

Generalization Bounds for Adversarial Contrastive Learning

Xin Zou

*School of Computer Science
Wuhan University
Wuhan, Hubei, China*

ZOUXIN2021@GMAIL.COM

Weiwei Liu*

*School of Computer Science
Wuhan University
Wuhan, Hubei, China*

LIUWEIWEI863@GMAIL.COM

Editor: Joan Bruna

Abstract

Deep networks are well-known to be fragile to adversarial attacks, and adversarial training is one of the most popular methods used to train a robust model. To take advantage of unlabeled data, recent works have applied adversarial training to contrastive learning (Adversarial Contrastive Learning; ACL for short) and obtain promising robust performance. However, the theory of ACL is not well understood. To fill this gap, we leverage the Rademacher complexity to analyze the generalization performance of ACL, with a particular focus on linear models and multi-layer neural networks under ℓ_p attack ($p \geq 1$). Our theory shows that the average adversarial risk of the downstream tasks can be upper bounded by the adversarial unsupervised risk of the upstream task. The experimental results validate our theory.

Keywords: Robustness, Adversarial learning, Contrastive learning, Rademacher complexity, Generalization bound.

1. Introduction

Deep neural networks (DNNs) have achieved state-of-the-art performance in many fields. However, several prior works (Szegedy et al., 2014; Goodfellow et al., 2015) have shown that DNNs may be vulnerable to imperceptibly changed adversarial examples, which causes a lot of focus on the robustness of the models (Madry et al., 2018; Mao et al., 2020; Li et al., 2022).

One of the most popular approaches to achieving adversarial robustness is adversarial training, which involves training the model with samples perturbed to maximize the loss on the target model (Goodfellow et al., 2015; Madry et al., 2018; Zhang et al., 2019). Schmidt et al. (2018) show that adversarial robust generalization requires a larger amount of data, while Yin et al. (2019); Awasthi et al. (2020) show that the Rademacher Complexity of adversarial training is strictly larger in theory than that of natural training, which implies that we need more data for adversarial training.

*. Corresponding Author.

Since the labeled data is limited and expensive to obtain, one option would be to use large-scale unlabeled data and apply self-supervised learning (Gidaris et al., 2018; Noroozi and Favaro, 2016), an approach that trains the model on unlabeled data in a supervised manner by utilizing self-generated labels from the data itself. Contrastive Learning (CL) (Chen et al., 2020; He et al., 2020), which aims to maximize the feature similarity of similar pairs and minimize the feature similarity of dissimilar ones, is a popular self-supervised learning technique.

Recently, Kim et al. (2020); Ho and Vasconcelos (2020); Jiang et al. (2020) apply adversarial training in CL and achieve state-of-the-art model robustness. They find that if adversarial training is conducted on the upstream contrastive task, the trained model will be robust on the downstream supervised task. However, their results are all empirical and lack theoretical analysis. To fill this gap, we here present a theoretical analysis of adversarial contrastive learning through the lens of Rademacher complexity, with a particular focus on linear models and multi-layer neural networks. Our theoretical results show that the average adversarial risk of the downstream tasks can be upper bounded by the adversarial unsupervised risk of the upstream task; this implies that if we train a robust feature extractor on the upstream task, we can obtain a model that is robust on downstream tasks.

The remainder of this article is structured as follows: §2 introduces some related works in this field. §3 provides some basic definitions and settings that will be used in the following sections. §4 presents our first part’s main results, which show a connection between the robust risk for the upstream task and the robust risk for the downstream tasks. §5 shows our second part’s results, it outlines our bounds for linear models and multi-layer neural networks by bounding the Rademacher complexity of each hypothesis class. §6 shows some experimental results that verify our theory; finally, the conclusions are presented in the last section.

2. Related Work

Adversarial Robustness. Szegedy et al. (2014) show that DNNs are fragile to imperceptible distortions in the input space. Subsequently, Goodfellow et al. (2015) propose the fast gradient sign method (FGSM), which perturbs a target sample towards its gradient direction to increase the loss and then uses the generated sample to train the model in order to improve the robustness. Following this line of research, Madry et al. (2018); Moosavi-Dezfooli et al. (2016); Kurakin et al. (2017); Carlini and Wagner (2017) propose iterative variants of the gradient attack with improved adversarial learning frameworks. Besides, Ma et al. (2022) analyze the trade-off between robustness and fairness, Li and Liu (2023) study the worst-class adversarial robustness in adversarial training. For the theoretical perspective, Montasser et al. (2019) study the PAC learnability of adversarial robust learning, Xu and Liu (2022) extend the work of Montasser et al. (2019) to multiclass case and Yin et al. (2019); Awasthi et al. (2020) give theoretical analyses to adversarial training by standard uniform convergence argumentation and giving a bound of the Rademacher complexity. Our work is quite different from Yin et al. (2019); Awasthi et al. (2020), firstly, they just analyze the linear models and two-layer neural networks, but we consider linear models and multi-layer deep neural networks; secondly, they consider the classification loss, which is much easier to analyze than our contrastive loss.

Contrastive Learning. Contrastive Learning is a popular self-supervised learning paradigm that attempts to learn a good feature representation by minimizing the feature distance of similar pairs and maximizing the feature distance of dissimilar pairs (Chen et al., 2020; Wang and Liu, 2022). SimCLR (Chen et al., 2020) learns representations by maximizing the agreement between differently augmented views of the same data, while MoCo (He et al., 2020) builds large and consistent dictionaries for unsupervised learning with a contrastive loss. From the theoretical perspective, Saunshi et al. (2019) first presents a framework to analyze CL. We generalize their framework to the adversarial CL. The key challenging issues in this work are how to define and rigorously analyze the adversarial CL losses. Moreover, we further analyze the Rademacher complexity of linear models and multi-layer deep neural networks composited with a complex adversarial contrastive loss. Neyshabur et al. (2015) show an upper bound of the Rademacher complexity for neural networks under group norm regularization. In fact, the Frobenius norm and the $\ell_{1,\infty}$ -norm in our theoretical analysis are also group norms, while the settings and proof techniques between Neyshabur et al. (2015) and our work are quite different: (1) they consider the standard setting while we consider a more difficult adversarial setting; (2) they consider the Rademacher complexity of neural networks with size 1 output layer, while we consider the case that the neural network is composited with a complex contrastive learning loss; (3) they prove their results by reduction with respect to the number of the layers, while motivated by the technique of Gao and Wang (2021), we use the covering number of the neural network to upper bound the Rademacher complexity in the adversarial case.

Adversarial Contrastive Learning. Several recent works (Kim et al., 2020; Ho and Vasconcelos, 2020; Jiang et al., 2020) apply adversarial training in the contrastive pre-training stage to improve the robustness of the models on the downstream tasks, achieving good robust performance in their experiments. Our work attempts to provide a theoretical explanation as to why models robustly trained on the upstream task can be robust on downstream tasks.

3. Problem Setup

We introduce the problem setups in this section.

3.1 Basic Contrastive Learning Settings

We first set up some notations and describe the contrastive learning framework.

Let $\mathcal{X} \in \mathbb{R}^m$ be the domain of all possible data points, this paper assumes that \mathcal{X} is bounded. Contrastive learning assumes that we get the similar data in the form of pairs (x, x^+) , which is drawn from a distribution \mathcal{D}_{sim} on \mathcal{X}^2 , and k independent and identically distributed (i.i.d.) negative samples $x_1^-, x_2^-, \dots, x_k^-$ drawn from a distribution \mathcal{D}_{neg} on \mathcal{X} . Given the training set $\mathcal{S} = \{(x_i, x_i^+, x_{i1}^-, \dots, x_{ik}^-)\}_{i=1}^M$, we aim to learn a representation f from \mathcal{F} that maps similar pairs (x, x^+) into similar points $(f(x), f(x^+))$, while at the same time keeping $f(x_i^-), \dots, f(x_k^-)$ away from $f(x)$, where \mathcal{F} is a class of representation functions $f : \mathcal{X} \rightarrow \mathbb{R}^n$.

Latent Classes. Let \mathcal{C} denote the set of all latent classes (Saunshi et al., 2019) that are all possible classes for points in \mathcal{X} ; for each class $c \in \mathcal{C}$, moreover, the probability \mathcal{D}_c

over \mathcal{X} captures the probability that a point belongs to class c . The distribution on \mathcal{C} is denoted by ρ .

To formalize the similarity among data in \mathcal{X} , assume we obtain i.i.d. similar data points x, x^+ from the same distribution \mathcal{D}_c , where c is randomly selected according to the distribution ρ on latent classes \mathcal{C} . We can then define \mathcal{D}_{sim} and \mathcal{D}_{neg} as follows:

Definition 1 (\mathcal{D}_{sim} and \mathcal{D}_{neg} , Saunshi et al., 2019) *For unsupervised tasks, we define the distribution of sampling similar samples $\mathcal{D}_{sim}(x, x^+)$ and negative sample $\mathcal{D}_{neg}(x^-)$ as follows:*

$$\mathcal{D}_{sim}(x, x^+) = \mathbb{E}_{c \sim \rho} \mathcal{D}_c(x) \mathcal{D}_c(x^+), \quad \mathcal{D}_{neg}(x^-) = \mathbb{E}_{c \sim \rho} \mathcal{D}_c(x^-).$$

Supervised Tasks. For supervised tasks, we focus on the tasks in which a representation function f will be tested on a $(k+1)$ -way supervised task \mathcal{T} consisting of distinct classes $\{c_1, c_2, \dots, c_{k+1}\} \subseteq \mathcal{C}$, while the labeled data set of task \mathcal{T} consists of M i.i.d. examples drawn according to the following process: A label $c \in \{c_1, \dots, c_{k+1}\}$ is selected according to a distribution $\mathcal{D}_{\mathcal{T}}$, and a sample x is drawn from \mathcal{D}_c . The distribution of the labeled pair (x, c) is defined as: $\mathcal{D}_{\mathcal{T}}(x, c) = \mathcal{D}_{\mathcal{T}}(c) \mathcal{D}_c(x)$.

3.2 Evaluation Metric for Representations

We evaluate the quality of a representation function f with reference to its performance on a multi-class classification task \mathcal{T} using a linear classifier.

Consider a task $\mathcal{T} = \{c_1, \dots, c_{k+1}\}$. A multi-class classifier for \mathcal{T} is a function $g: \mathcal{X} \rightarrow \mathbb{R}^{k+1}$, the output coordinates of which are indexed by the classes c in task \mathcal{T} .

Let $\{g(x)_y - g(x)_{y'}\}_{y' \neq y}$ be a k -dimensional vector of differences in the coordinates of the output of the classifier g . The loss function of g on a point $(x, y) \in \mathcal{X} \times \mathcal{T}$ is defined as $\ell(\{g(x)_y - g(x)_{y'}\}_{y' \neq y})$. For example, one often considers the standard hinge loss $\ell(\mathbf{v}) = \max\{0, 1 + \max_i \{-v_i\}\}$ and the logistic loss $\ell(\mathbf{v}) = \log_2(1 + \sum_i \exp(-v_i))$ for $\mathbf{v} \in \mathbb{R}^k$.

The supervised risk of classifier g is defined as follows:

$$L_{sup}(\mathcal{T}, g) := \mathbb{E}_{(x, c) \sim \mathcal{D}_{\mathcal{T}}} \left[\ell \left(\{g(x)_c - g(x)_{c'}\}_{c' \neq c} \right) \right].$$

The risk $L_{sup}(\mathcal{T}, g)$ of classifier g on task \mathcal{T} measures the quality of the outputs of g , take the hinge loss as an example, our goal is to get a classifier that has much higher confidence for the true class (i.e., the value $g(x)_c$) than others (i.e., the values $g(x)_{c'}$ for $c' \neq c$), so we want all the differences $g(x)_c - g(x)_{c'}$ for $c \neq c'$ to be as large as possible. If g wrongly classifies x , i.e., $\arg \max_{c'} g(x)_{c'} \neq c$, of course, ℓ on (x, c) is not smaller than 1; however, even though g correctly classifies x , the loss value can not decrease to zero unless $g(x)_c - g(x)_{c'} \geq 1$ for all $c' \neq c$.

Let $\{c_i\}_{i=1}^{k+1} = \{c_1, \dots, c_{k+1}\}$ be a set of classes from \mathcal{C} . Given the matrix $W \in \mathbb{R}^{(k+1) \times n}$, we have $g(x) = Wf(x)$ as a classifier that composite the feature extractor g and linear classifier W . The supervised risk of f is defined as the risk of g when the best W is chosen:

$$L_{sup}(\mathcal{T}, f) := \inf_{W \in \mathbb{R}^{(k+1) \times n}} L_{sup}(\mathcal{T}, Wf).$$

When training a feature extractor f on the upstream task, we do not make predictions on the examples, so we can't define the risk of the feature extractor f as we do for the classifier g . Our goal on the upstream task is to train a feature extractor that can perform well on the downstream tasks, so we consider the potential of the feature extractor here, i.e., the minimal possible classification risk of a linear classifier on the feature extracted by f . So we define the risk of the feature extractor f as above by taking the infimum over all linear classifiers.

Definition 2 (Mean Classifier, Saunshi et al., 2019) For a function f and task $\mathcal{T} = \{c_1, \dots, c_{k+1}\}$, the mean classifier is W^μ , whose c^{th} row is the mean $\mu_c := \mathbb{E}_{x \sim \mathcal{D}_c} [f(x)]$ and we define $L_{sup}^\mu(\mathcal{T}, f) := L_{sup}(\mathcal{T}, W^\mu f)$.

Since $L_{sup}(\mathcal{T}, f)$ involves taking infimum over all possible linear classifiers, it is difficult to analyze $L_{sup}(\mathcal{T}, f)$ and establish connections between the risk of the unsupervised upstream task and $L_{sup}(\mathcal{T}, f)$. We introduce the mean classifier to bridge them by bounding average $L_{sup}^\mu(\mathcal{T}, f)$ over the tasks with the risk of the unsupervised upstream risk (for more details, please refer to §4).

Definition 3 (Average Supervised Risk) The average supervised risk for a function f on $(k+1)$ -way tasks is defined as:

$$L_{sup}(f) := \mathbb{E}_{\{c_i\}_{i=1}^{k+1} \sim \rho^{k+1}} [L_{sup}(\{c_i\}_{i=1}^{k+1}, f) | c_i \neq c_j]$$

Given the mean classifier, the average supervised risk for a function f is defined as follows:

$$L_{sup}^\mu(f) := \mathbb{E}_{\{c_i\}_{i=1}^{k+1} \sim \rho^{k+1}} [L_{sup}^\mu(\{c_i\}_{i=1}^{k+1}, f) | c_i \neq c_j].$$

In contrastive learning, the feature extractor is trained as a pretrained model to be used on downstream tasks. There may be many different downstream tasks such as binary classification tasks with classes $\{c_1, c_2\} \subseteq \mathcal{C}$ and many other multi-class classification tasks, so here we consider the average error of $L_{sup}^\mu(\mathcal{T}, f)$ over all possible tasks \mathcal{T} sampled from ρ^{k+1} as the final performance measure of a feature extractor f .

3.3 Contrastive Learning Algorithm

We denote k as the number of negative samples used for training and $(x, x^+) \sim \mathcal{D}_{sim}$, $(x_1^-, \dots, x_k^-) \sim \mathcal{D}_{neg}^k$.

Definition 4 (Unsupervised Risk) The unsupervised risk is defined as follows:

$$L_{un}(f) := \mathbb{E}[\ell(\{f(x)^T(f(x^+) - f(x_i^-))\}_{i=1}^k)].$$

Given M samples $\{(x_j, x_j^+, x_{j1}^-, \dots, x_{jk}^-)\}_{j=1}^M$ from $\mathcal{D}_{sim} \times \mathcal{D}_{neg}^k$, the empirical counterpart of unsupervised risk is defined as follows:

$$\hat{L}_{un}(f) := \frac{1}{M} \sum_{j=1}^M \ell(\{f(x_j)^T(f(x_j^+) - f(x_{ji}^-))\}_{i=1}^k).$$

In the contrastive learning upstream task, we want to learn a feature extractor f such that f maps examples from the same class to similar features and makes the features of examples from different classes far away from each other. So for the examples $(x, x^+, x_1^-, \dots, x_k^-)$, we want $f(x)^T f(x^+)$ to be as large as possible while $f(x)^T f(x_i^-), i = 1, \dots, k$ to be as small as possible. The loss ℓ we defined before elegantly captures the aim of contrastive learning, if we set $v_i = f(x)^T f(x^+) - f(x)^T f(x_i^-)$, minimizing $\ell(\{f(x)^T(f(x^+) - f(x_i^-))\}_{i=1}^k)$ will yield large $f(x)^T f(x^+)$ and small $f(x)^T f(x_i^-), i = 1, \dots, k$. So optimizing over $L_{un}(f)$ can reach our goal for contrastive learning.

Following Saunshi et al. (2019), the unsupervised risk can be described by the following equation:

$$L_{un}(f) = \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \mathbb{E}_{x, x^+ \sim \mathcal{D}_{c^+}, x_i^- \sim \mathcal{D}_{c_i^-}} [\ell(\{f(x)^T(f(x^+) - f(x_i^-))\}_{i=1}^k)].$$

The Empirical Risk Minimization (ERM) algorithm is used to find a function $\hat{f}_{ERM} \in \arg \min_{f \in \mathcal{F}} \hat{L}_{un}(f)$ that minimizes the empirical unsupervised risk. The function \hat{f}_{ERM} can be subsequently used for supervised linear classification tasks.

3.4 Adversarial Setup

A key question in adversarial contrastive learning is that of how to define the adversary sample in contrastive learning. From a representation perspective, one can find a point \tilde{x} that is close to x and keeps the feature $f(\tilde{x})$ both as far from $f(x^+)$ as possible and as close to the feature of some negative sample $f(x_i^-)$ as possible. Inspired by this intuition, we define the Contrastive Adversary Sample as follows:

Definition 5 (\mathcal{U} -Contrastive Adversary Sample) *Given a neighborhood $\mathcal{U}(x)$ of $x, x^+ \sim \mathcal{D}_{c^+}, x_i^- \sim \mathcal{D}_{c_i^-}$ for $i = 1, \dots, k$, we define the \mathcal{U} -Contrastive Adversary Sample of x as follows:*

$$\tilde{x} := \arg \sup_{x' \in \mathcal{U}(x)} \mathbb{E}_{x^+ \sim \mathcal{D}_{c^+}, x_i^- \sim \mathcal{D}_{c_i^-}} [\ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k)].$$

In this article, we suppose that the loss function ℓ is convex. By the subadditivity of sup, we have, $\forall f \in \mathcal{F}$:

$$\sup_{\substack{x' \in \mathcal{U}(x) \\ x_i^- \sim \mathcal{D}_{c_i^-}}} \mathbb{E}_{x^+ \sim \mathcal{D}_{c^+}} [\ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k)] \leq \mathbb{E}_{\substack{x^+ \sim \mathcal{D}_{c^+} \\ x_i^- \sim \mathcal{D}_{c_i^-}}} [\sup_{x' \in \mathcal{U}(x)} \ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k)]. \quad (1)$$

In the following sections, we analyze the theoretical properties of the right-hand side of (1), which can be easily optimized in practice by Adversarial Empirical Risk Minimization (AERM).

The \mathcal{U} -Adversarial Unsupervised Risk and its surrogate risk can be defined as below.

Definition 6 (\mathcal{U} -Adversarial Unsupervised Risk) *Given a neighborhood $\mathcal{U}(x)$ of x , the \mathcal{U} -Adversarial Unsupervised Risk of a presentation function f is defined as follows:*

$$\tilde{L}_{un}(f) := \mathbb{E}_{\substack{c^+, c_i^- \\ \sim \rho^{k+1}}} \mathbb{E}_{x \sim \mathcal{D}_{c^+}} [\sup_{x' \in \mathcal{U}(x)} \mathbb{E}_{\substack{x^+ \sim \mathcal{D}_{c^+} \\ x_i^- \sim \mathcal{D}_{c_i^-}}} [\ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k)]]. \quad (2)$$

Moreover, the surrogate risk of (2) is as follows:

$$\tilde{L}_{sun}(f) = \mathbb{E}_{\substack{c^+, c_i^- \\ \sim \rho^{k+1}}} \mathbb{E}_{\substack{x, x^+ \sim \mathcal{D}_{c^+}^2 \\ x_i^- \sim \mathcal{D}_{c_i^-}}} \sup_{x' \in \mathcal{U}(x)} \ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k).$$

By (1), we have $\tilde{L}_{un}(f) \leq \tilde{L}_{sun}(f)$ for any $f \in \mathcal{F}$. The Adversarial Supervised Risk of a classifier g for a task \mathcal{T} is defined as follows:

$$\tilde{L}_{sup}(\mathcal{T}, g) := \mathbb{E}_{(x, c) \sim \mathcal{D}_{\mathcal{T}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(\{g(x)_c - g(x)_{c'}\}_{c' \neq c}) \right].$$

The Adversarial Supervised Risk of a representation function f for a task \mathcal{T} is defined as follows:

$$\tilde{L}_{sup}(\mathcal{T}, f) := \inf_{W \in \mathbb{R}^{(k+1) \times n}} \tilde{L}_{sup}(\mathcal{T}, Wf). \quad (3)$$

For the mean classifier, we define:

$$\tilde{L}_{sup}^{\mu}(\mathcal{T}, f) := \tilde{L}_{sup}(\mathcal{T}, W^{\mu}f)$$

The Average Adversarial Supervised Risk for a representation function is as defined below.

Definition 7 (Average \mathcal{U} -Adversarial Supervised Risk)

$$\tilde{L}_{sup}(f) := \mathbb{E}_{\{c_i\}_{i=1}^{k+1} \sim \rho^{k+1}} [\tilde{L}_{sup}(\{c_i\}_{i=1}^{k+1}, f) | c_i \neq c_j].$$

For the mean classifier, we have the following:

$$\tilde{L}_{sup}^{\mu}(f) := \mathbb{E}_{\{c_i\}_{i=1}^{k+1} \sim \rho^{k+1}} [\tilde{L}_{sup}^{\mu}(\{c_i\}_{i=1}^{k+1}, f) | c_i \neq c_j].$$

4. Theoretical Analysis for Adversarial Contrastive Learning

This section presents some theoretical results for adversarial contrastive learning.

