eta_2^\perp) \mathbf{u}_\perp$, identical steps can be implemented to get the result for on-manifold $\mathbf{z}_n = (\eta_1'' + \eta_2'') \mathbf{u}_n$ and total $\mathbf{z} = (\eta_1 + \eta_2) \mathbf{u}$ attacks. We assume $r \in Q_+$, the case where $r \in Q_-$ can be proved similarly. We denote $\tilde{\mathbf{x}}' = \tilde{\mathbf{x}} - \mathbf{z}_\perp$. By Lemma C.3.3, for every $j \in J_+$ we have

$$\begin{aligned} \mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j &= \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j - \mathbf{w}_j^\top (\eta_1^\perp + \eta_2^\perp) \mathbf{u}_\perp \\ &\leq \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j - (\eta_1^\perp + \eta_2^\perp) \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau). \end{aligned}$$

$$\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j \leq \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j. \quad (16)$$

Thus, in the neurons J_+ the input does not increase when moving from $\tilde{\mathbf{x}}$ to $\tilde{\mathbf{x}}'$.

Consider now $j \in J_-$. Let $\tilde{\mathbf{x}}_* = \mathbf{x} - \eta_1^\perp \mathbf{u}_\perp$. Setting $\eta_1^\perp = \frac{2\Delta' + p + 1}{c_6(g, \tau)}$ and using Lemma C.3.5, we have $\mathbf{w}_j^\top \tilde{\mathbf{x}}_* + b_j \geq 0$. Also, by Lemma C.3.3, we have

$$\begin{aligned} \mathbf{w}_j^\top \tilde{\mathbf{x}}_* + b_j &= \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j - \mathbf{w}_j^\top \eta_1^\perp \mathbf{u}_\perp \\ &\geq \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j - \eta_1^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau), \end{aligned}$$

the above is at least $\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j$. Thus, when moving from $\tilde{\mathbf{x}}$ to $\tilde{\mathbf{x}}_*$ the input to the neurons J_- can only increase, and at $\tilde{\mathbf{x}}_*$ it is non-negative.

Next, we move from $\tilde{\mathbf{x}}_*$ to $\tilde{\mathbf{x}}'$. We have

$$\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j = \mathbf{w}_j^\top \tilde{\mathbf{x}}_* + b_j - \eta_2^\perp \mathbf{w}_j^\top \mathbf{u}_\perp \geq \max \{0, \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j\} - \eta_2^\perp \mathbf{w}_j^\top \mathbf{u}_\perp.$$

By Lemma C.3.3, the above is at least

$$\max \{0, \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j\} - \eta_2^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) \geq 0, \quad (17)$$

Overall, we have

$$\begin{aligned}
 \mathcal{N}_\theta(\tilde{\mathbf{x}}') &= \left[\sum_{j \in J_+} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j) \right] + \left[\sum_{j \in J_-} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j) \right] \\
 &\stackrel{(i)}{=} \left[\sum_{j \in J_+} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j) \right] + \left[\sum_{j \in J_-} v_j (\mathbf{w}_j^\top \tilde{\mathbf{x}}' + b_j) \right] \\
 &\stackrel{(ii)}{\leq} \left[\sum_{j \in J_+} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) \right] + \\
 &\quad \left[\sum_{j \in J_-} v_j \left(\max \{0, \mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j\} - \eta_2^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) \right) \right],
 \end{aligned}$$

where in (i) we used Eq. (17), and in (ii) we used both Eq. (16) and Eq. (17). Now, the above equals

$$\begin{aligned}
 &\left[\sum_{j \in [w]} v_j \phi(\mathbf{w}_j^\top \tilde{\mathbf{x}} + b_j) \right] - \left[\sum_{j \in J_-} v_j \eta_2^\perp \sum_{i \in [n]} \lambda_i v_j \phi'_{i,j} c_6(g, \tau) \right] \\
 &= \mathcal{N}_\theta(\tilde{\mathbf{x}}) - \eta_2^\perp c_6(g, \tau) \left[\sum_{q' \in [k]} \sum_{i \in I(q')} \sum_{j \in J_-} \lambda_i v_j^2 \phi'_{i,j} \right].
 \end{aligned}$$

