

OUTLIER ROBUST ADVERSARIAL TRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

Supervised learning models are challenged by the intrinsic complexities of training data such as outliers and minority subpopulations and intentional attacks at inference time with adversarial samples. While traditional robust learning methods and the recent adversarial training approaches are designed to handle each of the two challenges, to date, no work has been done to develop models that are robust with regard to the low-quality training data and the potential adversarial attack at inference time simultaneously. It is for this reason that we introduce Outlier Robust Adversarial Training (ORAT) in this work. ORAT is based on a bi-level optimization formulation of adversarial training with a robust rank-based loss function. Theoretically, we show that the learning objective of ORAT satisfies the \mathcal{H} -consistency (Awasthi et al., 2021) in binary classification, which establishes it as a proper surrogate to adversarial 0/1 loss. Furthermore, we analyze its generalization ability and provide uniform convergence rates in high probability. ORAT can be optimized with a simple algorithm. Experimental evaluations on three benchmark datasets demonstrate the effectiveness and robustness of ORAT in handling outliers and adversarial attacks.

1 INTRODUCTION

In supervised learning, we obtain a parametric model $f_\theta(\mathbf{x})$ to predict the label (discrete or continuous) y from an input \mathbf{x} . The optimal value of the parameter θ is obtained with a set of labeled training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, by minimizing a loss function. However, supervised learning algorithms are challenged by two types of data degradation. First, the quality of training data is affected by erroneous samples due to mistakes in data collection or labeling (often known as the *outliers*) or isolated sub-populations of actual samples (Sukhbaatar et al., 2015). In addition, at inference time, an input data point \mathbf{x} could be intentionally modified to create an adversarial sample that misleads $f_\theta(\mathbf{x})$ to make a wrong prediction (Goodfellow et al., 2015).

To date, the resilience of supervised learning models with regard to unintentional non-ideal training data or intentional adversarial attacks has been studied separately in machine learning. The former is the topic of robust learning methods (Wang et al., 2018; Hu et al., 2020; Zhai et al., 2021), and the latter is addressed with adversarial training (AT) (Madry et al., 2018; Wang et al., 2020; Zhang et al., 2021). However, in practice, the two issues can occur in tandem. This has been noticed in Sanyal et al. (2021) and Zhu et al. (2021). They have demonstrated that the outlier problem (i.e., label noise) exists in AT and found the model performance degrades with the increase in noise level because outliers hurt the quality of training data. Only Zhu et al. (2021) provide a heuristic algorithm based on a label correction strategy to correct the noise label. However, their approach may introduce more extra noisy labels due to the imperfect classifier and cannot handle outliers with true labels that do not belong to the existing label list.

To reduce the influence of the outliers in AT, one may think of adapting the self-learning (Han et al., 2019) approach (i.e., a robust learning algorithm) to remove examples that are most likely outliers (e.g., examples with larger loss) from data before training and then conducting AT on the cleaner set. However, there are two drawbacks to this simple scheme. First, it is not an end-to-end approach, which will increase the training cost. Second, it is hard to remove outliers precisely, and an imperfect strategy may drop examples of clean data points. This may hurt the final model performance in the AT phase. An alternative solution is that combine robust learning and AT by using robust losses (e.g., Huber loss (Hastie et al., 2009), symmetric cross entropy loss (Wang et al., 2019), etc.) in AT. However, most existing robust losses cannot eliminate the influence of outliers. In addition, constructing the theoretical guarantee is a challenge for using a robust loss in AT, especially for

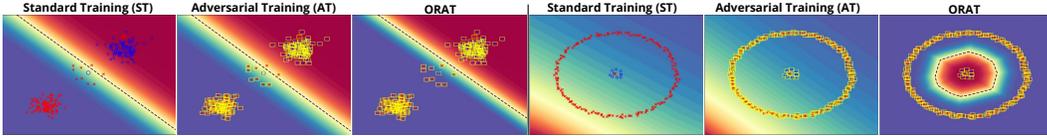


Figure 1: Illustrative examples of standard training (ST), adversarial training (AT), and ORAT for binary classification on a balanced but multi-modal synthetic dataset (left panel) with two outliers and an imbalanced synthetic dataset (right panel) with one outlier. outliers in the blue and red classes are shown in \times and \circ , respectively. The yellow squares around data samples represent the samples are perturbed within a l_∞ ball. The dash lines are the decision boundaries. Note that there are no decision boundaries in the right figures when using ST and AT because these two training strategies cannot obtain a classifier that can separate two classes.

\mathcal{H} -calibration and \mathcal{H} -consistency properties (Awasthi et al., 2021; Steinwart, 2007) of the designed adversarial loss with respect to the adversarial 0/1 loss in classification.