4.1 One Negative Sample Case

Let $\tau = \mathbb{P}_{c, c' \sim \rho^2}[c = c']$, σ be an M -dimensional Rademacher random vector with i.i.d. entries, define $(g_f)_{|\mathcal{S}} = (g_f(z_1), \dots, g_f(z_M))$ and $\mathcal{R}_{\mathcal{S}}(\mathcal{G}) := \mathbb{E}_{\sigma \sim \{\pm 1\}^M} \left[\sup_{f \in \mathcal{F}} \langle \sigma, (g_f)_{|\mathcal{S}} \rangle \right]$ where

$\mathcal{G} := \{g_f(x, x^+, x_1^-, \dots, x_k^-) = \sup_{x' \in \mathcal{U}(x)} \ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k) | f \in \mathcal{F}\}$, let $\hat{f} \in \arg \min_{f \in \mathcal{F}} \widehat{\tilde{L}}_{sun}(f)$ where $\widehat{\tilde{L}}_{sun}(f)$ is the empirical counterpart of $\tilde{L}_{sun}(f)$, we have:

Theorem 8 (The proof can be found in the Appendix A.1) *Let $\ell : \mathbb{R}^k \rightarrow \mathbb{R}$ be bounded by B . Then, for any $\delta \in (0, 1)$, with a probability of at least $1 - \delta$ over the choice of the training set $\mathcal{S} = \{(x_j, x_j^+, x_j^-)\}_{j=1}^M$, for any $f \in \mathcal{F}$:*

$$\tilde{L}_{sup}(\hat{f}) \leq \tilde{L}_{sup}^\mu(\hat{f}) \leq \frac{1}{1-\tau}(\tilde{L}_{sun}(f) - \tau\ell(0)) + \frac{1}{1-\tau}AG_M,$$

where

$$AG_M = O\left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B\sqrt{\frac{\log\frac{1}{\delta}}{M}}\right). \quad (4)$$

Remark 9 *Theorem 8 shows that when the hypothesis class \mathcal{F} is rich enough to contain some f with low surrogate adversarial unsupervised risk, the empirical minimizer of the surrogate adversarial unsupervised risk will then obtain good robustness on the supervised downstream task.*

Note that we can take $f = \hat{f}$ in the upper bound of Theorem 8 and get a bound $\tilde{L}_{sup}(\hat{f}) \leq \frac{1}{1-\tau}(\tilde{L}_{sun}(\hat{f}) - \tau\ell(0)) + \frac{1}{1-\tau}AG_M$. Then we can see that if the output of AERM (i.e., \hat{f}) achieves small unsupervised adversarial risk, then \hat{f} is a robust feature extractor such that it can achieve good robustness after fine-tuning on the downstream tasks.

Theorem 8 gives a relationship between the robustness of the contrastive (upstream) task and the robustness of the downstream classification tasks and explains why adversarial contrastive learning can help improve the robustness of the downstream task, as shown empirically in Kim et al. (2020); Ho and Vasconcelos (2020); Jiang et al. (2020).

4.2 Blocks of Similar Points

Saunshi et al. (2019) show a refined method that operates by using blocks of similar data and determine that the method achieves promising performance both theoretically and empirically. We adapt this method to adversarial contrastive learning.

Specifically, we sample $(b + 1)$ i.i.d. similar samples x, x_1^+, \dots, x_b^+ from $c^+ \sim \rho$ and b negative i.i.d. samples from $c^- \sim \rho$. The block adversarial contrastive learning risk is as follows:

$$\tilde{L}_{sun}^{block}(f) := \mathbb{E} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) \right],$$

and its empirical counterpart is as follows:

$$\hat{\tilde{L}}_{sun}^{block}(f) := \frac{1}{M} \sum_{i=1}^M \left[\sup_{x' \in \mathcal{U}(x_i)} \ell \left(f(x')^T \left(\frac{\sum_{j=1}^b f(x_{ij}^+)}{b} - \frac{\sum_{j=1}^b f(x_{ij}^-)}{b} \right) \right) \right],$$

Theorem 10 (The proof can be found in the Appendix A.2) *For any $f \in \mathcal{F}$, we have:*

$$\tilde{L}_{sup}(f) \leq \frac{1}{1-\tau} \left(\tilde{L}_{sun}^{block}(f) - \tau\ell(0) \right) \leq \frac{1}{1-\tau} \left(\tilde{L}_{sun}(f) - \tau\ell(0) \right).$$

Remark 11 *Theorem 10 shows that using blocks of similar data yields a tighter upper bound for the adversarial supervised risk than in the case for pairs of similar data. Theorem 10 implies that using the blocks in adversarial contrastive learning may improve the robust performance of the downstream tasks; this will be verified by the empirical results in §6.*

4.3 Multiple Negative Sample Case

This subsection extends our results to k negative samples. To achieve this, more definitions are required. Let $[k]$ denote the set $\{1, 2, \dots, k\}$.

Definition 12 *We define a distribution \mathcal{D} over the supervised tasks as follows: First, sample $k + 1$ classes (allow repetition) $c^+, c_1^-, \dots, c_k^- \sim \rho^{k+1}$, conditioned on the event that $c_i^- \neq c^+, \forall i \in [k]$. Then, set the task \mathcal{T} as the set of distinct classes in $\{c^+, c_1^-, \dots, c_k^-\}$.*

Definition 13 *The Average \mathcal{U} -Adversarial Supervised Risk of a representation function $f \in \mathcal{F}$ over \mathcal{D} is defined as follows:*

$$\tilde{\mathcal{L}}_{sup}(f) := \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\tilde{L}_{sup}(\mathcal{T}, f) \right].$$

Let $E_{distinct}$ be the event such that $\{c^+, c_1^-, \dots, c_k^-\}$ is distinct and $p = \mathbb{P}_{(c^+, c_1^-, \dots, c_k^-) \sim \mathcal{D}} [E_{distinct}]$. For any $f \in \mathcal{F}$, we have (The proof can be found in the Appendix A.3):

$$\tilde{L}_{sup}(f) \leq \frac{\tilde{\mathcal{L}}_{sup}(f)}{p}. \quad (5)$$

From (5), we can turn to analyze the relation between $\tilde{\mathcal{L}}_{sup}(f)$ and $\tilde{L}_{sun}(f)$ in the multiple negative sample case. Our Theorem 16 handles $\tilde{\mathcal{L}}_{sup}(f)$ instead of $\tilde{L}_{sup}(f)$.

Assumption 14 *Assume that $\forall I_1, I_2 \subseteq [d]$ such that $I_1 \cup I_2 = [d]$, ℓ satisfies the following inequations:*

$$\ell(\{v_i\}_{i \in I_1}) \leq \ell(\{v_i\}_{i \in [d]}) \leq \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}), \quad (6)$$

$$\ell(\{v_i\}_{i \in I_2}) \leq \ell(\{v_i\}_{i \in [d]}) \leq \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}). \quad (7)$$

Proposition 15 (The proof can be found in the Appendix A.4) *The hinge loss and the logistic loss satisfy Assumption 14.*

If \mathcal{C} is finite, we obtain a simple (and informal) bound of $\tilde{\mathcal{L}}_{sup}(\hat{f})$, for the more complex case that allows infinite \mathcal{C} and the formal form of Theorem 16 (Theorem 32), please refer to the Appendix A.5.

Theorem 16 (The proof can be found in the Appendix A.6) *Suppose \mathcal{C} is finite, for any $c \in \mathcal{C}$, $\rho(c) > 0$, and ℓ satisfies Assumption 14. Then, with a probability of at least $1 - \delta$ over the choice of the training set \mathcal{S} , $\forall f \in \mathcal{F}$:*

$$\tilde{\mathcal{L}}_{sup}(\hat{f}) \leq \alpha(\rho) \left(\tilde{L}_{sun}(f) + AG_M \right) - \beta.$$

5. Generalization Bounds for Example Hypothesis Classes

This section presents a concrete analysis of the Rademacher complexity for linear hypothesis class and multi-layer neural networks based on covering number (Wainwright, 2019, Definition 5.1). We first introduce some definitions and required lemmas.

Lemma 17 (Wainwright, 2019, Lemma 5.7, volume ratios and metric entropy) *Consider a pair of norms $\|\cdot\|$ and $\|\cdot\|'$ on \mathbb{R}^d , and let \mathbb{B} and \mathbb{B}' be their corresponding unit balls (i.e. $\mathbb{B} = \{\theta \in \mathbb{R}^d \mid \|\theta\| \leq 1\}$, with \mathbb{B}' similarly defined). The δ -covering number of \mathbb{B} in the $\|\cdot\|'$ -norm then obeys the following bounds (Wainwright, 2019, Lemma 5.7):*

$$\left(\frac{1}{\delta}\right)^d \frac{\text{vol}(\mathbb{B})}{\text{vol}(\mathbb{B}')} \leq \mathcal{N}(\delta; \mathbb{B}, \|\cdot\|') \leq \frac{\text{vol}(\frac{2}{\delta}\mathbb{B} + \mathbb{B}')}{\text{vol}(\mathbb{B}')},$$

where we define the Minkowski sum $A+B := \{a+b : a \in A, b \in B\}$, $\text{vol}(\mathbb{B}) := \int \mathbb{1}\{x \in \mathbb{B}\} dx$ is the volume of \mathbb{B} based on the Lebesgue measure, and $\mathcal{N}(\delta; \mathbb{B}, \|\cdot\|')$ is the δ -covering number of \mathbb{B} with respect to the norm $\|\cdot\|'$.

Lemma 18 (The proof can be found in the Appendix A.7) *Let $\mathbb{B}_p(r)$ be the p -norm ball in \mathbb{R}^d with radius r . The δ -covering number of $\mathbb{B}_p(r)$ with respect to $\|\cdot\|_p$ thus obeys the following bound:*

$$\mathcal{N}(\delta; \mathbb{B}_p(r), \|\cdot\|_p) \leq \left(1 + \frac{2r}{\delta}\right)^d.$$

Definition 19 (Wainwright, 2019, Definition 5.16, sub-Gaussian process) *A collection of zero-mean random variables $\{X_\theta, \theta \in \mathbb{T}\}$ is a sub-Gaussian process with respect to a metric ρ_X on \mathbb{T} if:*

$$\mathbb{E}[e^{\lambda(X_\theta - X_{\tilde{\theta}})}] \leq e^{\frac{\lambda^2 \rho_X^2(\theta, \tilde{\theta})}{2}}, \forall \theta, \tilde{\theta} \in \mathbb{T}, \lambda \in \mathbb{R}.$$

It is easy to prove that the Rademacher process (Wainwright, 2019, §5.2) satisfies the condition in Definition 19 with respect to the ℓ_2 -norm.

Lemma 20 (Wainwright, 2019, the Dudley's entropy integral bound) *Let $\{X_\theta, \theta \in \mathbb{T}\}$ be a zero-mean sub-Gaussian process with respect to the induced pseudometric ρ_X from Definition 19. Then, for any $\delta \in [0, D]$, we have:*

$$\mathbb{E} \left[\sup_{\theta, \tilde{\theta} \in \mathbb{T}} (X_\theta - X_{\tilde{\theta}}) \right] \leq 2 \mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathbb{T} \\ \rho_X(\gamma, \gamma') \leq \delta}} (X_\gamma - X_{\gamma'}) \right] + 32 \mathcal{J}(\delta/4; D). \quad (8)$$

Here, $D = \sup_{\theta, \tilde{\theta} \in \mathbb{T}} \rho_X(\theta, \tilde{\theta})$ and $\mathcal{J}(a; b) = \int_a^b \sqrt{\ln \mathcal{N}_X(u; \mathbb{T})} du$, where $\mathcal{N}_X(u; \mathbb{T})$ is the u -covering number of \mathbb{T} in the ρ_X -metric.

Remark 21 *Given $\theta_0 \in \mathbb{T}$, since $\mathbb{E}[X_{\theta_0}] := \mathbb{E}[\langle \theta_0, \sigma \rangle] = 0$, we have:*

$$\mathbb{E} \left[\sup_{\theta \in \mathbb{T}} X_\theta \right] = \mathbb{E} \left[\sup_{\theta \in \mathbb{T}} (X_\theta - X_{\theta_0}) \right] \leq \mathbb{E} \left[\sup_{\theta, \tilde{\theta} \in \mathbb{T}} (X_\theta - X_{\tilde{\theta}}) \right]. \quad (9)$$

Combining (8) with (9), we have:

$$\mathbb{E} \left[\sup_{\theta \in \mathbb{T}} X_\theta \right] \leq 2\mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathbb{T} \\ \rho_X(\gamma, \gamma') \leq \delta}} (X_\gamma - X_{\gamma'}) \right] + 32\mathcal{J}(\delta/4; D),$$

which can be used to draw the upper bound of $\mathcal{R}_S(\mathcal{G})$ by establishing a connection between the Rademacher process and the Rademacher complexity of the hypothesis classes when proper norm is chosen. For more details, please refer to the Appendix, details are in the proof of Theorem 22, Theorem 24 and Theorem 26.

To make Theorem 8 and Theorem 16 concrete, we need to upper bound $\mathcal{R}_S(\mathcal{G})$. Assume that loss ℓ is a non-increasing function; for example, hinge loss and logistic loss satisfy this assumption. Let $\mathcal{G} = \{g_f(x, x^+, x^-) = \sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F}\}$. Since ℓ is non-increasing, we have:

$$\mathcal{G} = \left\{ \ell \left(\min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \right) \mid f \in \mathcal{F} \right\}.$$

Let $\mathcal{H} = \left\{ \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F} \right\}$. Suppose ℓ is η -Lipschitz. By the Ledoux-Talagrand contraction inequality (Ledoux and Talagrand, 2013), we have:

$$\mathcal{R}_S(\mathcal{G}) \leq \eta \mathcal{R}_S(\mathcal{H}). \quad (10)$$

Thus, we only need to upper bound $\mathcal{R}_S(\mathcal{H})$. Let $\|A\|_{a,b}$ be the ℓ_b -norm of the ℓ_a -norm of the rows of A . Consider the training set $\mathcal{S} = \{(x_i, x_i^+, x_i^-)\}_{i=1}^M$. Let X be a matrix whose i th row is x_i . We define X^+ and X^- in a similar way. It is easy to see that $\forall p \geq 1, \forall i = 1, \dots, M, \|x_i\|_p \leq \|X\|_{p,\infty}$.

5.1 Linear Hypothesis Class

Let $\mathcal{F} = \{f: x \rightarrow Wx \mid W \in \mathbb{R}^{n \times m}, \|W\|_p \leq w\}$. To simplify notations, $\forall p \geq 1$ and $\frac{1}{p} + \frac{1}{p^*} = 1$, let

$$P = \max\{\|X\|_{p,\infty}, \|X^+\|_{p,\infty}, \|X^-\|_{p,\infty}\}, P^* = \max\{\|X\|_{p^*,\infty}, \|X^+\|_{p^*,\infty}, \|X^-\|_{p^*,\infty}\}. \quad (11)$$

We then have $\forall p \geq 1$,

$$\forall i = 1, \dots, M, \|x_i\|_p, \|x_i^+\|_p, \|x_i^-\|_p \leq P.$$

We now present the upper bound of $\mathcal{R}_S(\mathcal{H})$ under $\|\cdot\|_r$ attack.

Theorem 22 ($\mathcal{R}_S(\mathcal{H})$ under $\|\cdot\|_r$ attack for linear models) Consider the ℓ_r attack, i.e. let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_r \leq \epsilon\}$. We then have:

$$\mathcal{R}_S(\mathcal{H}) = O \left([PP^* + \epsilon R^* s(r^*, p, m)] \left[ms(p^*, p, m) w^2 \sqrt{M} \right] \right),$$

where $s(p, q, n) := n^{\max\{\frac{1}{p} - \frac{1}{q}, \frac{1}{q} - \frac{1}{p}\}}$, $\frac{1}{p} + \frac{1}{p^*} = 1$, $\frac{1}{r} + \frac{1}{r^*} = 1$, and R^* is defined similarly to (11).

The proof can be found in the Appendix A.8.

Remark 23 *Combining Theorem 22 with (4), we have:*

$$AG_M = O\left(\frac{[PP^* + \epsilon R^* s(r^*, p, m)] m \eta s(p^*, p, m) w^2 + B \sqrt{\log \frac{1}{\delta}}}{\sqrt{M}}\right). \quad (12)$$

5.2 Multi-layer Neural Network

In this section, we analyze fully connected multi-layer neural networks.

Suppose that $\mathcal{X} \subseteq \mathbb{R}^m$. Let $\mathcal{F} = \{W_d \sigma(W_{d-1} \sigma(\cdots \sigma(W_1 x))) \mid \|W_l\| \leq M_l, l = 1, \dots, d\}$, where $\|\cdot\|$ is the norm of the matrix and $\sigma(\cdot)$ is an elementwise L -Lipschitz function with $\sigma(0) = 0$ and $W_l \in \mathbb{R}^{h_l \times h_{l-1}}$, where $h_d = n, h_0 = m$. Assume ℓ is η -Lipschitz and non-increasing. From (10), we need only to bound the Rademacher complexity of \mathcal{H} .

We here consider two cases of the matrix norm $\|\cdot\|$. Let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$ for some $p \geq 1$.

5.2.1 FROBENIUS-NORM CASE.

We first consider using the Frobenius-Norm in the definition of the multi-layer neural networks hypothesis class \mathcal{F} .

Theorem 24 ($\mathcal{R}_S(\mathcal{H})$ under $\|\cdot\|_p$ attack for NNs under $\|\cdot\|_F$ constraint) *Let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$ (i.e. consider the ℓ_p attack), $\sigma(0) = 0$ with Lipschitz constant L and let $\mathcal{F} = \{W_d \sigma(W_{d-1} \sigma(\cdots \sigma(W_1 x))) \mid \|W_l\|_F \leq M_l^F, l = 1, \dots, d\}$. We then have:*

$$\mathcal{R}_S(\mathcal{H}) = O\left(\sqrt{\sum_{l=1}^d h_l h_{l-1} K \sqrt{d} \sqrt{M}}\right), \quad (13)$$

where

$$K = 2B_{X,\epsilon}^F \cdot (B_{X+}^F + B_{X-}^F),$$

where

$$B_{X,\epsilon}^F = L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2} - \frac{1}{p}}\right\} (\|X\|_{p,\infty} + \epsilon), \quad B_X^F = L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2} - \frac{1}{p}}\right\} \|X\|_{p,\infty}.$$

The proof can be found in the Appendix A.9.

Remark 25 *Combining Theorem 24 with (4), we have:*

$$AG_M = O\left(K \eta \sqrt{\frac{d \sum_{l=1}^d h_l h_{l-1}}{M}} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}}\right).$$

5.2.2 $\ell_{1,\infty}$ -NORM CASE

We consider the $\|\cdot\|_{1,\infty}$ norm constraint.

Theorem 26 ($\mathcal{R}_S(\mathcal{H})$ under $\|\cdot\|_p$ attack for NNs under $\|\cdot\|_{1,\infty}$ constraint) *Let $U(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$ (i.e. consider the ℓ_p attack), $\sigma(0) = 0$ with Lipschitz constant L ; moreover, let $\mathcal{F} = \{W_d \sigma(W_{d-1} \sigma(\dots \sigma(W_1 x))) \mid \|W_l\|_{1,\infty} \leq M_l^{1,\infty}, l=1, \dots, d\}$. We then have:*

$$\mathcal{R}_S(\mathcal{H}) = O \left(\sqrt{\sum_{l=1}^d h_l h_{l-1} \sqrt{d K_0 K_1} \sqrt{M}} \right),$$

where

$$K_0 = 2B_{X,\epsilon}^{1,\infty} \cdot (B'_{X^+} + B'_{X^-}), \quad K_1 = \frac{K_0}{2} + B'_{X,\epsilon} \cdot (B_{X^+}^{1,\infty} + B_{X^-}^{1,\infty}),$$

where

$$B'_{X,\epsilon} = L^{d-1} \prod_{l=1}^d h_l M_l^{1,\infty} m^{1-\frac{1}{p}} (\|X\|_{p,\infty} + \epsilon), \quad B'_X = L^{d-1} \prod_{l=1}^d h_l M_l^{1,\infty} m^{1-\frac{1}{p}} \|X\|_{p,\infty},$$

$$B_{X,\epsilon}^{1,\infty} = L^{d-1} \prod_{l=1}^d M_l^{1,\infty} (\|X\|_{p,\infty} + \epsilon), \quad B_X^{1,\infty} = L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \|X\|_{p,\infty}.$$

The proof can be found in the Appendix A.10.

Remark 27 *Combining Theorem 26 with (4), we have:*

$$AG_M = O \left(\eta \sqrt{\frac{d K_0 K_1 \sum_{l=1}^d h_l h_{l-1}}{M}} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}} \right).$$

Remark 28 *Our bound has important implications for the design of regularizers for adversarial contrastive learning. To achieve superior robust performance on the downstream tasks, the usual approach is to make $\|X\|_{p,\infty}$ small. For example, Pytorch scales the images to tensors with entries within the range $[0, 1]$. Moreover, Theorem 24 shows that we can take the norms of the layers as the regularizers to reduce the adversarial supervised risk.*

In our analysis for the Rademacher complexity, we consider models with norm-constrained weights, which means that the hypothesis class is uniformly Lipschitz, although the Lipschitz constant may be large (the product of maximal weight norms for the layers). One may wonder what will happen if we remove the constraints on the norm of the weights. For simplicity, let's consider a hypothesis class $\mathcal{H} \subseteq \{\pm 1\}^X$ for binary classification, we have:

$$\mathcal{R}_S(\mathcal{H}) = \mathbb{E}_{\sigma} \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \sigma_i h(x_i) \right],$$

where $\sigma_1, \dots, \sigma_n \in \{\pm 1\}$ are i.i.d. uniform random variables. We can regard $\sigma_1, \dots, \sigma_n$ as random labels that we need to fit by hypothesis from \mathcal{H} , so we can interpret $\mathcal{R}_S(\mathcal{H})$ as the

ATTACK	ϵ	TYPE	λ			
			0	0.002	0.05	0.2
PGD	0.01	CLEAN	75.73	74.8	74.5	75.67
		ADV	67.67	69.25	68.59	67.65
	0.02	CLEAN	53.11	55.72	55.73	55.72
		ADV	46.71	48.17	48.17	48.16
FGSM	0.01	CLEAN	74.42	76.13	76.12	76.11
		ADV	67.28	68.79	68.8	68.8
	0.02	CLEAN	54.28	54.29	54.28	66.68
		ADV	45.93	45.94	45.92	55.64

Table 1: Results of experiments on the regularizer. In this table, we list the clean accuracy (CLEAN) and adversarial accuracy (ADV) of the mean classifier under the PGD and FGSM attack with $\epsilon = 0.01$ and $\epsilon = 0.02$. λ is chosen from $\{0, 0.002, 0.05, 0.2\}$, and $\lambda = 0$ indicates no regularizer.

ability of \mathcal{H} to fit random ± 1 binary labels. Now let \mathcal{H} be the neural network, if we do not constrain the norm of the weights, theoretically, the universal approximation theorem (Maiorov and Pinkus, 1999) tells us that neural networks can fit any continuous function on a bounded input space, which means that in this case $\mathcal{R}_{\mathcal{S}}(\mathcal{H}) \approx 1$, leading to vacuous bounds in the binary classification case; experimentally, Zhang et al. (2017) show that deep neural networks easily fit random labels.

From another perspective, to derivate an upper bound for $\mathcal{R}_{\mathcal{S}}(\mathcal{H})$ by covering number, we need to find a δ -covering set for \mathcal{H} under some metric $\rho(\cdot, \cdot)$. If the weights of the layers are not bounded, we can not cover \mathcal{H} by a finite subset of \mathcal{H} , the δ -covering number of \mathcal{H} under $\rho(\cdot, \cdot)$ will be infinite, which means that Lemma 20 does not hold. So it is difficult to go beyond the Lipschitz network.