Combining the above with Lemma C.3.1, Lemma C.2.3, and Lemma C.3.3, we get

$$\begin{aligned}
 \mathcal{N}_\theta(\tilde{\mathbf{x}}') &\leq 2 - \eta_2^\perp c_6(g, \tau) \left[\sum_{q' \in Q_-} \sum_{i \in I(q')} \sum_{j \in J_-} \lambda_i v_j^2 \phi'_{i,j} \right] \\
 &\leq 2 - \eta_2^\perp c_6(g, \tau) |Q_-| \left(1 - \frac{c'}{1 - 2c'} \right) \frac{1}{3d + \Delta' + 1} \\
 &\leq 2 - \eta_2^\perp c_6(g, \tau) ck \left(\frac{1 - 3c'}{1 - 2c'} \right) \frac{1}{3d + \Delta' + 1}.
 \end{aligned}$$

For

$$\eta_2^\perp = \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_6(g, \tau) \cdot (1 - 3c')ck}$$

we conclude that $\mathcal{N}_\theta(\mathbf{x}')$ is at most -1 . When $r \in Q_-$ one can get an identical result with a perturbation $+\mathbf{z}_\perp$ instead of $-\mathbf{z}_\perp$.

Hence, an off-manifold perturbation $\mathbf{z}_\perp = (\eta_1^\perp + \eta_2^\perp)\mathbf{u}_\perp$ where ,

$$\eta_1^\perp = \frac{2\Delta' + p + 1}{c_6(g, \tau)}, \eta_2^\perp = \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_6(g, \tau) \cdot (1 - 3c')ck} \quad (18)$$

can flip the sign of the neuron accordingly.

If we follow the identical steps as above for on manifold perturbation $\mathbf{z}_\parallel = (\eta_1^\parallel + \eta_2^\parallel)\mathbf{u}_\parallel$, we get that:

$$\eta_1^\parallel = \frac{10}{9} \cdot \frac{2\Delta' + p + 1}{c_4(d, \tau, k, p)}, \eta_2^\parallel = \frac{10}{9} \cdot \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_4(d, \tau, k, p) \cdot (1 - 3c')ck} \quad (19)$$

Off manifold perturbation upper bounds

$$\begin{aligned}
 (\eta_1^\perp + \eta_2^\perp) &= \left(\frac{2\Delta' + p + 1}{c_6(g, \tau)} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_6(g, \tau) \cdot (1 - 3c')ck} \right) \\
 &\leq \left(\frac{2(p + \Delta' + 1)}{c_6(g, \tau)} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_6(g, \tau) \cdot (1 - 3c')ck} \right) \\
 &\stackrel{(i)}{=} \left(\frac{2c'(d - \Delta' + 1)/k}{\frac{18}{42} \cdot g\tau^2} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{\frac{18}{42} \cdot g\tau^2 \cdot (1 - 3c')ck} \right) \\
 &\stackrel{(ii)}{\leq} \left(\frac{1}{k} \cdot \frac{2}{10}(d + 1) + \frac{1}{ck} \cdot \frac{3(3d + \frac{d}{21} + 1)}{\frac{18}{42} \cdot g\tau^2(1 - \frac{3}{10})} \right) \\
 &\leq \left(\frac{1}{k} \cdot \mathcal{O}\left(\frac{d}{g\tau^2}\right) + \frac{1}{ck} \cdot \mathcal{O}\left(\frac{d}{g\tau^2}\right) \right) \\
 &\leq \mathcal{O}\left(\frac{d}{ckg\tau^2}\right), \tag{20}
 \end{aligned}$$

where in (i) we used $k = c' \cdot \frac{d - \Delta' + 1}{p + \Delta' + 1}$ and $\frac{18}{21} \cdot \frac{g\tau^2}{2} \leq c_6(g, \tau)$ from Lemma C.3.4. In (ii) we used Lemma C.2.1 and $c' < 1/10$.