In this work, we introduce Outlier Robust Adversarial Training (ORAT) to combine robust learning and adversarial training. Specifically, we develop an effective adversarial training algorithm for a rank-based learning objective that can exclude the influence of outliers from the training procedure. Figure 1 shows several illustrative examples. The learning objective in ORAT lends itself to an efficient numerical algorithm based on gradient descent methods after reformulation that removes the explicit ranking operation. To verify whether the optimal minimizers of the ORAT loss are close to or exactly the optimal minimizers of the adversarial 0/1 loss, we show that the ORAT loss satisfies the \mathcal{H} -calibration and consistency properties for classification under some moderate conditions, which establishes it as a proper surrogate to adversarial 0/1 loss. The notion of consistency has been studied in Awasthi et al. (2021). However, the ORAT loss is the first adversarial surrogate loss proven to satisfy \mathcal{H} -consistency and evaluated on real-world datasets. It encourages the applicability of the \mathcal{H} -consistent adversarial surrogate loss in real tasks. We further provide a quantitative error bound of the generalization gap between the training and testing performance of ORAT. Experimental evaluations on three benchmark datasets demonstrate the effectiveness and robustness of ORAT in handling outliers and adversarial attacks. Our contributions can be summarized as follows:

- We present outlier robust adversarial training (ORAT), which can handle outliers in adversarial training jointly.
- We show ORAT loss satisfies \mathcal{H} -calibration and \mathcal{H} -consistency. To the best of our knowledge, the ORAT loss is the first \mathcal{H} -consistent adversarial surrogate loss that is evaluated on real-world datasets.
- We provide a detailed theoretical analysis on the generalization error of training with ORAT loss.

2 BACKGROUND

Robust Learning. Training accurate machine learning models in the presence of outliers is of great practical importance. To combat outliers, the traditional methods are designed based on label correction (Wang et al., 2018), loss correction (Han et al., 2020), and refined training strategies (Yu et al., 2019). However, they require an extra clean dataset or potentially expensive detection process to estimate the outliers. A recent work (Hu et al., 2020) proposed the average of ranked range (AoRR) loss, which can eliminate the influence of the outliers if their proportion in training data is known. Let \mathcal{X} denote the input feature space, $\mathbf{x} \in \mathcal{X}$ is a training sample, $y \in \mathcal{Y} = \{1, \dots, C\}$ is its associated label, and $C \geq 2$. $f_\theta(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^C$ is the logits output of the predictor, and $\ell : \mathbb{R}^C \times \mathcal{Y} \rightarrow \mathbb{R}$ is a loss function. $\ell_{[i]}(f_\theta(\mathbf{x}_j), y_j)_{j=1}^n$ represents the i -th largest loss among the training sample set. For two integers k and m , $0 \leq m < k \leq n$, the AoRR loss is defined as follows,

$$\min_{\theta} \frac{1}{k-m} \sum_{i=m+1}^k \ell_{[i]}(f_\theta(\mathbf{x}_j), y_j)_{j=1}^n. \quad (1)$$

The AoRR loss excludes training samples with the top m -largest loss value, as well as samples with small losses. This is to reduce the influence of the outliers (larger losses) and to enhance the effect of the minority subgroup of the data because the samples with small loss values are most likely from the majority subgroup in the training set.

Adversarial Training. Recent studies have shown some surprising vulnerabilities of advanced supervised learning models, especially those based on deep neural networks, to specially designed adversarial samples (Goodfellow et al., 2015; Carlini & Wagner, 2017). To mitigate this issue, adversarial training (AT) (Madry et al., 2018) is proposed as a training approach against adversarial attacks. Let $\mathcal{B}_\epsilon(\mathbf{x}) = \{\mathbf{x}' \in \mathcal{X} \mid \|\mathbf{x} - \mathbf{x}'\|_p \leq \epsilon\}$ be the closed ball of radius $\epsilon > 0$ centered at \mathbf{x} ,

where p is usually chosen as 1, 2, or ∞ . Given a training dataset $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ independently drawn from a distribution \mathcal{D} , where $\mathbf{x}_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$. The objective function of AT is then

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \left[\max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_{\epsilon}(\mathbf{x}_i)} \ell(f_{\theta}(\tilde{\mathbf{x}}_i), y_i) \right], \quad (2)$$

where $\tilde{\mathbf{x}}$ is the most extreme adversarial sample within the ϵ -ball centered at \mathbf{x} . To generate adversarial data, The original AT applies projected gradient descent (PGD) using a fixed number of iterations P as a stopping criterion, namely the PGD ^{P} algorithm, to approximately solve the inner maximization problem of the Eq.(2). Other methods can also generate adversarial data, e.g., the fast gradient signed method (FGSM) (Goodfellow et al., 2015) and the CW attack (Carlini & Wagner, 2017).