6. Experiments

In this section, we conduct several experiments to **support our theory**. We emphasize that we are not proposing a method to try to get better robustness on the downstream tasks, we do the experiments to verify our two claims from our theoretical results: (1) As shown in Remark 11, using the blocks may improve the robust performance; (2) As shown in Remark 28, using the norms of the layers of the neural networks as the regularizer may help improve the robust performance.

Data sets. We use two data sets (Krizhevsky and Hinton, 2009) in our experiments: (1) the CIFAR-10 data set and (2) the CIFAR-100 data set. CIFAR-10 contains 50000/10000 train/test images with size 32×32 , which are categorized into 10 classes. CIFAR-100 contains 50000/10000 train/test images with size 32×32 , which are categorized into 100 classes.

Model. We use a neural network with two convolutional layers and one fully connected layer. Following He et al. (2020), we use the Stochastic Gradient Descent (SGD) optimizer




Figure 1: The effect of block size on the accuracy. In the figure, we show the clean accuracy and the adversarial accuracy of the mean classifier under PGD and FGSM attack with $\epsilon = 0.01$ and $\epsilon = 0.02$. The block size is chosen from $\{1, 2, 4, 6, 8, 10\}$. **(a) The influence on the clean accuracy; (b) The influence on the adversarial accuracy.**

with momentum 0.9 but set the weight decay to be 5×10^{-4} and the learning rate to be 0.001.

Evaluation of robustness. For representation f , we first calculate $\hat{u}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} f(x_i)$ to estimate the c -th row of the mean classifier, where x_1, \dots, x_{n_c} are the data points with label c in our training set. Denote \widehat{W}^μ as the estimator of W , we use the robustness of the classifier $\widehat{W}^\mu f$ as an evaluation of the robustness of f on the downstream task.

We show the results for **CIFAR-10** here; the results for CIFAR-100 can be found in the Appendix B.

6.1 Improvement from the regularizer

Inspired by our bound (13) in Theorem 24 and Theorem 8, the adversarial supervised risk can be upper bounded by the sum of the adversarial unsupervised loss and AG_M , which is related to the maximal Frobenius-norm of the network layers. We choose to simultaneously optimize the contrastive upstream pre-train risk and the Frobenius norm of the parameters of the model. We set the norm of the parameters for the layers as a regularizer and test the performance of the mean classifier; here, W^μ is calculated by averaging all features of the training data set as done in Nozawa et al. (2020). We use a hyper-parameter λ to balance the trade-off of the the contrastive upstream pre-train risk and the Frobenius norm of the parameters of the model. We choose to minimize the following regularized empirical risk:

$$L(f) = \widehat{L}_{sun}(f) + \lambda N(f) \tag{14}$$

where $N(f)$ is a regularizer that constrains the Frobenius norm of the parameters of the model f , here we choose $N(f) = \sum_{l=1}^d \|W_l\|_F$ where $\|W_l\|_F$ is the Frobenius norm of the

parameters for the l -th layer of f and:

$$\widehat{L}_{sun}(f) = \frac{1}{M} \sum_{j=1}^M \max_{x'_j \in \mathcal{U}(x_j)} \ell(\{f(x'_j)^T(f(x_j^+) - f(x_{ji}^-))\}_{i=1}^k).$$

More details about our algorithm are in Algorithm 1.

Algorithm 1: The AERM algorithm for adversarial contrastive learning

Input : The training data set $\mathcal{S} = \{(x_j, x_j^+, x_{j1}^-, \dots, x_{jk}^-)\}_{j=1}^M$ sampled from $\mathcal{D}_{sim} \times \mathcal{D}_{neg}^k$; the hyper-parameter λ in (14); the adversarial perturbation \mathcal{U} ; learning rate α ; the total iteration number T ;

- 1 initialize θ^0 to be randomized parameters;
- 2 $t \leftarrow 0$;
- 3 **while** $t < T$ **do**
- 4 randomly sample a batch of data with size N : $\mathcal{B} \subseteq \mathcal{S}$;
- 5 $\tilde{\mathcal{B}} \leftarrow \emptyset$;
- 6 **for** $(x, x^+, x_1^-, \dots, x_k^-)$ **in** \mathcal{B} **do**
- 7 calculate adversarial example
 $\tilde{x} \in \arg \sup_{x' \in \mathcal{U}(x)} \ell(\{f_{\theta^t}(x')^T(f_{\theta^t}(x^+) - f_{\theta^t}(x_i^-))\}_{i=1}^k)$;
- 8 $\tilde{\mathcal{B}} \leftarrow \tilde{\mathcal{B}} \cup \{(\tilde{x}, x^+, x_1^-, \dots, x_k^-)\}$;
- 9 **end**
- 10 $L(f_{\theta^t}) = \frac{1}{N} \sum_{(\tilde{x}, x^+, x_1^-, \dots, x_k^-) \in \tilde{\mathcal{B}}} \ell(\{f_{\theta^t}(\tilde{x})^T(f_{\theta^t}(x^+) - f_{\theta^t}(x_i^-))\}_{i=1}^k) + \lambda N(f_{\theta^t})$;
- 11 $\theta^{t+1} \leftarrow \theta^t - \alpha \nabla_{\theta^t} L(f_{\theta^t})$;
- 12 $t \leftarrow t + 1$;
- 13 **end**

Output : The feature extractor $f_{\theta^T} \in \mathcal{F}$ that tries to minimize (14);

The results are shown in Table 1. From Table 1, we can see that the F -norm regularizer can improve the adversarial accuracy (the prediction performance of a model on adversarial examples generated by attacker) of the mean classifier, which is in line with our Theorem 24.

6.2 Effect of block size

To verify Theorem 10, we analyze the effect of block size on the adversarial accuracy of the mean classifier. Figure 1 presents the results for clean accuracy and adversarial accuracy, respectively. From Figure 1, we can see that a larger block size will yield better adversarial accuracy. The results are consistent with Theorem 10: as the block size grows, Theorem 10 shows that we are optimizing a tighter bound, which leads to better performance as shown in Figure 1.

7. Conclusion

This paper studies the generalization performance of adversarial contrastive learning. We first extend the contrastive learning framework to the adversarial case, then we upper bound the average adversarial risk of the downstream tasks with the adversarial unsupervised risk of the upstream task and an adversarial Rademacher complexity term. Furthermore, we provide the upper bound of the adversarial Rademacher complexity for linear models and multi-layer neural networks. Finally, we conduct several experiments and the experimental results are consistent with our theory.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant 61976161.

Appendix A. Proofs

In this section, we display the proofs of our theorems, lemmas and corollaries. For reading convenience, we will restate the theorem before proving.

A.1 Proof of Theorem 8

In the below, we present some useful lemmas that will be used in the proofs of our main theorems.

Lemma 29 *For any $f \in \mathcal{F}$, we have:*

$$\tilde{L}_{sup}(f) \leq \tilde{L}_{sup}^\mu(f) \leq \frac{1}{1-\tau}(\tilde{L}_{un}(f) - \tau\ell(0)) \leq \frac{1}{1-\tau}(\tilde{L}_{sun}(f) - \tau\ell(0)). \quad (\text{A.1})$$

Remark 30 $\hat{L}_{sun}(f)$ denotes the empirical $\tilde{L}_{sun}(f)$ and $\hat{f} \in \arg \min_{f \in \mathcal{F}} \hat{L}_{sun}(f)$. Applying

Lemma 29 to \hat{f} shows that, if we can train a robust feature extractor with low surrogate adversarial unsupervised risk, we can obtain a robust classifier with low adversarial supervised risk on the downstream task.

Lemma 31 *Let $\ell : \mathbb{R}^k \rightarrow \mathbb{R}$ be bounded by B . Then, for any $\delta \in (0, 1)$, with a probability of at least $1 - \delta$ over the choice of the training set $\mathcal{S} = \{(x_j, x_j^+, x_j^-, \dots, x_{jk}^-)\}_{j=1}^M = \{z_j\}_{j=1}^M$, for any $f \in \mathcal{F}$:*

$$\tilde{L}_{sun}(\hat{f}) \leq \tilde{L}_{sun}(f) + AG_M,$$

where $AG_M = O\left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B\sqrt{\frac{\log \frac{1}{\delta}}{M}}\right)$, $\mathcal{R}_{\mathcal{S}}(\mathcal{G}) := \mathbb{E}_{\sigma \sim \{\pm 1\}^M} \left[\sup_{f \in \mathcal{F}} \langle \sigma, (g_f)_{|\mathcal{S}} \rangle \right]$, and $\mathcal{G} := \{g_f(x, x^+, x_1^-, \dots, x_k^-) = \sup_{x' \in \mathcal{U}(x)} \ell(\{f(x')^T(f(x^+) - f(x_i^-))\}_{i=1}^k) \mid f \in \mathcal{F}\}$, where σ is an M -dimensional Rademacher random vector with i.i.d. entries and $(g_f)_{|\mathcal{S}} = (g_f(z_1), \dots, g_f(z_M))$.

Proof [Proof of Lemma 29] By the definition of $\tilde{L}_{sup}(\mathcal{T}, f)$, i.e., (3), it's obvious that $\tilde{L}_{sup}(f) \leq \tilde{L}_{sup}^\mu(f)$, so we only need to prove the second part of (A.1). From definition (2), we have, $\forall f \in \mathcal{F}$:

$$\begin{aligned} \tilde{L}_{un}(f) &= \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \left(\mathbb{E}_{\substack{x^+ \sim \mathcal{D}_{c^+} \\ x^- \sim \mathcal{D}_{c^-}}} (\ell(f(x')^T(f(x^+) - f(x^-)))) \right) \right] \right\} \\ &\stackrel{(i)}{\geq} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^+} - \mu_{c^-})) \right] \right\} \\ &\stackrel{(ii)}{=} (1-\tau) \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^+} - \mu_{c^-})) \right] \middle| c^+ \neq c^- \right\} + \tau\ell(0), \end{aligned} \quad (\text{A.2})$$

where (i) comes from the convexity of ℓ and Jensen's Inequality, and (ii) comes from the property of conditional expectations. Then we have:

$$\begin{aligned}
 & \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^+} - \mu_{c^-})) \right] \middle| c^+ \neq c^- \right\} \\
 & \stackrel{(i)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ D_{\mathcal{T}}(c^+) \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^+} - \mu_{c^-})) \right] \right. \\
 & \quad \left. + D_{\mathcal{T}}(c^-) \mathbb{E}_{x \sim \mathcal{D}_{c^-}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^-} - \mu_{c^+})) \right] \middle| c^+ \neq c^- \right\} \quad (\text{A.3}) \\
 & \stackrel{(ii)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ D_{\mathcal{T}}(c^+) \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(g(x')_{c^+} - g(x')_{c^-}) \right] \right. \\
 & \quad \left. + D_{\mathcal{T}}(c^-) \mathbb{E}_{x \sim \mathcal{D}_{c^-}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(g(x')_{c^-} - g(x')_{c^+}) \right] \middle| c^+ \neq c^- \right\},
 \end{aligned}$$

where $\mathcal{T} = \{c^+, c^-\}$, $g(x) = \begin{bmatrix} \mu_{c^+}^T f(x) \\ \mu_{c^-}^T f(x) \end{bmatrix} = W^\mu f(x)$ and (i) comes from the symmetry of c^+, c^- ; (ii) is directly from some linear algebras. From (A.3) we know that:

$$\begin{aligned}
 & \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell(f(x')^T(\mu_{c^+} - \mu_{c^-})) \right] \middle| c^+ \neq c^- \right\} \\
 & \stackrel{(i)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{c \sim \mathcal{D}_{\mathcal{T}}} \mathbb{E}_{x \sim \mathcal{D}_c} \left[\sup_{x' \in \mathcal{U}(x)} \ell(\{g(x')_c - g(x')_{c'}\}_{c' \neq c}) \right] \middle| c^+ \neq c^- \right\} \\
 & \stackrel{(ii)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left[\tilde{L}_{sup}(\{c^+, c^-\}, W^\mu f) \middle| c^+ \neq c^- \right] \\
 & \stackrel{(iii)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left[\tilde{L}_{sup}^\mu(\{c^+, c^-\}, f) \middle| c^+ \neq c^- \right] \stackrel{(iv)}{=} \tilde{L}_{sup}^\mu(f),
 \end{aligned} \quad (\text{A.4})$$

where (i) is due to the tower property of expectation; (ii) is obvious by the definition of $\tilde{L}_{sup}(\mathcal{T}, g)$; (iii) comes from the definition of $\tilde{L}_{sup}^\mu(\mathcal{T}, f)$ and (iv) is from the definition of $\tilde{L}_{sup}^\mu(f)$. Combine (A.2) with (A.4), we conclude that:

$$\tilde{L}_{un}(f) \geq (1 - \tau) \tilde{L}_{sup}^\mu(f) + \tau \ell(0), \forall f \in \mathcal{F}.$$

So we have:

$$(1 - \tau) \tilde{L}_{sup}(f) + \tau \ell(0) \leq (1 - \tau) \tilde{L}_{sup}^\mu(f) + \tau \ell(0) \leq \tilde{L}_{un}(f), \forall f \in \mathcal{F}. \quad (\text{A.5})$$

Rearranging (A.5) yields (A.1). ■

Proof [Proof of Lemma 31] Denote $(x, x^+, x_1^-, \dots, x_k^-)$ by z , then by the Theorem 3.3 in Mohri et al. (2012), we have: With probability at least $1 - \frac{\delta}{2}$ over the choice of the training set \mathcal{S} ,

$$\mathbb{E}_z \left[\frac{1}{B} g_f(z) \right] \leq \frac{1}{M} \sum_{i=1}^M \frac{1}{B} g_f(z_i) + \frac{2}{M} \mathcal{R}_{\mathcal{S}} \left(\frac{\mathcal{G}}{B} \right) + 3 \sqrt{\frac{\log \frac{4}{\delta}}{M}},$$

which is equivalent to:

$$\tilde{L}_{sun}(f) \leq \widehat{L}_{sun}(f) + \frac{2}{M} \mathcal{R}_{\mathcal{S}}(\mathcal{G}) + 3B \sqrt{\frac{\log \frac{4}{\delta}}{M}}, \forall f \in \mathcal{F}. \quad (\text{A.6})$$

Let $f^* = \arg \min_{f \in \mathcal{F}} \tilde{L}_{sun}(f)$, since $\mathbb{E}_{\mathcal{S}} \left[\widehat{L}_{sun}(f) \right] = \tilde{L}_{sun}(f)$ and ℓ is bounded by B , Hoeffding's inequality tells us that: $\forall f \in \mathcal{F}, \forall t \in \mathbb{R}$:

$$\mathbb{P} \left[\widehat{L}_{sun}(f) - \tilde{L}_{sun}(f) \geq t \right] \leq e^{-\frac{2Mt^2}{B^2}}.$$

Set $t = B \sqrt{\frac{\log \frac{2}{\delta}}{2M}}$, we have: $\forall f \in \mathcal{F}$, with probability at least $1 - \frac{\delta}{2}$,

$$\widehat{L}_{sun}(f) - \tilde{L}_{sun}(f) \leq B \sqrt{\frac{\log \frac{2}{\delta}}{2M}}. \quad (\text{A.7})$$

Combine (A.6) with (A.7), the union bound tells us that: $\forall f \in \mathcal{F}$, with probability at least $1 - \delta$ over the choice of the training set \mathcal{S} ,

$$\begin{aligned} \tilde{L}_{sun}(\widehat{f}) &\stackrel{(i)}{\leq} \widehat{L}_{sun}(\widehat{f}) + O \left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}} \right) \stackrel{(ii)}{\leq} \widehat{L}_{sun}(f^*) + O \left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}} \right) \\ &\stackrel{(iii)}{\leq} \tilde{L}_{sun}(f^*) + O \left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}} \right) \stackrel{(iv)}{\leq} \tilde{L}_{sun}(f) + O \left(\frac{\mathcal{R}_{\mathcal{S}}(\mathcal{G})}{M} + B \sqrt{\frac{\log \frac{1}{\delta}}{M}} \right), \end{aligned}$$

where (i) comes from (A.6); (ii) is directly from the fact that $\widehat{L}_{sun}(\widehat{f}) \leq \widehat{L}_{sun}(f^*)$, which is from the definition of \widehat{f} ; (iii) is a result of (A.7) and (iv) is obvious by the definition of f^* . ■

Theorem 8 Let $\ell : \mathbb{R}^k \rightarrow \mathbb{R}$ be bounded by B . Then, for any $\delta \in (0, 1)$, with a probability of at least $1 - \delta$ over the choice of the training set $\mathcal{S} = \{(x_j, x_j^+, x_j^-)\}_{j=1}^M$, for any $f \in \mathcal{F}$:

$$\tilde{L}_{sup}(\widehat{f}) \leq \tilde{L}_{sup}^{\mu}(\widehat{f}) \leq \frac{1}{1-\tau} (\tilde{L}_{sun}(f) - \tau \ell(0)) + \frac{1}{1-\tau} AG_M.$$

Proof From Lemma 29 we know that:

$$\tilde{L}_{sup}(\widehat{f}) \leq \tilde{L}_{sup}^{\mu}(\widehat{f}) \leq \frac{1}{1-\tau} (\tilde{L}_{un}(\widehat{f}) - \tau \ell(0)) \leq \frac{1}{1-\tau} (\tilde{L}_{sun}(\widehat{f}) - \tau \ell(0)),$$

Then Lemma 31 directly yields the result we need. ■

A.2 Proof of Theorem 10

Theorem 10 For any $f \in \mathcal{F}$, we have:

$$\tilde{L}_{sup}(f) \leq \frac{1}{1-\tau} \left(\tilde{L}_{sun}^{block}(f) - \tau \ell(0) \right) \leq \frac{1}{1-\tau} \left(\tilde{L}_{sun}(f) - \tau \ell(0) \right).$$

Proof By the convexity of ℓ and Jensen's inequality, we have: $\forall x', x_i^+, x_i^-$:

$$\begin{aligned} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) &= \ell \left(\frac{1}{b} \sum_{i=1}^b f(x')^T (f(x_i^+) - f(x_i^-)) \right) \\ &\leq \frac{1}{b} \sum_{i=1}^b \ell (f(x')^T (f(x_i^+) - f(x_i^-))). \end{aligned}$$

Take maximization about x' both sides, we have:

$$\begin{aligned} \sup_{x' \in \mathcal{U}(x)} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) &\leq \frac{1}{b} \sup_{x' \in \mathcal{U}(x)} \left[\sum_{i=1}^b \ell (f(x')^T (f(x_i^+) - f(x_i^-))) \right] \\ &\leq \frac{1}{b} \sum_{i=1}^b \left[\sup_{x' \in \mathcal{U}(x)} \ell (f(x')^T (f(x_i^+) - f(x_i^-))) \right]. \end{aligned}$$

Taking expectations both sides yields:

$$\begin{aligned} \tilde{L}_{sun}^{block}(f) &= \mathbb{E}_{x, x_i^+, x_i^-} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) \right] \\ &\leq \mathbb{E}_{x, x_i^+, x_i^-} \left\{ \frac{1}{b} \sum_{i=1}^b \left[\sup_{x' \in \mathcal{U}(x)} \ell (f(x')^T (f(x_i^+) - f(x_i^-))) \right] \right\} \\ &= \mathbb{E}_{x, x^+, x^-} \left[\sup_{x' \in \mathcal{U}(x)} \ell (f(x')^T (f(x^+) - f(x^-))) \right] = \tilde{L}_{sun}(f), \end{aligned}$$

which proves the second inequality. For the first inequality, we have, $\forall f \in \mathcal{F}$:

$$\begin{aligned}
 \tilde{L}_{sun}^{block}(f) &= \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{\substack{x, x_i^+ \sim \mathcal{D}_{c^+}^{b+1} \\ x_i^- \sim \mathcal{D}_{c^-}^b}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) \right] \right\} \\
 &\stackrel{(i)}{\geq} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left\{ \sup_{\substack{x' \in \mathcal{U}(x) \\ x_i^+ \sim \mathcal{D}_{c^+}^b \\ x_i^- \sim \mathcal{D}_{c^-}^b}} \ell \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) \right\} \right\} \\
 &\stackrel{(ii)}{\geq} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left\{ \sup_{x' \in \mathcal{U}(x)} \ell \left[\mathbb{E}_{\substack{x_i^+ \sim \mathcal{D}_{c^+}^b \\ x_i^- \sim \mathcal{D}_{c^-}^b}} \left(f(x')^T \left(\frac{\sum_{i=1}^b f(x_i^+)}{b} - \frac{\sum_{i=1}^b f(x_i^-)}{b} \right) \right) \right] \right\} \right\} \\
 &\stackrel{(iii)}{=} \mathbb{E}_{c^+, c^- \sim \rho^2} \left\{ \mathbb{E}_{x \sim \mathcal{D}_{c^+}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(f(x')^T (\mu_{c^+} - \mu_{c^-}) \right) \right] \right\} \stackrel{(iv)}{=} (1 - \tau) \tilde{L}_{sup}^\mu(f) + \tau \ell(0),
 \end{aligned}$$

where (i) and (ii) are direct result of Jensen's Inequality and convexity of maximization function and ℓ ; (iii) is from the linearity of expectation and the last equality follows the same argumentation as in (A.4), which proves the Theorem. \blacksquare

A.3 Proof of (5)

Proof For any $f \in \mathcal{F}$, we have:

$$\begin{aligned}
 \tilde{\mathcal{L}}_{sup}(f) &= \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\tilde{L}_{sup}(\mathcal{T}, f) \right] \stackrel{(i)}{=} p \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\tilde{L}_{sup}(\mathcal{T}, f) \mid E_{distinct} \right] + (1-p) \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\tilde{L}_{sup}(\mathcal{T}, f) \mid \bar{E}_{distinct} \right] \\
 &\geq p \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\tilde{L}_{sup}(\mathcal{T}, f) \mid E_{distinct} \right] \stackrel{(ii)}{=} p \tilde{L}_{sup}(f),
 \end{aligned}$$

where $E_{distinct}$ is the event that $\{c^+, c_1^-, \dots, c_k^-\}$ is distinct and $p = \frac{P[E_{distinct}]}{(c^+, c_1^-, \dots, c_k^-) \sim \mathcal{D}}$

and (i) is from the property of conditional expectation and (ii) comes from the definition of $\tilde{L}_{sup}(f)$. \blacksquare

A.4 Proof of Proposition 15

Proposition 15 *The hinge loss and the logistic loss satisfy Assumption 14.*

Proof Since I_1 and I_2 are symmetric, We only need to prove (6). **For the Hinge Loss:**

$$\ell(\{v_i\}_{i \in I_1}) = \max \left\{ 0, 1 + \max_{i \in I_1} \{-v_i\} \right\} \stackrel{(i)}{\leq} \max \left\{ 0, 1 + \max_{i \in [d]} \{-v_i\} \right\} = \ell(\{v_i\}_{i \in [d]}),$$

where (i) is from the fact that $I_1 \subseteq [d]$, so the first inequality is proved, for the second one:

$$\ell(\{v_i\}_{i \in [d]}) = \max \left\{ 0, 1 + \max_{i \in [d]} \{-v_i\} \right\} = \max \left\{ 0, 1 + \max_{i \in I_1 \cup I_2} \{-v_i\} \right\},$$

where the last equality is directly from the definition of I_1 and I_2 .