$$\begin{aligned}
 \|\mathbf{u}_\perp\|^2 &= \left\| \sum_{q \in [k]} y^{(q)} \mathbf{P}\zeta^{(q)} \right\|^2 = \sum_{r \in [k]} \sum_{q \in [k]} y^{(r)} y^{(q)} \langle \zeta^{(r)}, \zeta^{(q)} \rangle_{\mathbf{P}} \\
 &= \sum_{r \in [k]} \left[\langle \zeta^{(r)}, \zeta^{(r)} \rangle_{\mathbf{P}} + \sum_{q \neq r} y^{(r)} y^{(q)} \langle \zeta^{(r)}, \zeta^{(q)} \rangle_{\mathbf{P}} \right] \\
 &\stackrel{(iii)}{\leq} k \cdot \frac{13g\tau^2}{8} + k^2 c_3(g, \tau, k) \\
 &\leq k \cdot \frac{13g\tau^2}{8} + k \frac{g\tau^2}{20} = \mathcal{O}(kg\tau^2).
 \end{aligned}$$

$$\|\mathbf{u}_\perp\| = \mathcal{O}\left(\sqrt{kg\tau^2}\right) \tag{21}$$

where in (iii) we used $\|\mathbf{P}\zeta^{(r)}\| < \frac{13g\tau^2}{8}$; $\langle \zeta^{(r)}, \zeta^{(q)} \rangle_{\mathbf{P}} < c_3(g, \tau, k)$ Lemma C.1.5 .

Combining Eq. (20) and Eq. (21) we have $\|\mathbf{z}_\perp\| = \mathcal{O}\left(\frac{d}{c\sqrt{kg\tau^2}}\right)$.

$$\begin{aligned}
 \|\mathbf{u}_\perp\|_\infty &= \left\| \sum_{q \in [k]} y^{(q)} \mathbf{P}\zeta^{(q)} \right\|_\infty \leq \sum_{q \in [k]} \|\mathbf{P}\zeta^{(q)}\|_\infty \\
 &\stackrel{(iv)}{\leq} 3k \cdot \tau \sqrt{2 \log(2g)} = \mathcal{O}\left(k\tau \sqrt{2 \log(2g)}\right). \tag{22}
 \end{aligned}$$

For (iv) use, $\|\mathbf{P}\zeta^{(q)}\|_\infty \leq 3\tau \sqrt{2 \log(2g)}$ (Lemma C.1.5).

Combining Eq. (20) and Eq. (22) we have $\|\mathbf{z}_\perp\|_\infty = \mathcal{O}\left(\frac{d\sqrt{2 \log(2g)}}{cg\tau}\right)$.

On manifold perturbation upper bounds

$$\begin{aligned}
 (\eta_1^{\text{II}} + \eta_2^{\text{II}}) &= \frac{10}{9} \left(\frac{2\Delta' + p + 1}{c_4(d, \tau, k, p)} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_4(d, \tau, k, p) \cdot (1 - 3c')ck} \right) \\
 &\leq \frac{10}{9} \left(\frac{2(p + \Delta' + 1)}{c_4(d, \tau, k, p)} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{c_4(d, \tau, k, p) \cdot (1 - 3c')ck} \right) \\
 &\stackrel{(i)}{=} \left(\frac{2c'(d - \Delta' + 1)/k}{\frac{18}{21} \cdot d(1 + \frac{\tau^2}{2})} + \frac{3(3d + \Delta' + 1)(1 - 2c')}{\frac{18}{21} \cdot d(1 + \frac{\tau^2}{2}) \cdot (1 - 3c')ck} \right) \\
 &\stackrel{(ii)}{\leq} \left(\frac{1}{k} \cdot \frac{\frac{2}{10}(d + 1)}{\frac{18}{21} \cdot d(1 + \frac{\tau^2}{2})} + \frac{1}{ck} \cdot \frac{3(3d + \frac{d}{21} + 1)}{\frac{18}{21} \cdot d(1 + \frac{\tau^2}{2})(1 - \frac{3}{10})} \right) \\
 &\leq \left(\frac{1}{k} \cdot \mathcal{O}\left(\frac{1}{1 + \frac{\tau^2}{2}}\right) + \frac{1}{ck} \cdot \mathcal{O}\left(\frac{1}{1 + \frac{\tau^2}{2}}\right) \right) \\
 &\leq \mathcal{O}\left(\frac{1}{ck(2 + \tau^2)}\right), \tag{23}
 \end{aligned}$$

where in (i) we used $k = c' \cdot \frac{d - \Delta' + 1}{p + \Delta' + 1}$ and $\frac{18}{21} \cdot d(1 + \frac{\tau^2}{2}) \leq c_4(d, \tau, k, p)$ from Lemma C.3.4. In (ii) we used Lemma C.2.1 and $c' < 1/10$.