Robust Learning under Adversarial Attacks. Currently, several works have realized the impact of label noise on adversarial training. For example, in Sanyal et al. (2021), the authors identified label noise as one of the causes for adversarial vulnerability but no defense methods are proposed to solve this problem. The work Zhu et al. (2021) empirically studies the efficacy of AT for mitigating the effect of label noise in training data. However, their proposed annotator algorithm (RAA) is based on the label correction strategy, which may introduce more extra noisy labels due to the bottleneck of the selected classifier. Furthermore, both of these methods are only considered label noise problems instead of outliers, especially the error that comes from the sample itself. Therefore, an outlier robust adversarial training method is urgently needed to fill this gap. Several works (Augustin et al., 2020; Bitterwolf et al., 2020) connect adversarial robustness to out-of-distribution (OOD) problems. Note that the notion of outliers is different from OOD points. One assumption for the OOD problem is that the training and test data distributions are mismatched. However, we do not have this assumption in this work. More related works are discussed in Appendix A.

3 METHOD

In ORAT, we combine the rank-based AoRR loss (Eq.(1)) and adversarial training (Eq.(2)) into the same framework. Let $L(\{(\mathbf{x}_j, y_j)\}_{j=1}^n) := \{\ell(f_{\theta}(\mathbf{x}_1), y_1), \dots, \ell(f_{\theta}(\mathbf{x}_n), y_n)\}$ be the set of all individual losses on the training samples. For notational brevity, we drop the explicit dependence on the individual loss ℓ , the parametric model f_{θ} , but it should be clear that the set changes with the training dataset $\{(\mathbf{x}_j, y_j)\}_{j=1}^n$. In addition, we denote $\ell_{[i]}(\{(\mathbf{x}_j, y_j)\}_{j=1}^n)$ as the i -th largest individual loss after sorting the elements in set $L(\{(\mathbf{x}_j, y_j)\}_{j=1}^n)$ (ties can be broken in any consistent way). With these definitions, we define the learning objective of ORAT as a bi-level optimization problem (Gould et al., 2016) as

$$\min_{\theta} \frac{1}{k-m} \sum_{i=m+1}^k \ell(f_{\theta}(\tilde{\mathbf{x}}_{[i]}), y_{[i]}), \quad \text{s.t. } (\tilde{\mathbf{x}}_{[i]}, y_{[i]}) = \operatorname{argmax}_{\tilde{\mathbf{x}}_j \in \mathcal{B}_{\epsilon}(\mathbf{x}_j)} \ell_{[i]}(\{\tilde{\mathbf{x}}_j, y_j\}_{j=1}^n), \quad (3)$$

where k and m are two integers such that $0 \leq m < k \leq n$. The notations here require some further explanation. First, $\tilde{\mathbf{x}}_j$ is the result of adversarial perturbation of an original data point \mathbf{x}_j with a perturbation strength ϵ , and $y_{[i]}$ is the corresponding label of the original \mathbf{x}_j . So the maximization sub-problem of Eq.(3) reads as follows. For each original training sample (\mathbf{x}_j, y_j) , we find the extreme adversarial input $\tilde{\mathbf{x}}_j$ within the ϵ -ball centered at \mathbf{x}_j . Then we calculate the individual loss of the perturbed samples $\ell(f_{\theta}(\tilde{\mathbf{x}}_j), y_j)$ and sort the losses to find the adversarial sample and label $(\tilde{\mathbf{x}}_{[i]}, y_{[i]})$ corresponding to the top- i individual loss. The outer minimization problem of Eq.(3) is the average of ranked range of $(m, k]$ individual losses found with such a procedure. Note that ORAT contains the original AT as a special case with $k = n$ and $m = 0$.

The robustness to outliers and adversarial attacks of the ORAT (Eq.(3)) can be more clearly understood as follows. The overall method can be viewed as a sample selection method and can filter incorrect data samples according to the top- k and top- m individual losses. Specifically, the perturbed samples with small losses are most likely come from the clean data. Therefore, it ignores m samples with a large loss (i.e., top-1 to top- m losses) during the training. On the other hand, the perturbed samples with an extremely low loss value are most likely very easy to be learned in the training procedure. They usually come from the majority group in a dataset. On the contrary, the perturbed samples from the minority subgroup can be also viewed as hard samples, which are very hard to be learned. In this case, ignoring the bottom $n - k$ would reduce the influence of the majority subgroup of data, which further prevents the effect of the imbalance data and enhance the impact of the minority subgroup of the data. Figure 1 also demonstrates this influence.