1. if $1 + \max_{i \in I_1 \cup I_2} \{-v_i\} \leq 0$, then $\ell(\{v_i\}_{i \in I_1 \cup I_2}) = 0$, since Hinge Loss is non-negative, we have:

$$\ell(\{v_i\}_{i \in I_1 \cup I_2}) = 0 + 0 \leq \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}).$$

2. if $1 + \max_{i \in I_1 \cup I_2} \{-v_i\} > 0$

- (a) if $1 + \max_{i \in I_1} \{-v_i\} \leq 0$, then $\ell(\{v_i\}_{i \in I_1}) = \max \left\{ 0, 1 + \max_{i \in I_1} \{-v_i\} \right\} = 0$. So we have:

$$\begin{aligned} \ell(\{v_i\}_{i \in I_1 \cup I_2}) &= \max \left\{ 0, 1 + \max_{i \in I_1 \cup I_2} \{-v_i\} \right\} = \max \left\{ 0, 1 + \max_{i \in I_2} \{-v_i\} \right\} \\ &= 0 + \ell(\{v_i\}_{i \in I_2}) = \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}). \end{aligned}$$

- (b) if $1 + \max_{i \in I_2} \{-v_i\} \leq 0$, by the same discussion of (a), we have:

$$\ell(\{v_i\}_{i \in I_1 \cup I_2}) \leq \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}).$$

- (c) if $1 + \max_{i \in I_1} \{-v_i\} > 0$ and $1 + \max_{i \in I_2} \{-v_i\} > 0$,

$$\begin{aligned} \ell(\{v_i\}_{i \in I_1 \cup I_2}) &= \max \left\{ 0, 1 + \max_{i \in I_1 \cup I_2} \{-v_i\} \right\} \leq \max \left\{ 0, 1 + \max_{i \in I_1} \{-v_i\} + 1 + \max_{i \in I_2} \{-v_i\} \right\} \\ &\leq \max \left\{ 0, 1 + \max_{i \in I_1} \{-v_i\} \right\} + \max \left\{ 0, 1 + \max_{i \in I_2} \{-v_i\} \right\} \\ &= \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}). \end{aligned}$$

So the second inequality is proved. **For the Logistic Loss:**

$$\ell(\{v_i\}_{i \in I_1}) = \log_2 \left(1 + \sum_{i \in I_1} e^{-v_i} \right) \leq \log_2 \left(1 + \sum_{i \in [d]} e^{-v_i} \right) = \ell(\{v_i\}_{i \in [d]}).$$

So the first inequality is proved. For the second one:

$$\begin{aligned}
 \ell(\{v_i\}_{i \in I_1 \cup I_2}) &= \log_2 \left(1 + \sum_{i \in I_1 \cup I_2} e^{-v_i} \right) \leq \log_2 \left(1 + \sum_{i \in I_1} e^{-v_i} + \sum_{i \in I_2} e^{-v_i} \right) \\
 &\leq \log_2 \left(1 + \sum_{i \in I_1} e^{-v_i} + \sum_{i \in I_2} e^{-v_i} + \left(\sum_{i \in I_1} e^{-v_i} \right) \left(\sum_{i \in I_2} e^{-v_i} \right) \right) \\
 &= \log_2 \left[\left(1 + \sum_{i \in I_1} e^{-v_i} \right) \left(1 + \sum_{i \in I_2} e^{-v_i} \right) \right] \\
 &= \log_2 \left(1 + \sum_{i \in I_1} e^{-v_i} \right) + \log_2 \left(1 + \sum_{i \in I_2} e^{-v_i} \right) = \ell(\{v_i\}_{i \in I_1}) + \ell(\{v_i\}_{i \in I_2}),
 \end{aligned}$$

which proves the second inequality. \blacksquare

A.5 Proof of Theorem 32

We here introduce some further notations, which will be used in our main results. For a tuple $(c^+, c_1^-, \dots, c_k^-)$, we define $I^+ := \{i \in [k] | c_i^- = c^+\}$ and Q as the set of distinct classes in $(c^+, c_1^-, \dots, c_k^-)$. We define $p_{\max}(\mathcal{T}) := \max_c \mathcal{D}_{\mathcal{T}}(c)$, $\tau_k := \mathbb{P}_{c^+, c_i^- \sim \rho^{k+1}} [I^+ \neq \emptyset]$, while $\ell_N(\vec{0})$ is defined as the loss of the N -dimensional zero vector. Let \mathcal{T} be a task sample from distribution \mathcal{D} and $\rho^+(\mathcal{T})$ be a distribution of c^+ when $(c^+, c_1^-, \dots, c_k^-)$ are sampled from ρ^{k+1} conditioned on $Q = \mathcal{T}$ and $I^+ = \emptyset$ and $\rho_{\min}^+(\mathcal{T}) := \min_{c \in \mathcal{T}} \rho^+(\mathcal{T})(c)$.

Theorem 32 *Assume that ℓ satisfies Assumption 14. With a probability of at least $1 - \delta$ over the choice of the training set \mathcal{S} , for any $f \in \mathcal{F}$, we have:*

$$\begin{aligned}
 \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{\text{sup}}(\mathcal{T}, \hat{f}) \right] &\leq \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{\text{sup}}^{\mu}(\mathcal{T}, \hat{f}) \right] \\
 &\leq \frac{1}{1 - \tau_k} \left(\tilde{L}_{\text{sun}}(f) + AG_M \right) - \frac{\tau_k}{1 - \tau_k} \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \left[\ell_{|I^+|}(\vec{0}) | I^+ \neq \emptyset \right],
 \end{aligned} \tag{A.8}$$

where $|I^+|$ is the cardinality of set I^+ .

Before proceeding with the proof of Theorem 32, we introduce some useful lemmas.

Lemma 33 *For any \mathcal{T} sampled from \mathcal{D} , we have: $\rho^+(\mathcal{T})(c) \geq \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \mathcal{D}_{\mathcal{T}}(c), \forall c \in \mathcal{T}$.*

Lemma 34 *Assume that f satisfies Assumption 14. For any $f \in \mathcal{F}$, we have:*

$$\begin{aligned}
 (1 - \tau_k) \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{\text{sup}}(\mathcal{T}, f) \right] &\leq (1 - \tau_k) \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{\text{sup}}^{\mu}(\mathcal{T}, f) \right] \\
 &\leq \tilde{L}_{\text{sun}}(f) - \tau_k \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \left[\ell_{|I^+|}(\vec{0}) | I^+ \neq \emptyset \right].
 \end{aligned}$$

Proof [Proof of Lemma 33] By the definition of $\rho^+(\mathcal{T})$ and $\rho_{\min}^+(\mathcal{T})$, we can find that:

$$\forall c \in \mathcal{T}, \rho_{\min}^+(\mathcal{T}) \leq \rho^+(\mathcal{T})(c).$$

So we have:

$$\forall c \in \mathcal{T}, \frac{\rho^+(\mathcal{T})(c)}{\rho_{\min}^+(\mathcal{T})} \geq 1. \quad (\text{A.9})$$

By the definition of $p_{\max}(\mathcal{T})$, we have:

$$\forall c \in \mathcal{T}, \frac{\mathcal{D}_{\mathcal{T}}(c)}{p_{\max}(\mathcal{T})} \leq 1. \quad (\text{A.10})$$

Combine (A.9) and (A.10), we have:

$$\forall c \in \mathcal{T}, \rho^+(\mathcal{T})(c) \geq \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \mathcal{D}_{\mathcal{T}}(c).$$

■

Proof [Proof of Lemma 34] By the definition of $\tilde{L}_{sun}(f)$, we have:

$$\begin{aligned} \tilde{L}_{sun}(f) &= \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x, x^+ \sim \mathcal{D}_{c^+}^2; x_i^- \sim \mathcal{D}_{c_i^-}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (f(x^+) - f(x_i^-)) \right\}_{i=1}^k \right) \right] \\ &= \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left\{ \mathbb{E}_{\substack{x^+ \sim \mathcal{D}_{c^+} \\ x_i^- \sim \mathcal{D}_{c_i^-}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (f(x^+) - f(x_i^-)) \right\}_{i=1}^k \right) \right] \right\} \\ &\stackrel{(i)}{\geq} \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left\{ \sup_{x' \in \mathcal{U}(x)} \left[\mathbb{E}_{\substack{x^+ \sim \mathcal{D}_{c^+} \\ x_i^- \sim \mathcal{D}_{c_i^-}}} \ell \left(\left\{ f(x')^T (f(x^+) - f(x_i^-)) \right\}_{i=1}^k \right) \right] \right\} \\ &\stackrel{(ii)}{\geq} \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_{c_i^-}) \right\}_{i=1}^k \right) \right] \triangleq R_1, \end{aligned} \quad (\text{A.11})$$

where (i), (ii) is from the Jensen's inequality and convexity of ℓ . Then we analyze lower bound of R_1 :

$$\begin{aligned}
 R_1 &\stackrel{(i)}{=} (1 - \tau_k) \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_{c_i^-}) \right\}_{i=1}^k \right) \middle| I^+ = \emptyset \right] \\
 &\quad + \tau_k \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_{c_i^-}) \right\}_{i=1}^k \right) \middle| I^+ \neq \emptyset \right] \\
 &\stackrel{(ii)}{\geq} (1 - \tau_k) \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim Q, c \neq c^+} \right) \middle| I^+ = \emptyset \right] \\
 &\quad + \tau_k \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_{c_i^-}) \right\}_{i=1}^k \right) \middle| I^+ \neq \emptyset \right] \\
 &\stackrel{(iii)}{\geq} (1 - \tau_k) \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim Q, c \neq c^+} \right) \middle| I^+ = \emptyset \right] \\
 &\quad + \tau_k \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\ell_{|I^+|}(\vec{0}) \middle| I^+ \neq \emptyset \right],
 \end{aligned} \tag{A.12}$$

where (i) comes from the property of conditional expectation, (ii) is a result of the fact that $Q \subseteq [k]$ and ℓ satisfies Assumption 14 and (iii) is from the fact that $[|I^+|] \subseteq [k]$ and ℓ satisfies

Assumption 14. Let $R_2 = \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim Q, c \neq c^+} \right) \middle| I^+ = \emptyset \right]$, we

have:

$$\begin{aligned}
 R_2 &= \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim Q, c \neq c^+} \right) \middle| I^+ = \emptyset \right] \\
 &\stackrel{(i)}{=} \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left\{ \mathbb{E}_{\substack{c^+, c_i^- \sim \rho^{k+1} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim Q, c \neq c^+} \right) \middle| Q = \mathcal{T}, I^+ = \emptyset \right] \right\} \\
 &\stackrel{(ii)}{=} \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left\{ \mathbb{E}_{\substack{c^+ \sim \rho^+(\mathcal{T}) \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim \mathcal{T}, c \neq c^+} \right) \right] \right\},
 \end{aligned} \tag{A.13}$$

where (i) is from the tower property of expectation and (ii) is directly obtained by the definition of $\rho^+(\mathcal{T})$. Let $R_3 = \mathbb{E}_{\substack{c^+ \sim \rho^+(\mathcal{T}) \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\left\{ f(x')^T (\mu_{c^+} - \mu_c) \right\}_{c \sim \mathcal{T}, c \neq c^+} \right) \right]$, we have:

$$\begin{aligned}
 R_3 &= \mathbb{E}_{\substack{c^+ \sim \rho^+(\mathcal{T}) \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\{f(x')^T (\mu_{c^+} - \mu_c)\}_{c \sim \mathcal{T}, c \neq c^+} \right) \right] \\
 &\stackrel{(i)}{\geq} \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \mathbb{E}_{\substack{c^+ \sim \mathcal{D}_{\mathcal{T}} \\ x \sim \mathcal{D}_{c^+}}} \left[\sup_{x' \in \mathcal{U}(x)} \ell \left(\{f(x')^T (\mu_{c^+} - \mu_c)\}_{c \sim \mathcal{T}, c \neq c^+} \right) \right] \quad (\text{A.14}) \\
 &\stackrel{(ii)}{=} \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{sup}^\mu(\mathcal{T}, f) \stackrel{(iii)}{\geq} \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{sup}(\mathcal{T}, f),
 \end{aligned}$$

where (i) is directly from Lemma 33; (ii) and (iii) are from the definition of $\tilde{L}_{sup}^\mu(\mathcal{T}, f)$ and $\tilde{L}_{sup}(\mathcal{T}, f)$, respectively. Combing (A.11) (A.12) (A.13) and (A.14) yields:

$$\begin{aligned}
 (1 - \tau_k) \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{sup}(\mathcal{T}, f) \right] &\leq (1 - \tau_k) \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{sup}^\mu(\mathcal{T}, f) \right] \\
 &\leq \tilde{L}_{sun}(f) - \tau_k \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \left[\ell_{|I^+|}(\vec{0}) | I^+ \neq \emptyset \right].
 \end{aligned}$$

■

Equipped with the above lemmas, now we can turn to the proof of Theorem 32.

Proof [Proof of Theorem 32] From Lemma 31 we know that with probability at least $1 - \delta$ over the choice of the training set \mathcal{S} , $\forall f \in \mathcal{F}$:

$$\tilde{L}_{sun}(\hat{f}) \leq \tilde{L}_{sun}(f) + AG_M.$$

Combing this with Lemma 34 directly yields (A.8). ■

A.6 Proof of Theorem 16

Theorem 16 *Suppose \mathcal{C} is finite. For any $c \in \mathcal{C}$, $\rho(c) > 0$, and ℓ satisfies Assumption 14. Then, with a probability of at least $1 - \delta$ over the choice of the training set \mathcal{S} , $\forall f \in \mathcal{F}$:*

$$\tilde{\mathcal{L}}_{sup}(\hat{f}) \leq \alpha(\rho) \left(\tilde{L}_{sun}(f) + AG_M \right) - \beta,$$

where $\alpha(\rho) = \frac{1}{1 - \tau_k} \max_{|\mathcal{T}| \leq k+1} \frac{p_{\max}(\mathcal{T})}{\rho_{\min}^+(\mathcal{T})}$ is a positive constant depending on ρ and \mathcal{T} distinct

$$\beta = \alpha(\rho) \tau_k \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \left[\ell_{|I^+|}(\vec{0}) | I^+ \neq \emptyset \right].$$

Proof Since \mathcal{C} is finite and $\rho(c) > 0$ for any $c \in \mathcal{C}$, we can see that:

$$\frac{1}{\alpha(1 - \tau_k)} = \min_{|\mathcal{T}| \leq k+1, \mathcal{T} \text{ distinct}} \left\{ \frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \right\} > 0. \quad (\text{A.15})$$

By Theorem 32 and the definition of $\tilde{\mathcal{L}}_{sup}(f)$, we have: $\forall f \in \mathcal{F}$, with probability at least $1 - \delta$ over the choice of the training set \mathcal{S} ,

$$\begin{aligned} \frac{1}{\alpha(1 - \tau_k)} \tilde{\mathcal{L}}_{sup}(\hat{f}) &\stackrel{(i)}{=} \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{1}{\alpha(1 - \tau_k)} \tilde{L}_{sup}(\mathcal{T}, \hat{f}) \right] \stackrel{(ii)}{\leq} \mathbb{E}_{\mathcal{T} \sim \mathcal{D}} \left[\frac{\rho_{\min}^+(\mathcal{T})}{p_{\max}(\mathcal{T})} \tilde{L}_{sup}(\mathcal{T}, \hat{f}) \right] \\ &\stackrel{(iii)}{\leq} \frac{1}{1 - \tau_k} \left(\tilde{L}_{sun}(f) + AG_M \right) - \frac{\tau_k}{1 - \tau_k} \mathbb{E}_{c^+, c_i^- \sim \rho^{k+1}} \left[\ell_{|I^+|}(\vec{0}) | I^+ \neq \emptyset \right], \end{aligned}$$

where (i) is from the definition of $\tilde{\mathcal{L}}_{sup}(\hat{f})$; (ii) is from (A.15) and (iii) is from Theorem 32. Rearranging yields the result. \blacksquare

A.7 Proof of Lemma 18

Lemma 18 *Let $\mathbb{B}_p(r)$ be the p -norm ball in \mathbb{R}^d with radius r . The δ -covering number of $\mathbb{B}_p(r)$ with respect to $\|\cdot\|_p$ thus obeys the following bound:*

$$\mathcal{N}(\delta; \mathbb{B}_p(r), \|\cdot\|_p) \leq \left(1 + \frac{2r}{\delta} \right)^d,$$

where $\mathcal{N}(\delta; \mathbb{B}, \|\cdot\|)$ is the δ -covering number of \mathbb{B} with respect to the norm $\|\cdot\|$.

Proof Set $\|\cdot\|, \|\cdot\|'$ in Lemma 17 to be $\|\cdot\|_p$, we have:

$$\begin{aligned} \mathcal{N}(\delta; \mathbb{B}_p(1), \|\cdot\|_p) &\leq \frac{\text{vol} \left(\frac{2}{\delta} \mathbb{B}_p(1) + \mathbb{B}_p(1) \right)}{\text{vol}(\mathbb{B}_p(1))} = \frac{\text{vol} \left(\left(1 + \frac{2}{\delta} \right) \mathbb{B}_p(1) \right)}{\text{vol}(\mathbb{B}_p(1))} \\ &\stackrel{(i)}{\leq} \left(1 + \frac{2}{\delta} \right)^d \frac{\text{vol}(\mathbb{B}_p(1))}{\text{vol}(\mathbb{B}_p(1))} = \left(1 + \frac{2}{\delta} \right)^d, \end{aligned}$$

where (i) is true because $\mathbb{B}_p(1) \subset \mathbb{R}^d$. Now suppose that $\{x^1, \dots, x^N\}$ is the minimal δ -covering of $\mathbb{B}_p(1)$, then:

$$\forall x \in \mathbb{B}_p(1), \exists x^i \in \{x^1, \dots, x^N\} \text{ s.t. } \|x - x^i\|_p \leq \delta.$$

So we have:

$$\forall rx \in \mathbb{B}_p(r), \exists rx^i \in \{rx^1, \dots, rx^N\} \text{ s.t. } \|rx - rx^i\|_p \leq r\delta.$$

So $\{rx^1, \dots, rx^N\}$ is a $r\delta$ -covering of $\mathbb{B}_p(r)$, so we have:

$$\mathcal{N}(r\delta; \mathbb{B}_p(r), \|\cdot\|_p) \leq \mathcal{N}(\delta; \mathbb{B}_p(1), \|\cdot\|_p) \leq \left(1 + \frac{2}{\delta} \right)^d.$$

Letting $\frac{\delta}{r}$ take place of δ , we have $\mathcal{N}(\delta; \mathbb{B}_p(r), \|\cdot\|_p) \leq \left(1 + \frac{2r}{\delta} \right)^d$. \blacksquare

A.8 Proof of Theorem 22

Theorem 22 Let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_r \leq \epsilon\}$ (i.e. consider the ℓ_r attack). We then have :

$$\mathcal{R}_S(\mathcal{H}) = O\left([PP^* + \epsilon R^* s(r^*, p, m)] \left[ms(p^*, p, m)w^2\sqrt{M}\right]\right),$$

where $s(p, q, n) := n^{\max\{\frac{1}{p} - \frac{1}{q}, \frac{1}{q} - \frac{1}{p}\}}$, $\frac{1}{p} + \frac{1}{p^*} = 1$, $\frac{1}{r} + \frac{1}{r^*} = 1$, and R^* is defined similarly to (11).

Before giving the proof of the theorem, we introduce some lemmas will be used in our proof.

Lemma 35 For any $x \in \mathbb{R}^n$, $0 < p_2 \leq p_1$, we have: $\|x\|_{p_1} \leq \|x\|_{p_2} \leq n^{\frac{1}{p_2} - \frac{1}{p_1}} \|x\|_{p_1}$.

Proof [Proof of Lemma 35] Firstly, we prove $\|x\|_{p_1} \leq \|x\|_{p_2}$. For any $x \in \mathbb{R}^n$, suppose $a_i = |x_i|$, let $f(p) = (\sum_{i=1}^n a_i^p)^{\frac{1}{p}}$. In order to prove $\|x\|_{p_1} \leq \|x\|_{p_2}$, it suffices to prove that: $\forall a_i \geq 0, i = 1, 2, \dots, n, f(p)$ is non-increasing on $(0, 1]$.

1. if $a_i = 0, \forall i$, f is a constant function, so f is non-increasing.
2. if $\exists i, a_i \neq 0$, suppose $\{a_i \mid a_i \neq 0\}$ has K elements, without loss of generality, suppose $\{a_i \mid a_i \neq 0\} = \{a_1, a_2, \dots, a_K\}$ and $a_1 = \max_{1 \leq k \leq K} \{a_k\}$, then we have:

$$0 < \left(\frac{a_k}{a_1}\right)^p \leq 1, \sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p \geq 1, \ln \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right) \geq 0, k = 1, 2, \dots, K. \quad (\text{A.16})$$

We can write $f(p) = (\sum_{i=1}^n a_i^p)^{\frac{1}{p}} = a_1 \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right)^{\frac{1}{p}} = a_1 \exp\left\{\frac{1}{p} \ln \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right)\right\}$, let $g(p) = \frac{1}{p} \ln \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right)$, by the monotone property of composite functions, to prove $f(p)$ is non-increasing, it suffices to prove $g(p)$ is non-increasing. Taking derivation of g yields:

$$g'(p) = \frac{\sum_{k=1}^K \left[\left(\frac{a_k}{a_1}\right)^p \ln \left(\frac{a_k}{a_1}\right)\right]}{p \sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p} - \frac{\ln \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right)}{p^2}.$$

From (A.16) we know that $\frac{\sum_{k=1}^K \left[\left(\frac{a_k}{a_1}\right)^p \ln \left(\frac{a_k}{a_1}\right)\right]}{p \sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p} \leq 0$ and $\frac{\ln \left(\sum_{k=1}^K \left(\frac{a_k}{a_1}\right)^p\right)}{p^2} \geq 0$, so we have $g'(p) \leq 0$, so g is non-increasing, which means that $\|x\|_{p_1} \leq \|x\|_{p_2}, \forall x \in \mathbb{R}^n$.