$$\begin{aligned}
 \|\mathbf{u}_{\text{II}}\|^2 &= \left\| \sum_{q \in [k]} y^{(q)}(\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)}) \right\|^2 \\
 &= \sum_{r \in [k]} \sum_{q \in [k]} y^{(r)} y^{(q)} \langle \boldsymbol{\zeta}^{(r)}, \boldsymbol{\zeta}^{(q)} \rangle_{\mathbf{P}_o} + \sum_{r \in [k]} \sum_{q \in [k]} y^{(r)} y^{(q)} \langle \tilde{\boldsymbol{\mu}}^{(r)}, \tilde{\boldsymbol{\mu}}^{(q)} \rangle + \sum_{r \in [k]} \sum_{q \in [k]} 2y^{(r)} y^{(q)} \langle \tilde{\boldsymbol{\mu}}^{(r)}, \boldsymbol{\zeta}^{(q)} \rangle_{\mathbf{P}_o} \\
 &= \sum_{r \in [k]} \left[\langle \boldsymbol{\zeta}^{(r)}, \boldsymbol{\zeta}^{(r)} \rangle_{\mathbf{P}_o} + d + \sum_{q \neq r} y^{(r)} y^{(q)} \left(\langle \boldsymbol{\zeta}^{(r)}, \boldsymbol{\zeta}^{(q)} \rangle_{\mathbf{P}_o} + \langle \tilde{\boldsymbol{\mu}}^{(r)}, \tilde{\boldsymbol{\mu}}^{(q)} \rangle \right) \right] + \sum_{r \in [k]} V_r \\
 &\stackrel{(iii)}{\leq} k \cdot \frac{13d\tau^2}{8} + kd + k^2 c_5(d, \tau, k) + k^2 p + \frac{k}{10} \cdot c_4(d, \tau, k, p) \\
 &\stackrel{(iv)}{\leq} k \cdot \frac{13d\tau^2}{8} + kd + k \frac{d\tau^2}{20} + k^2 p + \frac{k}{10} \cdot d(1 + \frac{\tau^2}{2}) \\
 &\stackrel{(v)}{\leq} k \cdot \frac{13d\tau^2}{8} + kd + k \frac{d\tau^2}{20} + \frac{kd + k}{10} + \frac{k}{10} \cdot d(1 + \frac{\tau^2}{2}) = kd \left(\frac{6}{5} + \frac{69\tau^2}{40} \right) + \frac{k}{10} \\
 &\leq \mathcal{O}(kd(2 + \tau^2)) \tag{24}
 \end{aligned}$$

For (iii) note that $\|\mathbf{P}_o \boldsymbol{\zeta}^{(r)}\| \leq \frac{13d\tau^2}{8}$, $\langle \boldsymbol{\zeta}^{(r)}, \boldsymbol{\zeta}^{(q)} \rangle_{\mathbf{P}_o} < c_5(d, \tau, k)$, $|V_r| < \frac{1}{10} \cdot c_4(d, \tau, k, p)$ (Lemma C.1.7).

For (iv), we know that $c_5(d, \tau, k) < \frac{d\tau^2}{20k}$ and $c_4(d, \tau, k, p) < d(1 + \tau^2/2)$ (Lemma C.3.4).

For (v), $k(p + \Delta' + 1) < 1/10 \cdot (d - \Delta' + 1) \implies kp < 1/10 \cdot (d + 1)$.

Combining Eq. (23) and Eq. (24) we have $\|\mathbf{z}_{\text{II}}\| = \mathcal{O}\left(\sqrt{\frac{d}{c^2 k(2 + \tau^2)}}\right)$

$$\begin{aligned}
 \|\mathbf{u}_{\text{II}}\|_{\infty} &= \left\| \sum_{q \in [k]} y^{(q)}(\tilde{\boldsymbol{\mu}}^{(q)} + \mathbf{P}_o \boldsymbol{\zeta}^{(q)}) \right\|_{\infty} \leq \sum_{q \in [k]} \underbrace{\|\tilde{\boldsymbol{\mu}}^{(q)}\|_{\infty}}_{\leq \|\tilde{\boldsymbol{\mu}}^{(q)}\|} + \sum_{q \in [k]} \|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|_{\infty} \\
 &\stackrel{(iv)}{\leq} k\sqrt{d} + 3k \cdot \tau \sqrt{2 \log(2d)} = \mathcal{O}\left(kd + k\tau \sqrt{2 \log(2d)}\right). \tag{25}
 \end{aligned}$$

For (iv) use, $\|\mathbf{P}_o \boldsymbol{\zeta}^{(q)}\|_{\infty} \leq 3\tau \sqrt{2 \log(2d)}$ (Lemma C.1.7).