Nextly, we prove that $\|x\|_{p_2} \leq n^{\frac{1}{p_2} - \frac{1}{p_1}} \|x\|_{p_1}$. By the definition of $\|\cdot\|_p$, we have: $\|x\|_{p_2}^{p_1} = (\sum_{i=1}^n |x_i|^{p_2})^{\frac{p_1}{p_2}} = \left(\frac{1}{n} \sum_{i=1}^n |x_i|^{p_2}\right)^{\frac{p_1}{p_2}} \cdot n^{\frac{p_1}{p_2}}$. Since $p_2 \leq p_1$, i.e. $\frac{p_1}{p_2} \geq 1$, we know the function $h(t) = t^{\frac{p_1}{p_2}}$ is convex. By Jensen's Inequality we know that:

$$\forall \lambda_i \text{ s.t. } \lambda_i \geq 0 \text{ and } \sum_{i=1}^n \lambda_i = 1, \forall t_i \in \mathbb{R} : h\left(\sum_{i=1}^n \lambda_i t_i\right) \leq \sum_{i=1}^n \lambda_i h(t_i). \quad (\text{A.17})$$

Set $\lambda_1 = \lambda_2 = \dots = \lambda_n = \frac{1}{n}$, $t_i = |x_i|^{p_2}$, $i = 1, 2, \dots, n$ in (A.17), we have:

$$\left(\frac{1}{n} \sum_{i=1}^n |x_i|^{p_2} \right)^{\frac{p_1}{p_2}} \leq \frac{1}{n} \sum_{i=1}^n (|x_i|^{p_2})^{\frac{p_1}{p_2}} = \frac{1}{n} \sum_{i=1}^n |x_i|^{p_1}.$$

Multiplying both sides by $n^{\frac{p_1}{p_2}}$ yields $(\frac{1}{n} \sum_{i=1}^n |x_i|^{p_2})^{\frac{p_1}{p_2}} \cdot n^{\frac{p_1}{p_2}} \leq \frac{1}{n} \sum_{i=1}^n |x_i|^{p_1} \cdot n^{\frac{p_1}{p_2}}$, i.e.

$$\|x\|_{p_2}^{p_1} \leq n^{\frac{p_1}{p_2}-1} \cdot \|x\|_{p_1}^{p_1}.$$

Taking both sides the p_1 th power yields $\|x\|_{p_2} \leq n^{\frac{1}{p_2}-\frac{1}{p_1}} \|x\|_{p_1}$. ■

Lemma 36 *Suppose that $\exists A, B \in \mathbb{R}^+$ s.t. $A\|x\|_p \leq \|x\|_q \leq B\|x\|_p$ ($p, q \geq 1$) for any $x \in \mathbb{R}^n$, then $\forall W \in \mathbb{R}^{m \times n}$: $\|W\|_p \leq \frac{B}{A} \|W\|_q$.*

Proof [Proof of Lemma 36] If we have $A\|x\|_p \leq \|x\|_q \leq B\|x\|_p$, then:

$$\|W\|_p = \max_{\|x\|_p \leq 1} \|Wx\|_p \stackrel{(i)}{\leq} \max_{\|x\|_p \leq B} \|Wx\|_p \stackrel{(ii)}{\leq} \max_{\|x\|_p \leq B} \frac{1}{A} \|Wx\|_q = \max_{\|x\|_p \leq 1} \frac{B}{A} \|Wx\|_q = \frac{B}{A} \|W\|_q,$$

where (i) is from the fact that $\|x\|_q \leq B\|x\|_p$, which means that $\{x \mid \|x\|_p \leq 1\} \subseteq \{x \mid \|x\|_q \leq 1\}$ and (ii) is from the condition that $A\|x\|_p \leq \|x\|_q$. ■

Lemma 37 *For any $p, q \geq 1$, then $\forall W \in \mathbb{R}^{m \times n}$: $\|W\|_q \leq n^{\max\{\frac{1}{p}-\frac{1}{q}, \frac{1}{q}-\frac{1}{p}\}} \|W\|_p$.*

Remark 38 *If we denote $n^{\max\{\frac{1}{p}-\frac{1}{q}, \frac{1}{q}-\frac{1}{p}\}}$ by $s(p, q, n)$, then we have: $\forall W \in \mathbb{R}^{m \times n}$, $\|W\|_q \leq s(p, q, n) \|W\|_p$.*

Proof [Proof of Lemma 37] We discuss it in two cases.

1. If $p \leq q$, by Lemma 35, we have: $\forall x \in \mathbb{R}^n$, $\|x\|_q \leq \|x\|_p \leq n^{\frac{1}{p}-\frac{1}{q}} \|x\|_q$. By Lemma 36, we know that $\forall W \in \mathbb{R}^{m \times n}$, $\|W\|_q \leq n^{\frac{1}{p}-\frac{1}{q}} \|W\|_p$.
2. If $p \geq q$, by Lemma 35, we have: $\forall x \in \mathbb{R}^n$, $\|x\|_p \leq \|x\|_q \leq n^{\frac{1}{q}-\frac{1}{p}} \|x\|_p$, i.e.

$$n^{\frac{1}{p}-\frac{1}{q}} \|x\|_q \leq \|x\|_p \leq \|x\|_q.$$

By Lemma 36, we know that $\forall W \in \mathbb{R}^{m \times n}$, $\|W\|_q \leq n^{\frac{1}{q}-\frac{1}{p}} \|W\|_p$.

So we conclude that $\forall W \in \mathbb{R}^{m \times n}$, $\|W\|_q \leq n^{\max\{\frac{1}{p}-\frac{1}{q}, \frac{1}{q}-\frac{1}{p}\}} \|W\|_p$. ■

Now we turn to bound the Rademacher complexity of \mathcal{H} , where

$$\mathcal{H} = \left\{ h(x, x^+, x^-) = \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F} \right\},$$

$$\mathcal{F} = \{f : x \rightarrow Wx \mid W \in \mathbb{R}^{n \times m}, \|W\|_p \leq w\}.$$

Let

$$\begin{aligned}\mathcal{H}_0 &= \{h(x, x^+, x^-) = f(x)^T (f(x^+) - f(x^-)) \mid f \in \mathcal{F}\} \\ &= \left\{ (x, x^+, x^-) \rightarrow x^T W^T W (x^+ - x^-) \mid W \in \mathbb{R}^{n \times m}, \|W\|_p \leq w \right\}.\end{aligned}$$

Proof [Proof of Theorem 22] Firstly, we drop the complicated and unwieldy min operate by directly solving the minimization problem.

$$\begin{aligned}h(x, x^+, x^-) &= \min_{\|x' - x\|_r \leq \epsilon} f(x')^T (f(x^+) - f(x^-)) = \min_{\|\delta\|_r \leq \epsilon} (W(x + \delta))^T (Wx^+ - Wx^-) \\ &= x^T W^T W (x^+ - x^-) + \min_{\|\delta\|_r \leq \epsilon} \delta^T W^T W (x^+ - x^-) \\ &\stackrel{(i)}{=} x^T W^T W (x^+ - x^-) - \epsilon \|W^T W (x^+ - x^-)\|_{r^*},\end{aligned}$$

where (i) is from Holder's Inequality and taking the value of δ that can get equality. Define:

$$\begin{aligned}\mathcal{H}_1 &= \left\{ (x, x^+, x^-) \rightarrow \|W^T W (x^+ - x^-)\|_{r^*} \mid W \in \mathbb{R}^{n \times m}, \|W\|_p \leq w \right\}, \\ \mathcal{H}_0^w &= \left\{ (x, x^+, x^-) \rightarrow x^T A (x^+ - x^-) \mid A \in \mathbb{R}^{m \times m}, \|A\|_p \leq w \right\}, \\ \mathcal{H}_1^w &= \left\{ (x, x^+, x^-) \rightarrow \|A (x^+ - x^-)\|_{r^*} \mid A \in \mathbb{R}^{m \times m}, \|A\|_p \leq w \right\}.\end{aligned}$$

Since $\forall W \in \mathbb{R}^{n \times m}$ s.t. $\|W\|_p \leq w$, we have:

$$\| \|W^T W\|_p \| \|W\|_p \leq \| \|W^T\|_p \cdot \|W\|_p = \| \|W\|_{p^*} \cdot \|W\|_p \leq s(p^*, p, m) \| \|W\|_p^2 \leq s(p^*, p, m) w^2,$$

where $\frac{1}{p} + \frac{1}{p^*} = 1$ and (i) comes from the Submultiplicativity of matrix norm and (ii) is from Lemma 37. So we have:

$$\left\{ W \mid \| \|W\|_p \leq w \right\} \subseteq \left\{ W \mid \| \|W^T W\|_p \leq s(p^*, p, m) w^2 \right\}.$$

And since $W^T W$ is positive semi-definite, it's easy to see that:

$$\left\{ W^T W \mid \| \|W^T W\|_p \leq s(p^*, p, m) w^2 \right\} \subseteq \left\{ A \mid \| \|A\|_p \leq s(p^*, p, m) w^2 \right\}.$$

So we have:

$$\begin{aligned}\mathcal{H}_0 &= \left\{ (x, x^+, x^-) \rightarrow x^T W^T W (x^+ - x^-) \mid W \in \mathbb{R}^{n \times m}, \|W\|_p \leq w \right\} \\ &\subseteq \left\{ (x, x^+, x^-) \rightarrow x^T W^T W (x^+ - x^-) \mid W \in \mathbb{R}^{n \times m}, \| \|W^T W\|_p \leq s(p^*, p, m) w^2 \right\} \\ &\subseteq \left\{ (x, x^+, x^-) \rightarrow x^T A (x^+ - x^-) \mid A \in \mathbb{R}^{m \times m}, \| \|A\|_p \leq s(p^*, p, m) w^2 \right\} = \mathcal{H}_0^{s(p^*, p, m) w^2}.\end{aligned}$$

Similarly, $\mathcal{H}_1 \subseteq \mathcal{H}_1^{s(p^*, p, m)w^2}$. Given the training set $\mathcal{S} = \{(x_i, x_i^+, x_i^-)\}_{i=1}^M$, we have:

$$\begin{aligned}
 \mathcal{R}_{\mathcal{S}}(\mathcal{H}) &= \mathbb{E}_{\boldsymbol{\sigma}} \left\{ \sup_{\|W\|_p \leq w} \sum_{i=1}^M \boldsymbol{\sigma}_i [x_i^T W^T W (x_i^+ - x_i^-) - \epsilon \|W^T W (x_i^+ - x_i^-)\|_{r^*}] \right\} \\
 &\stackrel{(i)}{\leq} \mathbb{E}_{\boldsymbol{\sigma}} \left\{ \sup_{\|W\|_p \leq w} \sum_{i=1}^M \boldsymbol{\sigma}_i [x_i^T W^T W (x_i^+ - x_i^-)] \right\} \\
 &\quad + \epsilon \mathbb{E}_{\boldsymbol{\sigma}} \left\{ \sup_{\|W\|_p \leq w} \sum_{i=1}^M [-\boldsymbol{\sigma}_i \|W^T W (x_i^+ - x_i^-)\|_{r^*}] \right\} \\
 &\stackrel{(ii)}{=} \mathbb{E}_{\boldsymbol{\sigma}} \left\{ \sup_{\|W\|_p \leq w} \sum_{i=1}^M \boldsymbol{\sigma}_i [x_i^T W^T W (x_i^+ - x_i^-)] \right\} \\
 &\quad + \epsilon \mathbb{E}_{\boldsymbol{\sigma}} \left\{ \sup_{\|W\|_p \leq w} \sum_{i=1}^M [\boldsymbol{\sigma}_i \|W^T W (x_i^+ - x_i^-)\|_{r^*}] \right\} \\
 &\stackrel{(iii)}{=} \mathcal{R}_{\mathcal{S}}(\mathcal{H}_0) + \epsilon \mathcal{R}_{\mathcal{S}}(\mathcal{H}_1) \stackrel{(iv)}{\leq} \mathcal{R}_{\mathcal{S}}(\mathcal{H}_0^{s(p^*, p, m)w^2}) + \epsilon \mathcal{R}_{\mathcal{S}}(\mathcal{H}_1^{s(p^*, p, m)w^2}),
 \end{aligned} \tag{A.18}$$

where $\boldsymbol{\sigma}$ is a Random Vector whose elements are i.i.d. Rademacher Random Variables and (i) comes from the subadditivity of sup function; (ii) is from the fact that $-\boldsymbol{\sigma}_i$ has the same distribution as $\boldsymbol{\sigma}_i$, $\forall i = 1, 2, \dots, M$; (iii) is by the definition of $\mathcal{R}_{\mathcal{S}}(\mathcal{H}_0)$ and $\mathcal{R}_{\mathcal{S}}(\mathcal{H}_1)$ and (iv) is from the monotone property of Rademacher Complexity and the fact that $\mathcal{H}_0 \subseteq \mathcal{H}_0^{s(p^*, p, m)w^2}$ and $\mathcal{H}_1 \subseteq \mathcal{H}_1^{s(p^*, p, m)w^2}$. Secondly, we upper bound the Rademacher Complexity of \mathcal{H}_0^w and \mathcal{H}_1^w . For $\mathcal{R}_{\mathcal{S}}(\mathcal{H}_0^w)$: Given the training set $\mathcal{S} = \{(x_i, x_i^+, x_i^-)\}_{i=1}^M = \{(z_i)\}_{i=1}^M$, define the ℓ_2 -norm for a function in \mathcal{H}_0^w as:

$$\forall h \in \mathcal{H}_0^w, \|h\|_2 := \sqrt{\sum_{i=1}^M [h(z_i)]^2},$$

Define $\mathcal{H}_0^w(\mathcal{S}) = \{(h(z_1), h(z_2), \dots, h(z_M)) \mid h \in \mathcal{H}_0^w\}$, we have that for any $h \in \mathcal{H}_0^w$ and $v_h = (h(z_1), h(z_2), \dots, h(z_M))$ be the corresponding vector in $\mathcal{H}_0^w(\mathcal{S})$, we have:

$$\|h\|_2 = \|v_h\|_2.$$

So we know that any δ -covering of \mathcal{H}_0^w ($\{h^1, h^2, \dots, h^N\}$) with respect to ℓ_2 norm in the functional space, corresponds to a δ -covering of $\mathcal{H}_0^w(\mathcal{S})$ with respect to the ℓ_2 norm in the Euclidean Space, i.e.

$$\left\{ \left(\begin{array}{c} h^1(z_1) \\ h^1(z_2) \\ \vdots \\ h^1(z_M) \end{array} \right), \left(\begin{array}{c} h^2(z_1) \\ h^2(z_2) \\ \vdots \\ h^2(z_M) \end{array} \right), \dots, \left(\begin{array}{c} h^N(z_1) \\ h^N(z_2) \\ \vdots \\ h^N(z_M) \end{array} \right) \right\}.$$

So we have:

$$\mathcal{N}(\delta; \mathcal{H}_0^w(\mathcal{S}), \|\cdot\|_2) = \mathcal{N}(\delta; \mathcal{H}_0^w, \|\cdot\|_2). \quad (\text{A.19})$$

By the definition of Rademacher Complexity, we know that $\mathcal{R}_{\mathcal{S}}(\mathcal{H}_0^w)$ is just the expectation of the Rademacher Process with respect to $\mathcal{H}_0^w(\mathcal{S})$, which is $\mathbb{E}[\sup_{\theta \in \mathcal{H}_0^w(\mathcal{S})} X_{\theta}]$. To use Lemma 20, we must show that Rademacher Process is a sub-Gaussian Process with respect to some metric ρ_X . Denote the Euclidean metric by ρ_2 , we have: for Rademacher Process $\{X_{\theta}, \theta \in \mathbb{T}\}$, $\forall \theta, \tilde{\theta} \in \mathbb{T}$ and $\lambda \in \mathbb{R}$:

$$\begin{aligned} \mathbb{E} \left[e^{\lambda(X_{\theta} - X_{\tilde{\theta}})} \right] &\stackrel{(i)}{=} \mathbb{E} \left[e^{\lambda(\langle \sigma, \theta \rangle - \langle \sigma, \tilde{\theta} \rangle)} \right] = \mathbb{E} \left[e^{\lambda(\langle \sigma, \theta - \tilde{\theta} \rangle)} \right] \stackrel{(ii)}{\leq} \prod_{i=1}^M \mathbb{E} \left[e^{\lambda \sigma_i (\theta_i - \tilde{\theta}_i)} \right] \\ &\stackrel{(iii)}{\leq} \prod_{i=1}^M e^{\frac{\lambda^2 (\theta_i - \tilde{\theta}_i)^2}{2}} = e^{\frac{\lambda^2}{2} \sum_{i=1}^M (\theta_i - \tilde{\theta}_i)^2} = e^{\frac{\lambda^2}{2} \rho_X^2(\theta, \tilde{\theta})}, \end{aligned}$$

where σ is a Random Vector whose elements are i.i.d. Rademacher Random Variables and (i) is from the definition of Rademacher Process; (ii) is from the expectation property of i.i.d. random variables and (iii) is from Example 2.3 in Wainwright (2019). So we proved that the Rademacher Process is a sub-Gaussian Process with respect to the Euclidean metric ρ_2 . So by Lemma 20 and (9), we know that: $\forall \delta \in (0, D]$,

$$\mathcal{R}_{\mathcal{S}}(\mathcal{H}_0^w) = \mathbb{E} \left[\sup_{\theta \in \mathcal{H}_0^w(\mathcal{S})} X_{\theta} \right] \leq \mathbb{E} \left[\sup_{\theta, \tilde{\theta} \in \mathcal{H}_0^w(\mathcal{S})} (X_{\theta} - X_{\tilde{\theta}}) \right] \leq 2 \mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}_0^w(\mathcal{S}) \\ \rho_X(\gamma, \gamma') \leq \delta}} (X_{\gamma} - X_{\gamma'}) \right] + 32 \mathcal{J}(\delta/4; D),$$

where

$$\begin{aligned} D &= \sup_{\theta, \theta' \in \mathcal{H}_0^w(\mathcal{S})} \|\theta - \theta'\|_2 \leq 2 \sup_{\theta \in \mathcal{H}_0^w(\mathcal{S})} \|\theta\|_2 = 2 \sup_{h \in \mathcal{H}_0^w} \|h\|_2 = 2 \sup_{h \in \mathcal{H}_0^w} \sqrt{\sum_{i=1}^M [h(z_i)]^2} \\ &\stackrel{(i)}{\leq} 2\sqrt{M} \sup_{\substack{\|A\|_p \leq w \\ 1 \leq i \leq M}} |x_i^T A(x_i^+ - x_i^-)| \stackrel{(ii)}{\leq} 2\sqrt{M} \sup_{\substack{\|A\|_p \leq w \\ 1 \leq i \leq M}} \|x_i\|_{p^*} \|A(x_i^+ - x_i^-)\|_p \\ &\stackrel{(iii)}{\leq} 2\sqrt{M} \sup_{\substack{\|A\|_p \leq w \\ 1 \leq i \leq M}} \|x_i\|_{p^*} \|A\|_p \|x_i^+ - x_i^-\|_p \stackrel{(iv)}{\leq} 4\sqrt{M} P^* P w, \end{aligned} \quad (\text{A.20})$$

and $\mathcal{J}(a; b) = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}_0^w(\mathcal{S}), \|\cdot\|_2)} du = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}_0^w, \|\cdot\|_2)} du$. Where (i) is from the definition of f ; (ii) is from the Holder's Inequality; (iii) is a result of properties of matrix norm and (iv) is from the definition of P^* and P . Similar to the discussion of upper bound

for D , for all $h_1, h_2 \in \mathcal{H}_0^w$, we have:

$$\begin{aligned}
 \|h_1 - h_2\|_2 &= \sqrt{\sum_{i=1}^M [h_1(z_i) - h_2(z_i)]^2} \leq \sqrt{M} \sup_{1 \leq i \leq M} |h_1(z_i) - h_2(z_i)| \\
 &= \sqrt{M} \sup_{1 \leq i \leq M} |x_i^T (A_1 - A_2)(x_i^+ - x_i^-)| \\
 &\stackrel{(i)}{\leq} \sqrt{M} \sup_{1 \leq i \leq M} \|x_i\|_{p^*} \|A_1 - A_2\|_p \|x_i^+ - x_i^-\|_p \stackrel{(ii)}{\leq} 2\sqrt{M} P^* P \|A_1 - A_2\|_p,
 \end{aligned} \tag{A.21}$$

where A_1, A_2 are the matrices corresponding to h_1, h_2 , respectively and (i) is from the same argument as (ii), (iii) in (A.20) and (ii) is from the definition of P^* and P . Suppose $N_A = \{A^1, \dots, A^N\}$ is a $\frac{\delta}{2PP^*\sqrt{M}}$ -covering of $S_A = \{A \mid \|A\|_p \leq w\}$ with respect to $\|\cdot\|_p$, i.e.:

$$\forall A \in S_A, \exists A^j \in N_A, \text{ s.t. } \|A - A^j\|_p \leq \frac{\delta}{2PP^*\sqrt{M}}.$$

Combine this with (A.21), let A be the matrix corresponding to h and A^j be the matrix corresponding to h^j and let $N_h = \{h^1, h^2, \dots, h^N\}$, we have:

$$\forall h \in \mathcal{H}_0^w, \exists h^j \in N_h, \text{ s.t. } \|h - h^j\|_2 \leq 2\sqrt{M} P^* P \frac{\delta}{2PP^*\sqrt{M}} = \delta.$$

So N_h is a δ -covering of \mathcal{H}_0^w with respect to $\|\cdot\|_2$, so we have:

$$\mathcal{N}(\delta; \mathcal{H}_0^w, \|\cdot\|_2) \leq \mathcal{N}\left(\frac{\delta}{2PP^*\sqrt{M}}; S_A, \|\cdot\|_p\right).$$

By Lemma 18, we know that:

$$\mathcal{N}(\delta; \mathcal{H}_0^w, \|\cdot\|_2) \leq \mathcal{N}\left(\frac{\delta}{2PP^*\sqrt{M}}; S_A, \|\cdot\|_p\right) \leq \left(1 + \frac{4PP^*w\sqrt{M}}{\delta}\right)^{m^2}. \tag{A.22}$$

So we have:

$$\begin{aligned}
 \mathcal{J}(0; D) &= \int_0^D \sqrt{\ln \mathcal{N}(u; \mathcal{H}_0^w, \|\cdot\|_2)} du \stackrel{(i)}{\leq} \int_0^D \sqrt{m^2 \ln \left(1 + \frac{4PP^*w\sqrt{M}}{u}\right)} du \\
 &\stackrel{(ii)}{\leq} m \int_0^D \sqrt{\frac{4PP^*w\sqrt{M}}{u}} du = 2m\sqrt{PP^*w} \sqrt[4]{M} \int_0^D u^{-\frac{1}{2}} du \\
 &= 4m\sqrt{PP^*wD} \sqrt[4]{M} \stackrel{(iii)}{\leq} 8mPP^*w\sqrt{M},
 \end{aligned}$$

where (i) is from (A.22); (ii) is because that for any $x \geq 0$, $\ln(1+x) \leq x$ and (iii) is from (A.20). So take $\delta \rightarrow 0^+$, we have:

$$\mathcal{R}_S(\mathcal{H}_0^w) \leq 32\mathcal{J}(0; D) \leq 256mPP^*w\sqrt{M}.$$

For $\mathcal{R}_S(\mathcal{H}_1^w)$: The same as $\mathcal{R}_S(\mathcal{H}_0^w)$, we consider ℓ_2 norm for a function $h \in \mathcal{H}_1^w$ and define $\mathcal{H}_1^w(\mathcal{S}) = \left\{ (h(z_1), h(z_2), \dots, h(z_M)) \mid h \in \mathcal{H}_1^w \right\}$, by similar argument in (A.19), we have:

$$\mathcal{N}(\delta; \mathcal{H}_1^w(\mathcal{S}), \|\cdot\|_2) = \mathcal{N}(\delta; \mathcal{H}_1^w, \|\cdot\|_2).$$