Combining Eq. (23) and Eq. (25) we have $\|\mathbf{z}_{\text{II}}\|_{\infty} = \mathcal{O}\left(\frac{d + \tau \sqrt{2 \log(2d)}}{c(2 + \tau^2)}\right)$. The corresponding probability with all of the bounds holding true for the off-manifold and on-manifold perturbation is $1 - \delta_1, 1 - \delta_2$, respectively. Where δ_1, δ_2 are defined in Lemma C.1.5, Lemma C.1.7 respectively. \square

CIFAR Class	0	1	2	3	4	5	6	7	8	9
LPCA	8	11	8	11	9	13	7	14	10	16
MLE(k=5)	18.34	21.20	21.35	21.08	21.43	22.11	22.96	22.57	19.80	24.42
TwoNN	21.77	25.98	25.53	26.09	24.60	27.32	26.19	25.95	24.77	29.11

Table 4: Intrinsic dimension of CIFAR-10 Dataset.

MNIST class	0	1	2	3	4	5	6	7	8	9
LPCA	19	8	31	29	27	23	21	21	32	22
MLE (k=5)	13.01	9.65	14.14	15.14	13.65	14.70	12.65	12.08	15.65	12.79
TwoNN	15.33	12.98	15.29	16.35	14.71	15.90	14.02	13.39	16.60	14.03

Table 5: Estimated intrinsic dimension of MNIST Dataset.

D Additional Experiment Results

D.1 Simulation

D.1.1 Training setting

We use a two-layer ReLU network to train the model. The width w of the neural network is set to be 2000. The initialization of the two layers follows *Kaiming Normal distribution* (He et al., 2015). We use gradient descent to train the neural network, i.e., the batch size is $n = 1000$, and take the learning rate as 0.1 to train 1000 epochs. For all the settings, the clean training accuracy is 100. We repeat the training process 10 times for an average result and the error bar.

D.1.2 Minimal strength attack generation

In the experiments, we track the attack strength of an unconstrained attack, *on-manifold* and *off-manifold* attack. The way we generate these are described as follows:

- Usual \mathbb{R}^D attack: PGD-20 (Madry et al., 2017) attack in the whole space.
- On-manifold attack: PGD-20 attack on the intrinsic data space. We achieve this by projecting the PGD algorithm’s gradients to the underlying d dimensional intrinsic space by applying the matrix $\mathbf{P}_o = \mathbf{M}(\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T$.
- Off-manifold attack: PGD-20 attack on the off-manifold space. We achieve this by projecting the PGD algorithm’s gradients to the co-kernel g dimensional space by applying the matrix $\mathbf{P} = \mathbf{I} - \mathbf{P}_o$.

D.1.3 Additional Figures

Effect of Ambient-Intrinsic Dimension Gap on Adversarial Vulnerability

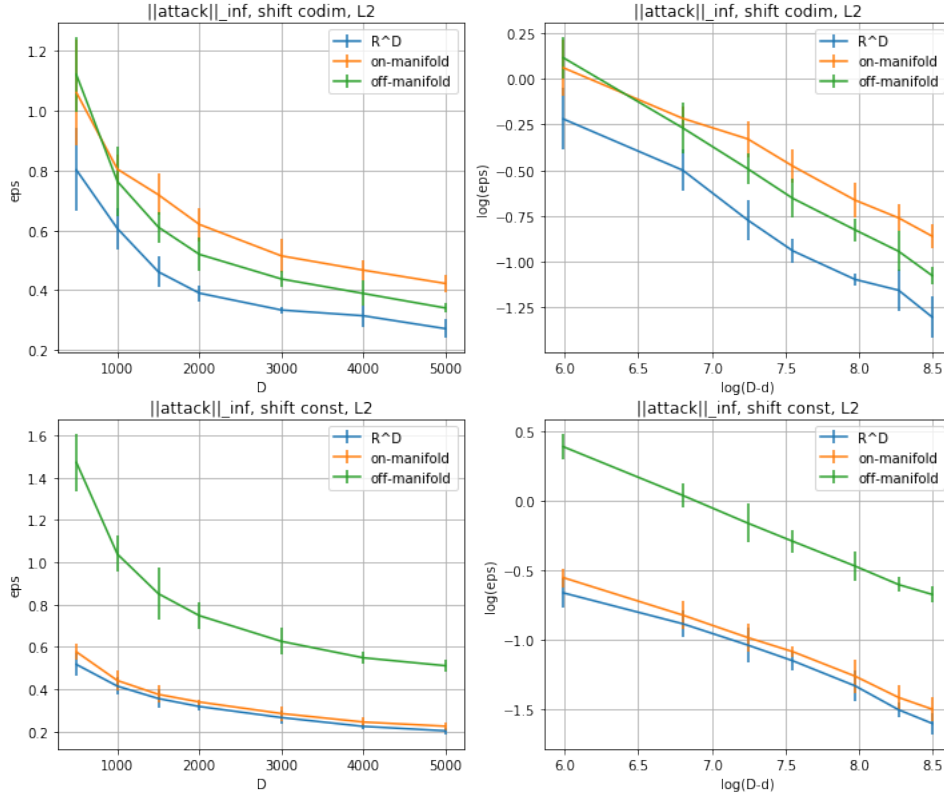


Figure 9: The ℓ_∞ norm of the ℓ_2 attack in Figure 3. The top and bottom rows correspond to the cases $\|\zeta\| = \Theta(\sqrt{d}), \Theta(1)$ respectively.

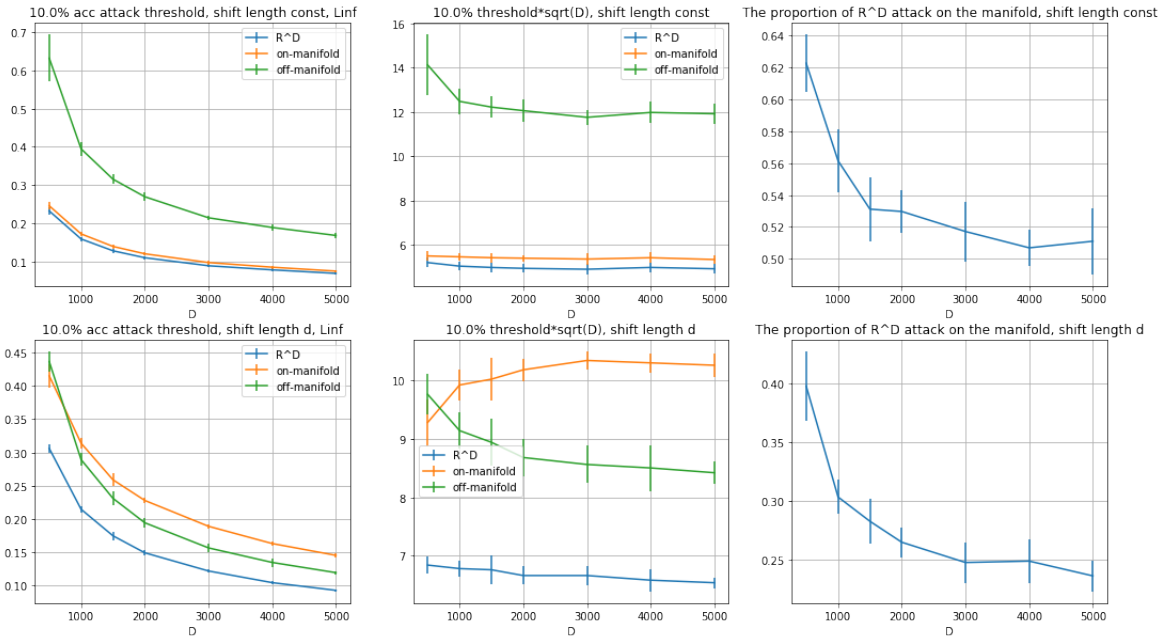


Figure 10: Attack strength threshold associating with 10% robust test accuracy (left), its scaled value (middle), and the on-manifold proportion in \mathbb{R}^D attack. The top and bottom rows correspond to the cases $\|\zeta\| = \Theta(1), \Theta(\sqrt{d})$ respectively.