By the definition of Rademacher Complexity, we know that $\mathcal{R}_S(\mathcal{H}_1^w)$ is just the expectation of the Rademacher Process with respect to $\mathcal{H}_1^w(\mathcal{S})$, which is $\mathbb{E}[\sup_{\theta \in \mathcal{H}_1^w(\mathcal{S})} X_\theta]$. By Lemma 20 and (9), we know that: $\forall \delta \in (0, D]$,

$$\mathcal{R}_S(\mathcal{H}_1^w) = \mathbb{E}[\sup_{\theta \in \mathcal{H}_1^w(\mathcal{S})} X_\theta] \leq \mathbb{E}\left[\sup_{\theta, \tilde{\theta} \in \mathcal{H}_1^w(\mathcal{S})} (X_\theta - X_{\tilde{\theta}})\right] \leq 2\mathbb{E}\left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}_1^w(\mathcal{S}) \\ \rho_X(\gamma, \gamma') \leq \delta}} (X_\gamma - X_{\gamma'})\right] + 32\mathcal{J}(\delta/4; D),$$

where

$$\begin{aligned} D &= \sup_{\theta, \theta' \in \mathcal{H}_1^w(\mathcal{S})} \|\theta - \theta'\|_2 \leq 2 \sup_{\theta \in \mathcal{H}_1^w(\mathcal{S})} \|\theta\|_2 = 2 \sup_{h \in \mathcal{H}_1^w} \|h\|_2 = 2 \sup_{h \in \mathcal{H}_1^w} \sqrt{\sum_{i=1}^M [h(z_i)]^2} \\ &\stackrel{(i)}{\leq} 2\sqrt{M} \sup_{\substack{\|A\|_p \leq w \\ 1 \leq i \leq M}} \|A(x_i^+ - x_i^-)\|_{r^*} \stackrel{(ii)}{\leq} 2\sqrt{M} \sup_{\substack{\|A\|_p \leq w \\ 1 \leq i \leq M}} \|A\|_{r^*} \|x_i^+ - x_i^-\|_{r^*} \\ &\stackrel{(iii)}{\leq} 4\sqrt{M}R^* \sup_{\|A\|_p \leq w} \|A\|_{r^*} \stackrel{(iv)}{\leq} 4\sqrt{M}R^* \sup_{\|A\|_p \leq w} s(r^*, p, m) \|A\|_p \leq 4wR^* s(r^*, p, m) \sqrt{M}, \end{aligned} \tag{A.23}$$

and $\mathcal{J}(a; b) = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}_1^w(\mathcal{S}), \|\cdot\|_2)} du = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}_1^w, \|\cdot\|_2)} du$. Where (i) is from the definition of f ; (ii) is a result of the properties of matrix norm; (iii) is from the definition of R^* and (iv) comes from Lemma 37. Similar to the discussion of upper bound for D , for all $h_1, h_2 \in \mathcal{H}_1^w$, we have:

$$\begin{aligned} \|h_1 - h_2\|_2 &= \sqrt{\sum_{i=1}^M [h_1(z_i) - h_2(z_i)]^2} \leq \sqrt{M} \sup_{1 \leq i \leq M} |h_1(z_i) - h_2(z_i)| \\ &= \sqrt{M} \sup_{1 \leq i \leq M} \|(A_1 - A_2)(x_i^+ - x_i^-)\|_{r^*} \stackrel{(i)}{\leq} \sqrt{M} \sup_{1 \leq i \leq M} \|A_1 - A_2\|_{r^*} \|x_i^+ - x_i^-\|_{r^*} \\ &\stackrel{(ii)}{\leq} 2\sqrt{M}R^* \|A_1 - A_2\|_{r^*} \stackrel{(iii)}{\leq} 2R^* s(r^*, p, m) \sqrt{M} \|A_1 - A_2\|_p, \end{aligned} \tag{A.24}$$

where A_1, A_2 are the matrices corresponding to h_1, h_2 , respectively and (i) from the properties of matrix norm and (ii) is from the definition of R^* and (iii) comes from Lemma 37. Suppose $N_A = \{A^1, \dots, A^N\}$ is a $\frac{\delta}{2R^* s(r^*, p, m) \sqrt{M}}$ -covering of $S_A = \{A \mid \|A\|_p \leq w\}$ with respect to $\|\cdot\|_p$, i.e.:

$$\forall A \in S_A, \exists A^j \in N_A, \text{ s.t. } \|A - A^j\|_p \leq \frac{\delta}{2R^* s(r^*, p, m) \sqrt{M}}.$$

Combine this with (A.24), let A be the matrix corresponding to h and A^j be the matrix corresponding to h^j and let $N_h = \{h^1, h^2, \dots, h^N\}$, we have:

$$\forall h \in \mathcal{H}_1^w, \exists h^j \in N_h, s.t. \|h - h^j\|_2 \leq 2R^*s(r^*, p, m)\sqrt{M} \frac{\delta}{2R^*s(r^*, p, m)\sqrt{M}} = \delta.$$

So N_h is a δ -covering of \mathcal{H}_1^w with respect to $\|\cdot\|_2$, so we have:

$$\mathcal{N}(\delta; \mathcal{H}_1^w, \|\cdot\|_2) \leq \mathcal{N}\left(\frac{\delta}{2R^*s(r^*, p, m)\sqrt{M}}; S_A, \|\cdot\|_p\right).$$

By Lemma 18, we know that:

$$\mathcal{N}(\delta; \mathcal{H}_1^w, \|\cdot\|_2) \leq \mathcal{N}\left(\frac{\delta}{2R^*s(r^*, p, m)\sqrt{M}}; S_A, \|\cdot\|_p\right) \leq \left(1 + \frac{4R^*s(r^*, p, m)w\sqrt{M}}{\delta}\right)^{m^2}. \quad (\text{A.25})$$

So we have:

$$\begin{aligned} \mathcal{J}(0; D) &= \int_0^D \sqrt{\ln \mathcal{N}(u; \mathcal{H}_1^w, \|\cdot\|_2)} du \stackrel{(i)}{\leq} \int_0^D \sqrt{m^2 \ln \left(1 + \frac{4R^*s(r^*, p, m)w\sqrt{M}}{u}\right)} du \\ &\stackrel{(ii)}{\leq} m \int_0^D \sqrt{\frac{4R^*s(r^*, p, m)w\sqrt{M}}{u}} du = 2m\sqrt{R^*s(r^*, p, m)w^4M} \int_0^D u^{-\frac{1}{2}} du \\ &= 4m\sqrt{R^*s(r^*, p, m)wD^4M} \stackrel{(iii)}{\leq} 8mR^*s(r^*, p, m)w\sqrt{M}, \end{aligned}$$

where (i) is from (A.25); (ii) is because that for any $x \geq 0$, $\ln(1+x) \leq x$ and (iii) is from (A.23). So take $\delta \rightarrow 0^+$, we have:

$$\mathcal{R}_S(\mathcal{H}_1^w) \leq 32\mathcal{J}(0; D) \leq 256mR^*s(r^*, p, m)w\sqrt{M}.$$

Combine upper bounds of $\mathcal{R}_S(\mathcal{H}_0^w)$ and $\mathcal{R}_S(\mathcal{H}_1^w)$ with (A.18), we have:

$$\begin{aligned} \mathcal{R}_S(\mathcal{H}) &\leq \mathcal{R}_S(\mathcal{H}_0^{s(p^*, p, m)w^2}) + \epsilon \mathcal{R}_S(\mathcal{H}_1^{s(p^*, p, m)w^2}) \\ &\leq 256mPP^*s(p^*, p, m)w^2\sqrt{M} + \epsilon 256mR^*s(r^*, p, m)s(p^*, p, m)w^2\sqrt{M} \\ &= 256m s(p^*, p, m)w^2\sqrt{M} (PP^* + \epsilon R^*s(r^*, p, m)) \\ &= O\left((PP^* + \epsilon R^*s(r^*, p, m)) \left(m s(p^*, p, m)w^2\sqrt{M}\right)\right). \end{aligned}$$

■

A.9 Proof of Theorem 24

Theorem 24 Let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$ (i.e. consider the ℓ_p attack), $\sigma(0) = 0$ with Lipschitz constant L and let $\mathcal{F} = \{W_d \sigma(W_{d-1} \sigma(\dots \sigma(W_1 x))) \mid \|W_l\|_F \leq M_l^F, l = 1, \dots, d\}$. We then have:

$$\mathcal{R}_S(\mathcal{H}) = O \left(\sqrt{\sum_{l=1}^d h_l h_{l-1} K \sqrt{d} \sqrt{M}} \right),$$

where $K = 2B_{X,\epsilon}^F \cdot (B_{X^+}^F + B_{X^-}^F)$, where

$$B_{X,\epsilon}^F = L^{d-1} \prod_{l=1}^d M_l^F \max \left\{ 1, m^{\frac{1}{2} - \frac{1}{p}} \right\} (\|X\|_{p,\infty} + \epsilon), \quad B_X^F = L^{d-1} \prod_{l=1}^d M_l^F \max \left\{ 1, m^{\frac{1}{2} - \frac{1}{p}} \right\} \|X\|_{p,\infty}.$$

Before giving the proof, we firstly introduce some useful lemmas.

Lemma 39 If $x_i^* \in \mathcal{U}(x_i) = \{x'_i \mid \|x_i - x'_i\|_p \leq \epsilon\}$, then for $\frac{1}{r^*} + \frac{1}{r} = 1$, we have:

$$\|x_i^*\|_{r^*} \leq \max \left\{ 1, m^{1 - \frac{1}{r} - \frac{1}{p}} \right\} (\|X\|_{p,\infty} + \epsilon).$$

Proof [Proof of Lemma 39] We divide it into two cases.

1. If $p \geq r^*$, then by Holder's Inequality with $\frac{1}{r^*} = \frac{1}{p} + \frac{1}{s}$, we have:

$$\|x_i^*\|_{r^*} \leq \sup \|\mathbf{1}\|_s \cdot \|x_i^*\|_p = \|\mathbf{1}\|_s \cdot \|x_i^*\|_p = m^{\frac{1}{s}} \|x_i^*\|_p = m^{1 - \frac{1}{r} - \frac{1}{p}} \|x_i^*\|_p,$$

where the equality holds when all the entries are equal.

2. If $p < r^*$, by Lemma 35, we have $\|x_i^*\|_{r^*} \leq \|x_i^*\|_p$, where the equality holds when one of the entries of x_i^* equals to one and the others equal to zero.

Then we have:

$$\begin{aligned} \|x_i^*\|_{r^*} &\leq \max \left\{ 1, m^{1 - \frac{1}{r} - \frac{1}{p}} \right\} \|x_i^*\|_p \leq \max \left\{ 1, m^{1 - \frac{1}{r} - \frac{1}{p}} \right\} (\|x_i\|_p + \|x_i - x_i^*\|_p) \\ &\leq \max \left\{ 1, m^{1 - \frac{1}{r} - \frac{1}{p}} \right\} (\|X\|_{p,\infty} + \epsilon). \end{aligned}$$

■

Lemma 40 Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^n$, then we have: $\|A \cdot b\|_2 \leq \|A\|_F \cdot \|b\|_2$.

Proof [Proof of Lemma 40] Let A_i be the rows of A , $i = 1, 2, \dots, m$, we have:

$$\|A \cdot b\|_2 = \sqrt{\sum_{i=1}^m (A_i b)^2} \stackrel{(i)}{\leq} \sqrt{\sum_{i=1}^m \|A_i\|_2^2 \cdot \|b\|_2^2} = \sqrt{\sum_{i=1}^m \|A_i\|_2^2} \cdot \sqrt{\|b\|_2^2} = \|A\|_F \cdot \|b\|_2,$$

where (i) is from Holder's Inequality. ■

Lemma 41 *Suppose σ is a L -Lipschitz function, then the elementwise vector map corresponding to σ is also L -Lipschitz with respect to $\|\cdot\|_2$.*

Proof [Proof of Lemma 41]

$$\begin{aligned} \|\sigma(x) - \sigma(y)\|_2 &= \sqrt{\sum_{i=1}^n (\sigma(x)_i - \sigma(y)_i)^2} = \sqrt{\sum_{i=1}^n (\sigma(x_i) - \sigma(y_i))^2} \\ &\stackrel{(i)}{\leq} \sqrt{\sum_{i=1}^n L^2(x_i - y_i)^2} = L \sqrt{\sum_{i=1}^n (x_i - y_i)^2} = L \cdot \|x - y\|_2, \end{aligned}$$

where (i) is because σ is L -Lipschitz. ■

Now we can turn to the proof of Theorem 24.

Proof [Proof of Theorem 24] In this case, let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$, we have:

$$\begin{aligned} \mathcal{F} &= \{x \rightarrow W_d \sigma(W_{d-1} \sigma(\cdots \sigma(W_1 x))) \mid \|W_l\|_F \leq M_l^F, l = 1, \dots, d\}, \\ \mathcal{H} &= \left\{ h(x, x^+, x^-) = \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F} \right\}. \end{aligned}$$

Let $S_l^F = \{W_l \in \mathbb{R}^{h_l \times h_{l-1}} \mid \|W_l\|_F \leq M_l^F\}$, $l = 1, 2, \dots, d$. Let \mathcal{C}_l^F be the δ_l -covering of S_l^F and define:

$$\begin{aligned} \mathcal{F}^c &= \left\{ f^c : x \rightarrow W_d^c \sigma(W_{d-1}^c \sigma(\cdots \sigma(W_1^c x))) \mid W_l^c \in \mathcal{C}_l^F, l = 1, 2, \dots, d \right\} \subseteq \mathcal{F}, \\ \mathcal{H}^c &= \left\{ h^c(x, x^+, x^-) = \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F}^c \right\} \subseteq \mathcal{H}. \end{aligned}$$

Similar to the proof of Theorem 22, we know that the Rademacher Process is a sub-Gaussian Process with respect to the Euclidean metric, which induces the ℓ_2 norm.

Given the training set $\mathcal{S} = \{(x_i, x_i^+, x_i^-)\}_{i=1}^M \triangleq \{z_i\}_{i=1}^M$, define the ℓ_2 -norm for a function in \mathcal{H} as:

$$\forall h \in \mathcal{H}, \|h\|_2 := \sqrt{\sum_{i=1}^M [h(z_i)]^2}.$$

Define $\mathcal{H}(\mathcal{S}) = \left\{ (h(z_1), h(z_2), \dots, h(z_M)) \mid h \in \mathcal{H} \right\}$, we have that for any $h \in \mathcal{H}$ and $v_h = (h(z_1), h(z_2), \dots, h(z_M))$ be the corresponding vector in $\mathcal{H}(\mathcal{S})$, we have $\|h\|_2 = \|v_h\|_2$. So we know that any δ -covering of \mathcal{H} ($\{h^1, h^2, \dots, h^N\}$) with respect to ℓ_2 norm in the functional space, corresponds to a δ -covering of $\mathcal{H}(\mathcal{S})$ with respect to the ℓ_2 norm in the Euclidean Space, i.e.

$$\left\{ \left(\begin{array}{c} h^1(z_1) \\ h^1(z_2) \\ \vdots \\ h^1(z_M) \end{array} \right), \left(\begin{array}{c} h^2(z_1) \\ h^2(z_2) \\ \vdots \\ h^2(z_M) \end{array} \right), \dots, \left(\begin{array}{c} h^N(z_1) \\ h^N(z_2) \\ \vdots \\ h^N(z_M) \end{array} \right) \right\}.$$

So we have:

$$\mathcal{N}(\delta; \mathcal{H}(\mathcal{S}), \|\cdot\|_2) = \mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2).$$

By the definition of Rademacher Complexity, we know that $\mathcal{R}_{\mathcal{S}}(\mathcal{H})$ is just the expectation of the Rademacher Process with respect to $\mathcal{H}(\mathcal{S})$, which is $\mathbb{E}[\sup_{\theta \in \mathcal{H}(\mathcal{S})} X_{\theta}]$.

So by Lemma 20 and (9), we know that: $\forall \delta \in (0, D]$:

$$\mathcal{R}_{\mathcal{S}}(\mathcal{H}) = \mathbb{E}[\sup_{\theta \in \mathcal{H}(\mathcal{S})} X_{\theta}] \leq \mathbb{E}\left[\sup_{\theta, \tilde{\theta} \in \mathcal{H}(\mathcal{S})} (X_{\theta} - X_{\tilde{\theta}})\right] \leq 2\mathbb{E}\left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}(\mathcal{S}) \\ \|\gamma' - \gamma\|_2 \leq \delta}} (X_{\gamma} - X_{\gamma'})\right] + 32\mathcal{J}(\delta/4; D),$$

where

$$D = \sup_{\theta, \theta' \in \mathcal{H}(\mathcal{S})} \|\theta - \theta'\|_2 \leq 2 \sup_{\theta \in \mathcal{H}(\mathcal{S})} \|\theta\|_2 = 2 \sup_{h \in \mathcal{H}} \|h\|_2 = 2 \sup_{h \in \mathcal{H}} \sqrt{\sum_{i=1}^M [h(z_i)]^2} \leq 2\sqrt{M} \sup_{\substack{h \in \mathcal{H} \\ 1 \leq i \leq M}} |h(z_i)|, \quad (\text{A.26})$$

and $\mathcal{J}(a; b) = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}(\mathcal{S}), \|\cdot\|_2)} du = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}, \|\cdot\|_2)} du$.

For any $f \in \mathcal{F}$, $x \in \mathcal{X}$, let x^l be the output of x passing through the first $l-1$ layers, we have:

$$\begin{aligned} \|f(x)\|_2 &= \|W_d \sigma(W_{d-1} x^{d-1})\|_2 \stackrel{(i)}{\leq} \| \|W_d\|_F \cdot \|\sigma(W_{d-1} x^{d-1})\|_2 \stackrel{(ii)}{=} \| \|W_d\|_F \cdot \|\sigma(W_{d-1} x^{d-1}) - \sigma(\mathbf{0})\|_2 \\ &\stackrel{(iii)}{\leq} LM_d^F \|W_{d-1} x^{d-1}\|_2 \leq \dots \leq L^{d-1} \prod_{l=2}^d M_l^F \|W_1 x\|_2 \leq L^{d-1} \prod_{l=1}^d M_l^F \|x\|_2 \\ &\stackrel{(iv)}{\leq} L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2}-\frac{1}{p}}\right\} \|x\|_p \leq L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2}-\frac{1}{p}}\right\} \|X\|_{p, \infty}, \end{aligned}$$

where (i) is from Lemma 40; (ii) is from the fact that $\sigma(\mathbf{0}) = 0$; (iii) comes from the assumption that σ is L -Lipschitz and $\| \|W_d\|_F \leq M_d^F$ and (iv) is attained by setting $r = r^* = 2$ in the proof of Lemma 39.

To simplify the notations, we define:

$$B_{X, \epsilon}^F = L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2}-\frac{1}{p}}\right\} (\|X\|_{p, \infty} + \epsilon), \quad B_X^F = L^{d-1} \prod_{l=1}^d M_l^F \max\left\{1, m^{\frac{1}{2}-\frac{1}{p}}\right\} \|X\|_{p, \infty}.$$

So we have:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \|f(x)\|_2 \leq B_X^F. \quad (\text{A.27})$$

Similarly, we have:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \forall \|x' - x\|_p \leq \epsilon, \|f(x')\|_2 \leq B_{X, \epsilon}^F. \quad (\text{A.28})$$

For any $h \in \mathcal{H}$, $z \in \mathcal{X}^3$, let $x^* = \arg \min_{\|x' - x\|_p \leq \epsilon} f(x')^T (f(x^+) - f(x^-))$ and let x^l be the output of x^* passing through the first $l - 1$ layers, we have:

$$\begin{aligned} |h(z)| &= \left| \inf_{\|x' - x\|_p \leq \epsilon} f(x')^T (f(x^+) - f(x^-)) \right| = |f(x^*)^T (f(x^+) - f(x^-))| \\ &\leq \|f(x^*)\|_2 \cdot \|f(x^+) - f(x^-)\|_2 \stackrel{(i)}{\leq} B_{X,\epsilon}^F \cdot (B_{X^+}^F + B_{X^-}^F), \end{aligned}$$

where (i) is from (A.27) and (A.28). So we have:

$$D \stackrel{(i)}{\leq} 2\sqrt{M} \sup_{\substack{h \in \mathcal{H} \\ 1 \leq i \leq M}} |h(z_i)| \leq 2\sqrt{M} B_{X,\epsilon}^F \cdot (B_{X^+}^F + B_{X^-}^F) \triangleq \sqrt{M} K, \quad (\text{A.29})$$

where (i) is from (A.26). Now, we need to find the smallest distance between \mathcal{H} and \mathcal{H}^c , i.e.

$$\sup_{h \in \mathcal{H}} \inf_{h^c \in \mathcal{H}^c} \|h - h^c\|_2.$$

By the discussion in (A.26), we have $\|h - h^c\|_2 \leq \sqrt{M} \max_{1 \leq i \leq M} |h(z_i) - h^c(z_i)|$. For any $z_i = (x_i, x_i^+, x_i^-)$, $i = 1, 2, \dots, M$, given h and h^c such that $\|W_l - W_l^c\|_F \leq \delta_l$, $l = 1, 2, \dots, d$, we have:

$$|h(z_i) - h^c(z_i)| = \left| \inf_{\|x'_i - x_i\|_p \leq \epsilon} f(x'_i)^T (f(x_i^+) - f(x_i^-)) - \inf_{\|x'_i - x_i\|_p \leq \epsilon} f^c(x'_i)^T (f^c(x_i^+) - f^c(x_i^-)) \right|.$$

Let $x_i^* = \arg \inf_{\|x'_i - x_i\|_p \leq \epsilon} f(x'_i)^T (f(x_i^+) - f(x_i^-))$ and $x_i^c = \arg \inf_{\|x'_i - x_i\|_p \leq \epsilon} f^c(x'_i)^T (f^c(x_i^+) - f^c(x_i^-))$, and let

$$y_i = \begin{cases} x_i^c & f(x_i^*)^T (f(x_i^+) - f(x_i^-)) \geq f^c(x_i^c)^T (f^c(x_i^+) - f^c(x_i^-)) \\ x_i^* & \text{otherwise} \end{cases}.$$

Then we have:

$$\begin{aligned} |h(z_i) - h^c(z_i)| &= |f(x_i^*)^T (f(x_i^+) - f(x_i^-)) - f^c(x_i^c)^T (f^c(x_i^+) - f^c(x_i^-))| \\ &\stackrel{(i)}{\leq} |f(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f^c(x_i^+) - f^c(x_i^-))| \\ &= |f(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f(x_i^+) - f(x_i^-)) \\ &\quad + f^c(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f^c(x_i^+) - f^c(x_i^-))| \\ &\stackrel{(ii)}{\leq} |(f(y_i) - f^c(y_i))^T (f(x_i^+) - f(x_i^-))| \\ &\quad + |f^c(y_i)^T (f(x_i^+) - f^c(x_i^+))| + |f^c(y_i)^T (f(x_i^-) - f^c(x_i^-))| \\ &\stackrel{(iii)}{\leq} (B_{X^+}^F + B_{X^-}^F) \|f(y_i) - f^c(y_i)\|_2 + B_{X,\epsilon}^F \|f(x_i^+) - f^c(x_i^+)\|_2 \\ &\quad + B_{X,\epsilon}^F \|f(x_i^-) - f^c(x_i^-)\|_2, \end{aligned} \quad (\text{A.30})$$

where (i) is easily verified by the definition of y_i ; (ii) is from the triangle inequality and (iii) is from (A.27) and (A.28).