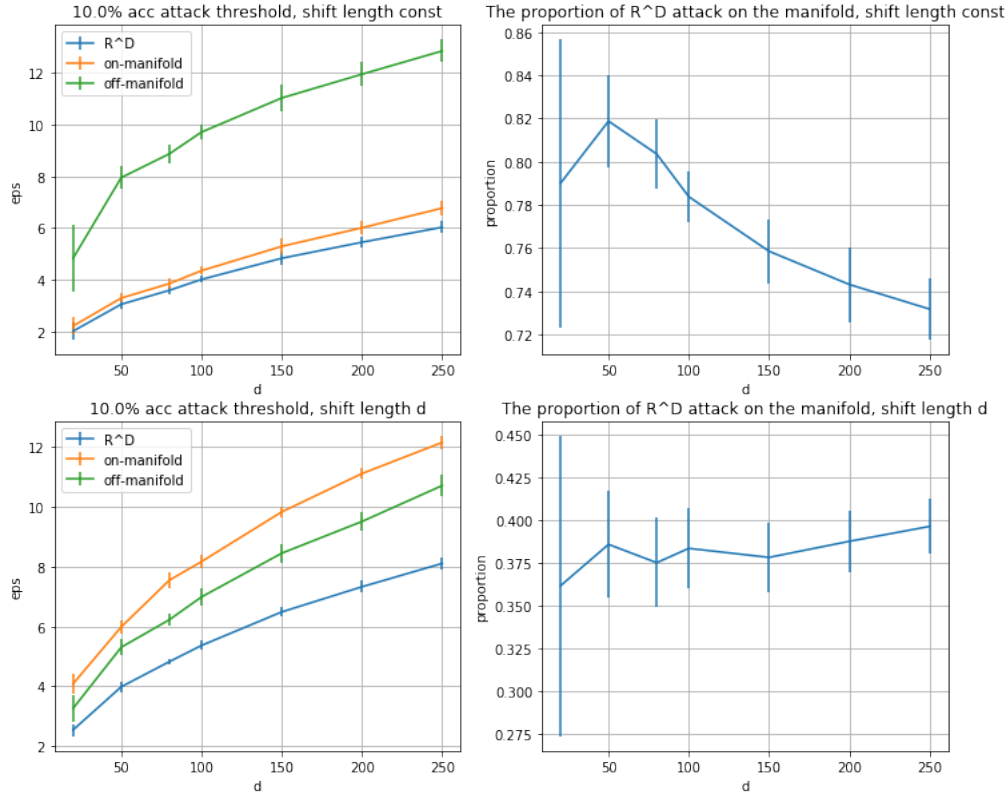


Figure 11: Attack strength threshold associating with 10% robust test accuracy (left) and the on-manifold proportion in \mathbb{R}^D attack when changing d . The top and bottom rows correspond to the cases $\|\zeta\| = \Theta(1), \Theta(\sqrt{d})$ respectively.




Figure 12: Attack strength threshold associating with 10% robust test accuracy (left) and the on-manifold proportion in \mathbb{R}^D attack when changing d . The initialization in this setting is 0.01 of the initialization used in all other figures.




Figure 13: Attack strength threshold associating with 10% robust test accuracy (left) and the on-manifold proportion in \mathbb{R}^D attack when changing d . The optimizer takes `weight_decay` as 0.1. The number of epochs is changed to 500 in this experiment to save the computation time for large D cases. Since we change the number of epochs in this experiment, we also conduct experiment for the scenario of Figure 3, and there is no change in the result.

D.2 MNIST

D.2.1 Architecture and training setting

We use Adam optimizer with learning rate 10^{-3} to train the model. The batch size is 256.

Layer
Conv2d(1, 16, 4, stride=2, padding=1), ReLU
Conv2d(16, 32, 4, stride=2, padding=1), ReLU
Linear($32 \cdot (7 + P/2) \cdot (7 + P/2)$, 100), ReLU
Linear(100, 10)

Table 6: Neural network architecture.

We take 0.02 as the stopping criteria because when $P = 0$, the training loss is around 0.02 after 10 epochs. We repeat the experiment 30 times to get a mean and standard deviation.

D.2.2 Additional Figures




Figure 14: The relationship between attack strength and the padding number P in Table 6 (left) and the relationship between robust test accuracy and attack strength under different P . ℓ_2 attack. Train loss is 0.02.




Figure 15: The relationship between attack strength and the padding number P in Table 6 (left) and the relationship between robust test accuracy and attack strength under different P . ℓ_∞ attack. Train loss is 0.02.