Define $g_b^a(\cdot)$ as:

$$g_b^a(y) = W_b \sigma(W_{b-1} \sigma(\cdots W_{a+1} \sigma(W_a^c \cdots \sigma(W_1^c y)))) .$$

Then we have:

$$\begin{aligned} \|f(y_i) - f^c(y_i)\|_2 &= \|g_d^0(y_i) - g_d^d(y_i)\|_2 \\ &= \|g_d^0(y_i) - g_d^1(y_i) + g_d^1(y_i) - g_d^2(y_i) + \cdots + g_d^{d-1}(y_i) - g_d^d(y_i)\|_2 \\ &\stackrel{(i)}{\leq} \|g_d^0(y_i) - g_d^1(y_i)\|_2 + \cdots + \|g_d^{d-1}(y_i) - g_d^d(y_i)\|_2, \end{aligned}$$

where (i) is from the triangle inequality.

Then we calculate $\|g_d^{l-1}(y_i) - g_d^l(y_i)\|_2, l = 1, 2, \dots, d$:

$$\begin{aligned} \|g_d^{l-1}(y_i) - g_d^l(y_i)\|_2 &= \|W_d \sigma(g_{d-1}^{l-1}(y_i)) - W_d \sigma(g_{d-1}^l(y_i))\|_2 \\ &\stackrel{(i)}{\leq} \|W_d\|_F \cdot \|\sigma(g_{d-1}^{l-1}(y_i)) - \sigma(g_{d-1}^l(y_i))\|_2 \\ &\stackrel{(ii)}{\leq} L M_d^F \|g_{d-1}^{l-1}(y_i) - g_{d-1}^l(y_i)\|_2 \leq \cdots \\ &\stackrel{(iii)}{\leq} L^{d-l} \cdot \prod_{j=l+1}^d M_j^F \cdot \|W_l \sigma(g_{l-1}^{l-1}(y_i)) - W_l^c \sigma(g_{l-1}^{l-1}(y_i))\|_2 \quad (\text{A.31}) \\ &= L^{d-l} \cdot \prod_{j=l+1}^d M_j^F \cdot \|(W_l - W_l^c) \sigma(g_{l-1}^{l-1}(y_i))\|_2 \\ &\stackrel{(iv)}{\leq} L^{d-l} \cdot \prod_{j=l+1}^d M_j^F \cdot \delta_l \cdot \|\sigma(g_{l-1}^{l-1}(y_i))\|_2, \end{aligned}$$

where (i) is from Lemma 40; (ii) comes from the assumption that σ is L -Lipschitz and $\|W_d\|_F \leq M_d^F$; (iii) is from the definition of $g_b^a(\cdot)$ and (iv) is from Lemma 40 and the choice of h^c when h is fixed, which means that $\|W_l - W_l^c\|_F \leq \delta_l$.

Next, we upper bound $\|\sigma(g_{l-1}^{l-1}(y_i))\|_2$:

$$\begin{aligned} \|\sigma(g_{l-1}^{l-1}(y_i))\|_2 &= \|\sigma(g_{l-1}^{l-1}(y_i)) - \sigma(\mathbf{0})\|_2 \stackrel{(i)}{\leq} L \cdot \|g_{l-1}^{l-1}(y_i)\|_2 = L \cdot \|W_{l-1}^c \sigma(g_{l-2}^{l-2}(y_i))\|_2 \\ &\stackrel{(ii)}{\leq} L \cdot \|W_{l-1}^c\|_F \cdot \|\sigma(g_{l-2}^{l-2}(y_i))\|_2 \stackrel{(iii)}{\leq} L M_{l-1}^F \|\sigma(g_{l-2}^{l-2}(y_i))\|_2 \\ &\leq \cdots \leq L^{l-1} \cdot \prod_{j=1}^{l-1} M_j^F \cdot \|y_i\|_2, \end{aligned} \tag{A.32}$$

where (i) is because σ is L -Lipschitz; (ii) is from Lemma 40 and (iii) is because $\| \|W_{l-1}^c\| \|_F \leq M_{l-1}^F$.

From (A.31) and (A.32) we have:

$$\begin{aligned} \|g_d^{l-1}(y_i) - g_d^l(y_i)\|_2 &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \|y_i\|_2 \\ &\stackrel{(i)}{\leq} L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \max \left\{ 1, m^{\frac{1}{2} - \frac{1}{p}} \right\} (\|X\|_{p,\infty} + \epsilon) = B_{X,\epsilon}^F \frac{\delta_l}{M_l^F}, \end{aligned}$$

where (i) is from Lemma 39.

Similarly:

$$\begin{aligned} \|g_d^{l-1}(x_i^+) - g_d^l(x_i^+)\|_2 &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \|x_i^+\|_2 \\ &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \max \left\{ 1, m^{\frac{1}{2} - \frac{1}{p}} \right\} \|X^+\|_{p,\infty} = B_{X^+}^F \frac{\delta_l}{M_l^F}, \\ \|g_d^{l-1}(x_i^-) - g_d^l(x_i^-)\|_2 &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \|x_i^-\|_2 \\ &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^F}{M_l^F} \delta_l \max \left\{ 1, m^{\frac{1}{2} - \frac{1}{p}} \right\} \|X^-\|_{p,\infty} = B_{X^-}^F \frac{\delta_l}{M_l^F}. \end{aligned}$$

Combining the above with (A.30) yields:

$$\begin{aligned} |h(z_i) - h^c(z_i)| &\leq (B_{X^+}^F + B_{X^-}^F) \left(\|g_d^0(y_i) - g_d^1(y_i)\|_2 + \cdots + \|g_d^{d-1}(y_i) - g_d^d(y_i)\|_2 \right) \\ &\quad + B_{X,\epsilon}^F \left(\|g_d^0(x_i^+) - g_d^1(x_i^+)\|_2 + \cdots + \|g_d^{d-1}(x_i^+) - g_d^d(x_i^+)\|_2 \right) \\ &\quad + B_{X,\epsilon}^F \left(\|g_d^0(x_i^-) - g_d^1(x_i^-)\|_2 + \cdots + \|g_d^{d-1}(x_i^-) - g_d^d(x_i^-)\|_2 \right) \\ &= (B_{X^+}^F + B_{X^-}^F) B_{X,\epsilon}^F \sum_{l=1}^d \frac{\delta_l}{M_l^F} + B_{X,\epsilon}^F B_{X^+}^F \sum_{l=1}^d \frac{\delta_l}{M_l^F} + B_{X,\epsilon}^F B_{X^-}^F \sum_{l=1}^d \frac{\delta_l}{M_l^F} \\ &= 2B_{X,\epsilon}^F (B_{X^+}^F + B_{X^-}^F) \sum_{l=1}^d \frac{\delta_l}{M_l^F} = K \sum_{l=1}^d \frac{\delta_l}{M_l^F}. \end{aligned}$$

So $\|h - h^c\|_2 = \sqrt{M} \max_{1 \leq i \leq M} |h(z_i) - h^c(z_i)| \leq \sqrt{M} \sum_{l=1}^d \frac{K \delta_l}{M_l^F}$. Let $\delta_l = \frac{M_l^F \delta}{dK\sqrt{M}}$, we have:

$$\|h - h^c\|_2 \leq \sqrt{M} \sum_{l=1}^d \frac{K}{M_l^F} \cdot \frac{M_l^F \delta}{dK\sqrt{M}} \leq \delta.$$

Then: $\forall h \in \mathcal{H}, \exists h^c \in \mathcal{H}^c$ s.t. $\|h - h^c\|_2 \leq \delta$, which means that $\sup_{h \in \mathcal{H}} \inf_{h^c \in \mathcal{H}^c} \|h - h^c\|_2 \leq \delta$ when

choosing $\delta_l = \frac{M_l^F \delta}{dK\sqrt{M}}$.

So \mathcal{H}^c is a δ -covering of \mathcal{H} , and $\mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2) \leq |\mathcal{H}^c| = \prod_{l=1}^d |\mathcal{C}_l^F|$. By Lemma 18 we know that $|\mathcal{C}_l^F| = \mathcal{N}(\frac{M_l^F \delta}{dK\sqrt{M}}; S_l^F, \|\cdot\|_F) \leq \left(1 + \frac{2dK\sqrt{M}}{\delta}\right)^{h_l \times h_{l-1}}$. So we have:

$$\mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2) \leq |\mathcal{H}^c| = \prod_{l=1}^d |\mathcal{C}_l^F| \leq \left(1 + \frac{2dK\sqrt{M}}{\delta}\right)^{\sum_{l=1}^d h_l \cdot h_{l-1}}. \quad (\text{A.33})$$

So we can conclude that:

$$\begin{aligned} \mathcal{J}(0; D) &= \int_0^D \sqrt{\ln \mathcal{N}(u; \mathcal{H}, \|\cdot\|_2)} du \stackrel{(i)}{\leq} \int_0^D \sqrt{\left(\sum_{l=1}^d h_l \cdot h_{l-1}\right) \ln \left(1 + \frac{2dK\sqrt{M}}{u}\right)} du \\ &\stackrel{(ii)}{\leq} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1}} \int_0^D \sqrt{\frac{2dK\sqrt{M}}{u}} du = \sqrt{2dK \sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{M}} \int_0^D u^{-\frac{1}{2}} du \\ &= 2\sqrt{2dKD \sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{M}} \stackrel{(iii)}{\leq} 2\sqrt{2} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{dK\sqrt{M}}}, \end{aligned}$$

where (i) is from (A.33); (ii) comes from the fact that $\ln(1+x) \leq x, \forall x \geq 0$ and (iii) comes from (A.29).

Since we shows that $\mathcal{R}_S(\mathcal{H}) \leq 2\mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}(S) \\ \|\gamma' - \gamma\|_2 \leq \delta}} (X_\gamma - X_{\gamma'}) \right] + 32\mathcal{J}(\delta/4; D)$ before, take

$\delta \rightarrow 0^+$, we have:

$$\mathcal{R}_S(\mathcal{H}) \leq 32\mathcal{J}(0; D) \leq 64\sqrt{2} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{dK\sqrt{M}}} = O \left(\sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{dK\sqrt{M}}} \right).$$

■

A.10 Proof of Theorem 26

Theorem 26 Let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$ (i.e. consider the ℓ_p attack), $\sigma(0) = 0$ with Lipschitz constant L ; moreover, let $\mathcal{F} = \left\{W_d \sigma(W_{d-1} \sigma(\dots \sigma(W_1 x))) \mid \|W_l\|_{1, \infty} \leq M_l^{1, \infty}, l=1, \dots, d\right\}$. We then have:

$$\mathcal{R}_S(\mathcal{H}) = O \left(\sqrt{\sum_{l=1}^d h_l h_{l-1} \sqrt{dK_0 K_1 \sqrt{M}}} \right),$$

where

$$K_0 = 2B_{X, \epsilon}^{1, \infty} \cdot (B'_{X^+} + B'_{X^-}), \quad K_1 = \frac{K_0}{2} + B'_{X, \epsilon} \cdot (B_{X^+}^{1, \infty} + B_{X^-}^{1, \infty}),$$

where

$$B'_{X,\epsilon} = L^{d-1} \prod_{l=1}^d h_l M_l^{1,\infty} m^{1-\frac{1}{p}} (\|X\|_{p,\infty} + \epsilon), \quad B'_X = L^{d-1} \prod_{l=1}^d h_l M_l^{1,\infty} m^{1-\frac{1}{p}} \|X\|_{p,\infty},$$

$$B_{X,\epsilon}^{1,\infty} = L^{d-1} \prod_{l=1}^d M_l^{1,\infty} (\|X\|_{p,\infty} + \epsilon), \quad B_X^{1,\infty} = L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \|X\|_{p,\infty}.$$

Before giving the proof, we firstly introduce some useful lemmas.

Lemma 42 *Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^n$, then we have: $\|A \cdot b\|_\infty \leq \|A\|_{1,\infty} \cdot \|b\|_\infty$.*

Proof [Proof of Lemma 42] Let A_i be the rows of A , $i = 1, 2, \dots, m$, we have:

$$\|A \cdot b\|_\infty = \max_{1 \leq i \leq m} |A_i b| \stackrel{(i)}{\leq} \max_{1 \leq i \leq m} (\|A_i\|_1 \cdot \|b\|_\infty) = \|A\|_{1,\infty} \cdot \|b\|_\infty,$$

where (i) is from the Holder's Inequality. ■

Lemma 43 *Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^n$, then we have:*

$$\|A \cdot b\|_1 \leq \|A\|_{\infty,1} \cdot \|b\|_1 \leq m \|A\|_{1,\infty} \cdot \|b\|_1.$$

Proof [Proof of Lemma 43] Let A_i be the rows of A , $i = 1, 2, \dots, m$, we have:

$$\|A \cdot b\|_1 = \sum_{i=1}^m |A_i b| \stackrel{(i)}{\leq} \sum_{i=1}^m (\|A_i\|_\infty \cdot \|b\|_1) = \|A\|_{\infty,1} \cdot \|b\|_1,$$

where (i) is from the Holder's Inequality. And we have:

$$\begin{aligned} \|A\|_{\infty,1} &= \|(\|A_1\|_\infty, \dots, \|A_m\|_\infty)\|_1 \stackrel{(i)}{\leq} \|(\|A_1\|_1, \dots, \|A_m\|_1)\|_1 \\ &\stackrel{(ii)}{\leq} m \|(\|A_1\|_1, \dots, \|A_m\|_1)\|_\infty = m \|A\|_{1,\infty}, \end{aligned}$$

where (i) is from the fact that $\|x\|_\infty \leq \|x\|_1$ and (ii) is from the fact that for all $x \in \mathcal{R}^m$, $\|x\|_1 \leq m \|x\|_\infty$. So we have: $\|A \cdot b\|_1 \leq \|A\|_{\infty,1} \cdot \|b\|_1 \leq m \|A\|_{1,\infty} \cdot \|b\|_1$. ■

Lemma 44 *Suppose σ is a L -Lipschitz function, then the elementwise vector map corresponding to σ is also L -Lipschitz with respect to $\|\cdot\|_\infty$.*

Proof [Proof of Lemma 44]

$$\begin{aligned} \|\sigma(x) - \sigma(y)\|_\infty &= \max_{1 \leq i \leq n} |\sigma(x)_i - \sigma(y)_i| = \max_{1 \leq i \leq n} |\sigma(x_i) - \sigma(y_i)| \\ &\stackrel{(i)}{\leq} \max_{1 \leq i \leq n} L|x_i - y_i| = L \cdot \|x - y\|_\infty, \end{aligned}$$

where (i) is because σ is L -Lipschitz. ■

Now we can turn to the proof of Theorem 26.

Proof [Proof of Theorem 26] In this case, let $\mathcal{U}(x) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$, we have:

$$\mathcal{F} = \left\{ x \rightarrow W_d \sigma(W_{d-1} \sigma(\cdots \sigma(W_1 x))) \mid \|W_l\|_{1,\infty} \leq M_l^{1,\infty}, l = 1, \dots, d \right\},$$

$$\mathcal{H} = \left\{ h(x, x^+, x^-) = \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F} \right\}.$$

Let $S_l^{1,\infty} = \{W_l \in \mathbb{R}^{h_l \times h_{l-1}} \mid \|W_l\|_{1,\infty} \leq M_l^{1,\infty}\}$, $l = 1, 2, \dots, d$. Let $\mathcal{C}_l^{1,\infty}$ be the δ_l -covering of $S_l^{1,\infty}$ and define:

$$\mathcal{F}^c = \left\{ f^c : x \rightarrow W_d^c \sigma(W_{d-1}^c \sigma(\cdots \sigma(W_1^c x))) \mid W_l^c \in \mathcal{C}_l^{1,\infty}, l = 1, 2, \dots, d \right\} \subseteq \mathcal{F},$$

$$\mathcal{H}^c = \left\{ h^c(x, x^+, x^-) = \min_{x' \in \mathcal{U}(x)} (f(x')^T (f(x^+) - f(x^-))) \mid f \in \mathcal{F}^c \right\} \subseteq \mathcal{H}.$$

Similar to the proof of Theorem 22, we know that the Rademacher Process is a sub-Gaussian Process with respect to the Euclidean metric, which induces the ℓ_2 norm.

Similar to the proof of Theorem 24, given the training set $\mathcal{S} = \{(x_i, x_i^+, x_i^-)\}_{i=1}^M \triangleq \{z_i\}_{i=1}^M$, define the ℓ_2 -norm for a function in \mathcal{H} as:

$$\forall h \in \mathcal{H}, \|h\|_2 := \sqrt{\sum_{i=1}^M [h(z_i)]^2}.$$

Define $\mathcal{H}(\mathcal{S}) = \{(h(z_1), h(z_2), \dots, h(z_M)) \mid h \in \mathcal{H}\}$, with the same argument as in proof of Theorem 24, we have:

$$\mathcal{N}(\delta; \mathcal{H}(\mathcal{S}), \|\cdot\|_2) = \mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2),$$

and

$$\mathcal{R}_{\mathcal{S}}(\mathcal{H}) = \mathbb{E}[\sup_{\theta \in \mathcal{H}(\mathcal{S})} X_{\theta}].$$

So by Lemma 20 and (9), we know that: $\forall \delta \in (0, D]$:

$$\mathcal{R}_{\mathcal{S}}(\mathcal{H}) = \mathbb{E}[\sup_{\theta \in \mathcal{H}(\mathcal{S})} X_{\theta}] \leq \mathbb{E} \left[\sup_{\theta, \tilde{\theta} \in \mathcal{H}(\mathcal{S})} (X_{\theta} - X_{\tilde{\theta}}) \right] \leq 2 \mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}(\mathcal{S}) \\ \|\gamma - \gamma'\|_2 \leq \delta}} (X_{\gamma} - X_{\gamma'}) \right] + 32 \mathcal{J}(\delta/4; D),$$

where

$$D = \sup_{\theta, \theta' \in \mathcal{H}(\mathcal{S})} \|\theta - \theta'\|_2 \leq 2 \sup_{\theta \in \mathcal{H}(\mathcal{S})} \|\theta\|_2 = 2 \sup_{h \in \mathcal{H}} \|h\|_2 = 2 \sup_{h \in \mathcal{H}} \sqrt{\sum_{i=1}^M [h(z_i)]^2} \leq 2\sqrt{M} \sup_{\substack{h \in \mathcal{H} \\ 1 \leq i \leq M}} |h(z_i)|, \quad (\text{A.34})$$

and $\mathcal{J}(a; b) = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}(\mathcal{S}), \|\cdot\|_2)} du = \int_a^b \sqrt{\ln \mathcal{N}(u; \mathcal{H}, \|\cdot\|_2)} du$

Then for any $f \in \mathcal{F}, x \in \mathcal{X}$, let x^l be the output of x passing through the first $l - 1$ layers, we have:

$$\begin{aligned}
 \|f(x)\|_\infty &= \|W_d \sigma(W_{d-1} x^{d-1})\|_\infty \stackrel{(i)}{\leq} \|W_d\|_{1,\infty} \cdot \|\sigma(W_{d-1} x^{d-1})\|_\infty \\
 &\stackrel{(ii)}{=} \|W_d\|_{1,\infty} \cdot \|\sigma(W_{d-1} x^{d-1}) - \sigma(\mathbf{0})\|_\infty \stackrel{(iii)}{\leq} L M_d^{1,\infty} \|W_{d-1} x^{d-1}\|_\infty \leq \dots \\
 &\leq L^{d-1} \prod_{l=2}^d M_l^{1,\infty} \|W_1 x\|_\infty \leq L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \|x\|_\infty \stackrel{(iv)}{\leq} L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \max\{1, m^{-p}\} \|x\|_p \\
 &= L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \|x\|_p \leq L^{d-1} \prod_{l=1}^d M_l^{1,\infty} \|X\|_{p,\infty},
 \end{aligned}$$

where (i) is from Lemma 42; (ii) is from the fact that $\sigma(0) = 0$; (iii) comes from the assumption that σ is L -Lipschitz and $\|W_d\|_{1,\infty} \leq M_d^{1,\infty}$ and (iv) is attained by setting $r = 1, r^* = \infty$ in the proof of Lemma 39.

To simplify the notations, we define:

$$\begin{aligned}
 B_X^{1,\infty} &= L^{d-1} \cdot \prod_{l=1}^d M_l^{1,\infty} \cdot \|X\|_{p,\infty}, B'_X = L^{d-1} \cdot \left(\prod_{l=1}^d h_l \cdot M_l^{1,\infty} \right) m^{1-\frac{1}{p}} \cdot \|X\|_{p,\infty} \\
 B_{X,\epsilon}^{1,\infty} &= L^{d-1} \cdot \prod_{l=1}^d M_l^{1,\infty} \cdot (\|X\|_{p,\infty} + \epsilon), B'_{X,\epsilon} = L^{d-1} \cdot \left(\prod_{l=1}^d h_l \cdot M_l^{1,\infty} \right) m^{1-\frac{1}{p}} \cdot (\|X\|_{p,\infty} + \epsilon).
 \end{aligned}$$

So:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \|f(x)\|_\infty \leq B_X^{1,\infty}.$$

Similarly, we have:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \forall \|x' - x\|_p \leq \epsilon, \|f(x')\|_\infty \leq B_{X,\epsilon}^{1,\infty}. \quad (\text{A.35})$$

Similarly,

$$\begin{aligned}
 \|f(x)\|_1 &= \|W_d \sigma(W_{d-1} x^{d-1})\|_1 \stackrel{(i)}{\leq} h_d \cdot \|W_d\|_{1,\infty} \cdot \|\sigma(W_{d-1} x^{d-1})\|_1 \\
 &\stackrel{(ii)}{=} h_d \cdot \|W_d\|_{1,\infty} \cdot \|\sigma(W_{d-1} x^{d-1}) - \sigma(\mathbf{0})\|_1 \stackrel{(iii)}{\leq} h_d \cdot L \cdot M_d^{1,\infty} \|W_{d-1} x^{d-1}\|_1 \leq \dots \\
 &\leq L^{d-1} \prod_{l=2}^d h_l M_l^{1,\infty} \|W_1 x\|_1 \leq L^{d-1} \prod_{l=1}^d h_l M_l^{1,\infty} \|x\|_1 \\
 &\stackrel{(iv)}{\leq} L^{d-1} \left(\prod_{l=1}^d h_l M_l^{1,\infty} \right) \max\{1, m^{1-\frac{1}{p}}\} \|x\|_p \\
 &= L^{d-1} \left(\prod_{l=1}^d h_l M_l^{1,\infty} \right) m^{1-\frac{1}{p}} \|x\|_p \leq L^{d-1} \left(\prod_{l=1}^d h_l M_l^{1,\infty} \right) m^{1-\frac{1}{p}} \|X\|_{p,\infty},
 \end{aligned}$$

where (i) is from Lemma 43; (ii) is from the fact that $\sigma(0) = 0$; (iii) comes from the assumption that σ is L -Lipschitz and $\|W_d\|_{1,\infty} \leq M_d^{1,\infty}$ and (iv) is attained by setting $r = \infty, r^* = 1$ in the proof of Lemma 39.

So we know that:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \|f(x)\|_1 \leq B'_X. \quad (\text{A.36})$$

Similarly, we have:

$$\forall x \in \mathcal{X}, f \in \mathcal{F}, \forall \|x' - x\|_p \leq \epsilon, \|f(x')\|_1 \leq B'_{X,\epsilon}. \quad (\text{A.37})$$

For any $h \in \mathcal{H}, z \in \mathcal{X}^3$, let $x^* = \arg \min_{\|x' - x\|_p \leq \epsilon} f(x')^T (f(x^+) - f(x^-))$ and let x^l be the output of x^* passing through the first $l - 1$ layers, we have:

$$\begin{aligned} |h(z)| &= \left| \inf_{\|x' - x\|_p \leq \epsilon} f(x')^T (f(x^+) - f(x^-)) \right| = |f(x^*)^T (f(x^+) - f(x^-))| \\ &\leq \|f(x^*)\|_\infty \cdot \|f(x^+) - f(x^-)\|_1 \stackrel{(i)}{\leq} B_{X,\epsilon}^{1,\infty} \cdot (B'_{X^+} + B'_{X^-}), \end{aligned}$$

where (i) is from (A.35) and (A.36). So we get:

$$D \stackrel{(i)}{\leq} 2\sqrt{M} \sup_{\substack{h \in \mathcal{H} \\ 1 \leq i \leq M}} |h(z_i)| \leq 2\sqrt{M} B_{X,\epsilon}^{1,\infty} \cdot (B'_{X^+} + B'_{X^-}) \triangleq \sqrt{M} K_0, \quad (\text{A.38})$$

where (i) is from (A.34). Now, we need to find the smallest distance between \mathcal{H} and \mathcal{H}^c , i.e.

$$\sup_{h \in \mathcal{H}} \inf_{h^c \in \mathcal{H}^c} \|h - h^c\|_2.$$

By the discussion in (A.34), we have $\|h - h^c\|_2 \leq \sqrt{M} \max_{1 \leq i \leq M} |h(z_i) - h^c(z_i)|$. For any $z_i = (x_i, x_i^+, x_i^-), i = 1, 2, \dots, M$, given h and h^c such that $\|W_l - W_l^c\|_{1,\infty} \leq \delta_l, l = 1, 2, \dots, d$, we have:

$$|h(z_i) - h^c(z_i)| = \left| \inf_{\|x'_i - x_i\|_p \leq \epsilon} f(x'_i)^T (f(x_i^+) - f(x_i^-)) - \inf_{\|x'_i - x_i\|_p \leq \epsilon} f^c(x'_i)^T (f^c(x_i^+) - f^c(x_i^-)) \right|.$$

Let $x_i^* = \arg \inf_{\|x'_i - x_i\|_p \leq \epsilon} f(x'_i)^T (f(x_i^+) - f(x_i^-))$ and $x_i^c = \arg \inf_{\|x'_i - x_i\|_p \leq \epsilon} f^c(x'_i)^T (f^c(x_i^+) - f^c(x_i^-))$,

and let

$$y_i = \begin{cases} x_i^c & f(x_i^*)^T (f(x_i^+) - f(x_i^-)) \geq f^c(x_i^c)^T (f^c(x_i^+) - f^c(x_i^-)) \\ x_i^* & \text{otherwise} \end{cases}.$$

Then we have:

$$\begin{aligned}
 |h(z_i) - h^c(z_i)| &= |f(x_i^*)^T (f(x_i^+) - f(x_i^-)) - f^c(x_i^c)^T (f^c(x_i^+) - f^c(x_i^-))| \\
 &\stackrel{(i)}{\leq} |f(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f^c(x_i^+) - f^c(x_i^-))| \\
 &= |f(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f(x_i^+) - f(x_i^-)) \\
 &\quad + f^c(y_i)^T (f(x_i^+) - f(x_i^-)) - f^c(y_i)^T (f^c(x_i^+) - f^c(x_i^-))| \\
 &\stackrel{(ii)}{\leq} |(f(y_i) - f^c(y_i))^T (f(x_i^+) - f(x_i^-))| \\
 &\quad + |f^c(y_i)^T (f(x_i^+) - f^c(x_i^+))| + |f^c(y_i)^T (f(x_i^-) - f^c(x_i^-))| \\
 &\stackrel{(iii)}{\leq} (B'_{X^+} + B'_{X^-}) \|f(y_i) - f^c(y_i)\|_\infty + B'_{X,\epsilon} \|f(x_i^+) - f^c(x_i^+)\|_\infty \\
 &\quad + B'_{X,\epsilon} \|f(x_i^-) - f^c(x_i^-)\|_\infty,
 \end{aligned} \tag{A.39}$$

where (i) is easily verified by the definition of y_i ; (ii) is from the triangle inequality and (iii) is from (A.36) and (A.37). Again we define $g_b^a(\cdot)$ as:

$$g_b^a(y) = W_b \sigma(W_{b-1} \sigma(\cdots W_{a+1} \sigma(W_a^c \cdots \sigma(W_1^c y)))) .$$

Then:

$$\begin{aligned}
 \|f(y_i) - f^c(y_i)\|_\infty &= \|g_d^0(y_i) - g_d^d(y_i)\|_\infty \\
 &= \|g_d^0(y_i) - g_d^1(y_i) + g_d^1(y_i) - g_d^2(y_i) + \cdots + g_d^{d-1}(y_i) - g_d^d(y_i)\|_\infty \\
 &\stackrel{(i)}{\leq} \|g_d^0(y_i) - g_d^1(y_i)\|_\infty + \cdots + \|g_d^{d-1}(y_i) - g_d^d(y_i)\|_\infty,
 \end{aligned}$$

where (i) is from the triangle inequality.

Then we calculate $\|g_d^{l-1}(y_i) - g_d^l(y_i)\|_\infty, l = 1, 2, \dots, d$:

$$\begin{aligned}
 \|g_d^{l-1}(y_i) - g_d^l(y_i)\|_\infty &= \|W_d \sigma(g_{d-1}^{l-1}(y_i)) - W_d \sigma(g_{d-1}^l(y_i))\|_\infty \\
 &\stackrel{(i)}{\leq} \|W_d\|_{1,\infty} \cdot \|\sigma(g_{d-1}^{l-1}(y_i)) - \sigma(g_{d-1}^l(y_i))\|_\infty \\
 &\stackrel{(ii)}{\leq} LM_d^{1,\infty} \|g_{d-1}^{l-1}(y_i) - g_{d-1}^l(y_i)\|_\infty \leq \cdots \\
 &\stackrel{(iii)}{\leq} L^{d-l} \cdot \prod_{j=l+1}^d M_j^{1,\infty} \cdot \|W_l \sigma(g_{l-1}^{l-1}(y_i)) - W_l^c \sigma(g_{l-1}^{l-1}(y_i))\|_\infty \tag{A.40} \\
 &= L^{d-l} \cdot \prod_{j=l+1}^d M_j^{1,\infty} \cdot \|(W_l - W_l^c) \sigma(g_{l-1}^{l-1}(y_i))\|_\infty \\
 &\stackrel{(iv)}{\leq} L^{d-l} \cdot \prod_{j=l+1}^d M_j^{1,\infty} \cdot \delta_l \cdot \|\sigma(g_{l-1}^{l-1}(y_i))\|_\infty,
 \end{aligned}$$

where (i) is from Lemma 42; (ii) comes from the assumption that σ is L -Lipschitz and $\|W_d\|_{1,\infty} \leq M_d^{1,\infty}$; (iii) is from the definition of $g_b^a(\cdot)$ and (iv) is from Lemma 42 and the choice of h^c when h is fixed, which means that $\|W_l - W_l^c\|_{1,\infty} \leq \delta_l$.

Next, we upper bound $\|\sigma(g_{l-1}^{l-1}(y_i))\|_\infty$:

$$\begin{aligned} \|\sigma(g_{l-1}^{l-1}(y_i))\|_\infty &= \|\sigma(g_{l-1}^{l-1}(y_i)) - \sigma(\mathbf{0})\|_\infty \stackrel{(i)}{\leq} L \cdot \|g_{l-1}^{l-1}(y_i)\|_\infty \\ &= L \cdot \|W_{l-1}^c \sigma(g_{l-2}^{l-2}(y_i))\|_\infty \stackrel{(ii)}{\leq} L \cdot \|W_{l-1}^c\|_{1,\infty} \cdot \|\sigma(g_{l-2}^{l-2}(y_i))\|_\infty \quad (\text{A.41}) \\ &\stackrel{(iii)}{\leq} L M_{l-1}^{1,\infty} \|\sigma(g_{l-2}^{l-2}(y_i))\|_\infty \leq \dots \leq L^{l-1} \cdot \prod_{j=1}^{l-1} M_j^{1,\infty} \cdot \|y_i\|_\infty, \end{aligned}$$

where (i) is because σ is L -Lipschitz; (ii) is from Lemma 42 and (iii) is because $\|W_{l-1}^c\|_{1,\infty} \leq M_{l-1}^{1,\infty}$.

(A.40) and (A.41) show that:

$$\begin{aligned} \|g_d^{l-1}(y_i) - g_d^l(y_i)\|_\infty &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l \|y_i\|_\infty \\ &\stackrel{(i)}{\leq} L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l (\|X\|_{p,\infty} + \epsilon) = B_{X,\epsilon}^{1,\infty} \frac{\delta_l}{M_l^{1,\infty}}, \end{aligned}$$

where (i) is from Lemma 39.

Similarly, we have:

$$\begin{aligned} \|g_d^{l-1}(x_i^+) - g_d^l(x_i^+)\|_\infty &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l \|x_i^+\|_\infty \\ &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l \|X^+\|_{p,\infty} = B_{X^+}^{1,\infty} \frac{\delta_l}{M_l^{1,\infty}}, \end{aligned}$$

$$\begin{aligned} \|g_d^{l-1}(x_i^-) - g_d^l(x_i^-)\|_\infty &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l \|x_i^-\|_\infty \\ &\leq L^{d-1} \frac{\prod_{j=1}^d M_j^{1,\infty}}{M_l^{1,\infty}} \delta_l \|X^-\|_{p,\infty} = B_{X^-}^{1,\infty} \frac{\delta_l}{M_l^{1,\infty}}. \end{aligned}$$

Combine the above with (A.39):

$$\begin{aligned}
 |h(z_i) - h^c(z_i)| &\leq (B'_{X^+} + B'_{X^-}) \left(\|g_d^0(y_i) - g_d^1(y_i)\|_\infty + \cdots + \|g_d^{d-1}(y_i) - g_d^d(y_i)\|_\infty \right) \\
 &\quad + B'_{X,\epsilon} \left(\|g_d^0(x_i^+) - g_d^1(x_i^+)\|_\infty + \cdots + \|g_d^{d-1}(x_i^+) - g_d^d(x_i^+)\|_\infty \right) \\
 &\quad + B'_{X,\epsilon} \left(\|g_d^0(x_i^-) - g_d^1(x_i^-)\|_\infty + \cdots + \|g_d^{d-1}(x_i^-) - g_d^d(x_i^-)\|_\infty \right) \\
 &= (B'_{X^+} + B'_{X^-}) B_{X,\epsilon}^{1,\infty} \sum_{l=1}^d \frac{\delta_l}{M_l^{1,\infty}} + B'_{X,\epsilon} B_{X^+}^{1,\infty} \sum_{l=1}^d \frac{\delta_l}{M_l^{1,\infty}} + B'_{X,\epsilon} B_{X^-}^{1,\infty} \sum_{l=1}^d \frac{\delta_l}{M_l^{1,\infty}} \\
 &= \left[B_{X,\epsilon}^{1,\infty} (B'_{X^+} + B'_{X^-}) + B'_{X,\epsilon} (B_{X^+}^{1,\infty} + B_{X^-}^{1,\infty}) \right] \sum_{l=1}^d \frac{\delta_l}{M_l^{1,\infty}} \triangleq K_1 \sum_{l=1}^d \frac{\delta_l}{M_l^{1,\infty}}.
 \end{aligned}$$

So $\|h - h^c\|_2 = \sqrt{M} \max_{1 \leq i \leq M} |h(z_i) - h^c(z_i)| \leq \sqrt{M} \sum_{l=1}^d \frac{K_1 \delta_l}{M_l^{1,\infty}}$. Let $\delta_l = \frac{M_l^{1,\infty} \delta}{d K_1 \sqrt{M}}$, then:

$$\|h - h^c\|_2 \leq \sqrt{M} \sum_{l=1}^d \frac{K_1}{M_l^{1,\infty}} \cdot \frac{M_l^{1,\infty} \delta}{d K_1 \sqrt{M}} \leq \delta,$$

which means that: $\forall h \in \mathcal{H}, \exists h^c \in \mathcal{H}^c$ s.t. $\|h - h^c\|_2 \leq \delta$, so we have $\sup_{h \in \mathcal{H}} \inf_{h^c \in \mathcal{H}^c} \|h - h^c\|_2 \leq \delta$

when choosing $\delta_l = \frac{M_l^{1,\infty} \delta}{d K_1 \sqrt{M}}$.

So \mathcal{H}^c is a δ -covering of \mathcal{H} , and $\mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2) \leq |\mathcal{H}^c| = \prod_{l=1}^d |\mathcal{C}_l^{1,\infty}|$. By Lemma 18 we know that: $|\mathcal{C}_l^{1,\infty}| = \mathcal{N}\left(\frac{M_l^{1,\infty} \delta}{d K_1 \sqrt{M}}; S_l^{1,\infty}, \|\cdot\|_{1,\infty}\right) \leq \left(1 + \frac{2d K_1 \sqrt{M}}{\delta}\right)^{h_l \times h_{l-1}}$.

This means:

$$\mathcal{N}(\delta; \mathcal{H}, \|\cdot\|_2) \leq |\mathcal{H}^c| = \prod_{l=1}^d |\mathcal{C}_l^{1,\infty}| \leq \left(1 + \frac{2d K_1 \sqrt{M}}{\delta}\right)^{\sum_{l=1}^d h_l \cdot h_{l-1}}. \quad (\text{A.42})$$

So we can conclude that:

$$\begin{aligned}
 \mathcal{J}(0; D) &= \int_0^D \sqrt{\ln \mathcal{N}(u; \mathcal{H}, \|\cdot\|_2)} du \stackrel{(i)}{\leq} \int_0^D \sqrt{\left(\sum_{l=1}^d h_l \cdot h_{l-1}\right) \ln \left(1 + \frac{2d K_1 \sqrt{M}}{u}\right)} du \\
 &\stackrel{(ii)}{\leq} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1}} \int_0^D \sqrt{\frac{2d K_1 \sqrt{M}}{u}} du = \sqrt{2d K_1 \sum_{l=1}^d h_l \cdot h_{l-1}} \sqrt{M} \int_0^D u^{-\frac{1}{2}} du \\
 &= 2 \sqrt{2d K_1 D \sum_{l=1}^d h_l \cdot h_{l-1}} \sqrt{M} \stackrel{(iii)}{\leq} 2\sqrt{2} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{d K_0 K_1} \sqrt{M}},
 \end{aligned}$$

where (i) is from (A.42); (ii) comes from the fact that $\ln(1+x) \leq x, \forall x \geq 0$ and (iii) comes from (A.38).

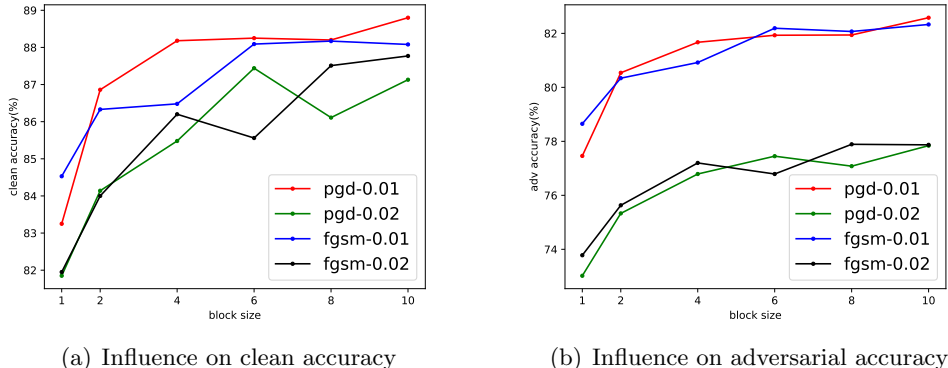


Figure B.1: The effect of block size on the accuracy. In the figure, we show the clean accuracy and the adversarial accuracy of the mean classifier under PGD and FGSM attack with $\epsilon = 0.01$ and $\epsilon = 0.02$. The block size is chosen from $\{1, 2, 4, 6, 8, 10\}$. **(a) The influence on the clean accuracy; (b) The influence on the adversarial accuracy.**

Since we shows that $\mathcal{R}_{\mathcal{S}}(\mathcal{H}) \leq 2\mathbb{E} \left[\sup_{\substack{\gamma, \gamma' \in \mathcal{H}(\mathcal{S}) \\ \|\gamma' - \gamma\|_2 \leq \delta}} (X_{\gamma} - X_{\gamma'}) \right] + 32\mathcal{J}(\delta/4; D)$ before, take $\delta \rightarrow 0^+$, we have:

$$\mathcal{R}_{\mathcal{S}}(\mathcal{H}) \leq 32\mathcal{J}(0; D) \leq 64\sqrt{2} \sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{dK_0K_1} \sqrt{M}} = O\left(\sqrt{\sum_{l=1}^d h_l \cdot h_{l-1} \sqrt{dK_0K_1} \sqrt{M}}\right).$$

■

Appendix B. Extra Experimental Results

In this section, we present our experimental results for **CIFAR-100**. The basic settings are the same as §6.

B.1 Improvement from Regularizer

Table B.1 shows that the F -norm regularizer helps to improve the adversarial robustness of the model, which agrees with our Theorem 24.

ATTACK	ϵ	TYPE	λ				
			0	0.002	0.005	0.01	0.02
PGD	0.01	CLEAN	84.07	85.65	85.48	85.52	58.53
		ADV	78.49	79.55	79.39	79.41	79.44
	0.02	CLEAN	81.87	81.91	81.93	82.08	81.97
		ADV	73.01	73.05	73.07	73.16	73.07
FGSM	0.01	CLEAN	84.61	84.53	84.95	84.58	84.55
		ADV	78.70	78.64	79.18	78.70	78.63
	0.02	CLEAN	80.47	80.48	80.39	80.44	82.09
		ADV	72.45	72.45	72.36	72.41	73.87

Table B.1: Results of experiments on the regularizer on data set **CIFAR-100**. In this table, we list the clean accuracy (CLEAN) and adversarial accuracy (ADV) of the mean classifier under the PGD and FGSM attack with $\epsilon = 0.01$ and $\epsilon = 0.02$. λ is chosen from $\{0, 0.002, 0.005, 0.01, 0.02\}$, and $\lambda = 0$ indicates no regularizer.

B.2 Effect of Block Size

Figure B.1 records the influence of block size on the clean accuracy and the adversarial accuracy of the model, from which we can see that a larger block size will yield better adversarial accuracy.

References

- Pranjal Awasthi, Natalie Frank, and Mehryar Mohri. Adversarial learning guarantees for linear hypotheses and neural networks. In *ICML*, volume 119, pages 431–441, 2020.
- Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *SP*, pages 39–57, 2017.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119, pages 1597–1607, 2020.
- Qingyi Gao and Xiao Wang. Theoretical investigation of generalization bounds for adversarial learning of deep neural networks. *Journal of Statistical Theory and Practice*, 15(2):51, 2021.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020.

- Chih-Hui Ho and Nuno Vasconcelos. Contrastive learning with adversarial examples. In NeurIPS, 2020.
- Ziyu Jiang, Tianlong Chen, Ting Chen, and Zhangyang Wang. Robust pre-training by adversarial contrastive learning. In NeurIPS, 2020.
- Minseon Kim, Jihoon Tack, and Sung Ju Hwang. Adversarial self-supervised contrastive learning. In NeurIPS, 2020.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In ICLR, 2017.
- Michel Ledoux and Michel Talagrand. Probability in Banach Spaces: isoperimetry and processes. Springer Science & Business Media, 2013.
- Boqi Li and Weiwei Liu. Wat: Improve the worst-class robustness in adversarial training, 2023.
- Xiyuan Li, Xin Zou, and Weiwei Liu. Defending against adversarial attacks via neural dynamic system. In NeurIPS, 2022.
- Xinsong Ma, Zekai Wang, and Weiwei Liu. On the tradeoff between robustness and fairness. In NeurIPS, 2022.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In ICLR, 2018.
- Vitaly Maiorov and Allan Pinkus. Lower bounds for approximation by MLP neural networks. Neurocomputing, 25(1-3):81–91, 1999.
- Chengzhi Mao, Amogh Gupta, Vikram Nitin, Baishakhi Ray, Shuran Song, Junfeng Yang, and Carl Vondrick. Multitask learning strengthens adversarial robustness. In ECCV, pages 158–174, 2020.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. Foundations of Machine Learning. Adaptive computation and machine learning. MIT Press, 2012.
- Omar Montasser, Steve Hanneke, and Nathan Srebro. VC classes are adversarially robustly learnable, but only improperly. In COLT, pages 2512–2530, 2019.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: A simple and accurate method to fool deep neural networks. In CVPR, pages 2574–2582, 2016.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In COLT, volume 40, pages 1376–1401, 2015.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In ECCV, volume 9910, pages 69–84, 2016.

- Kento Nozawa, Pascal Germain, and Benjamin Guedj. Pac-bayesian contrastive unsupervised representation learning. In UAI, volume 124, pages 21–30, 2020.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In ICML, volume 97, pages 5628–5637, 2019.
- Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In NeurIPS, pages 5019–5031, 2018.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In ICLR, 2014.
- Martin J. Wainwright. High-dimensional statistics: A non-asymptotic viewpoint, volume 48. Cambridge University Press, 2019.
- Zekai Wang and Weiwei Liu. Robustness verification for contrastive learning. In ICML, pages 22865–22883, 2022.
- Jingyuan Xu and Weiwei Liu. On robust multiclass learnability. In NeurIPS, 2022.
- Dong Yin, Kannan Ramchandran, and Peter L. Bartlett. Rademacher complexity for adversarially robust generalization. In ICML, volume 97, pages 7085–7094, 2019.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In ICLR, 2017.
- Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically principled trade-off between robustness and accuracy. In ICML, volume 97, pages 7472–7482, 2019.