Although ORAT is designed by combining AoRR and AT, this combination has not been explored in the existing literature. More importantly, we will demonstrate and validate both theoretically (Section 4) and experimentally (Section 5) that this combination makes sense. To the best of our knowledge, no robust losses are combined with AT that can handle both low-quality training data (with outliers) and adversarial attacks in inference. In addition, the ρ -margin loss (Awasthi et al., 2021), a generalization of the ramp loss, is proved to satisfy the calibration and consistency. Inspired by this loss, we naturally thought that the truncated robust loss could be used to design a new adversarial surrogate loss that satisfies \mathcal{H} -calibration and consistency. Therefore, we choose AoRR loss to create ORAT because it is a well-defined truncated loss. The ranking operation in ORAT is the main obstacle to using Eq.(3) as a learning objective in an efficient way. However, we can substitute the ranking operation by introducing two auxiliary variables λ and $\hat{\lambda}$ and use the equivalent form of ORAT in the following result.

Theorem 1. Denote $[a]_+ = \max\{0, a\}$ as the hinge function. Eq.(3) is equivalent to

$$\frac{1}{k-m} \min_{\theta, \lambda} \max_{\hat{\lambda}} \sum_{i=1}^n \hat{\mathcal{L}}(f_{\theta}, \lambda, \hat{\lambda}) := \left[\frac{k-m}{n} \lambda + \frac{n-m}{n} \hat{\lambda} - [\hat{\lambda} - [f_{\theta}(\tilde{\mathbf{x}}_i), y_i] - \lambda]_+ \right], \quad (4)$$

s.t. $\tilde{\mathbf{x}}_i = \arg \max_{\tilde{\mathbf{x}} \in \mathcal{B}_{\epsilon}(\mathbf{x}_i)} \ell(f_{\theta}(\tilde{\mathbf{x}}), y_i)$.

Furthermore, $\hat{\lambda} > \lambda$, when the optimal solution is achieved.

The proof of Theorem 1 can be found in Appendix C.1. Using Theorem 1, we can develop a learning algorithm based on stochastic (mini-batch) gradient descent to optimize Eq.(3). Specifically, with initial choice for the values of $\theta^{(0)}$, $\lambda^{(0)}$, and $\hat{\lambda}^{(0)}$, at the t -th iteration, a mini-batch set \mathcal{S}_b of training samples is chosen uniformly at random from the training set and used to estimate the (sub)gradient of the objective. Since the optimal solution of parameters of θ , λ , and $\hat{\lambda}$ do not depend on the factor of $\frac{1}{k-m}$, we can replace it to $\frac{1}{|\mathcal{S}_b|}$ for mini-batch optimization, where $|\mathcal{S}_b|$ is the size of \mathcal{S}_b . Following the original AT (Madry et al., 2018), we use the projected gradient descent approach to approximately solve the constraint. $\Pi_{\mathcal{B}_{\epsilon}(\mathbf{x}_i)}(\cdot)$ is the projection function that projects the adversarial data back into the ϵ -ball centered at \mathbf{x}_i if necessary. $\partial_{\theta} \hat{\mathcal{L}}$, $\partial_{\lambda} \hat{\mathcal{L}}$, and $\partial_{\hat{\lambda}} \hat{\mathcal{L}}$ are the (sub)gradients of $\hat{\mathcal{L}}$ with respect to θ , λ , and $\hat{\lambda}$, respectively. Their explicit forms can be found in Appendix B. The pseudo-code of optimizing ORAT is described in Algorithm 1. Note that our optimization process is similar to the traditional adversarial training algorithms (Madry et al., 2018; Zhang et al., 2021; Liu et al., 2021) with one additional minimization with respect to λ and one additional maximization for $\hat{\lambda}$. Therefore, our algorithm has the same time complexity to original AT. We will show that our algorithm can converge in experiments.

Algorithm 1: Outlier Robust Adversarial Training

Input: A training dataset \mathcal{S} of size n

Output: A robust model with parameters θ^*

Initialization: $\theta^{(0)}$, $\lambda^{(0)}$, $\hat{\lambda}^{(0)}$, $t=0$, η , k , m , ϵ , α , and P

for $e = 1$ to num_epoch **do**

for $b = 1$ to num_batch **do**

 Sample a mini-batch $\mathcal{S}_b = \{(\mathbf{x}_i, y_i)\}_{i=1}^{|\mathcal{S}_b|}$ from \mathcal{S}

for $i = 1$ to batch_size **do**

$\tilde{\mathbf{x}}_i \leftarrow \mathbf{x}_i$

while $P > 0$ **do**

$\tilde{\mathbf{x}}_i = \Pi_{\mathcal{B}_{\epsilon}(\mathbf{x}_i)}(\tilde{\mathbf{x}}_i + \alpha \text{sign}(\nabla_{\tilde{\mathbf{x}}_i} \ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i)))$

$P \leftarrow P - 1$

end

end

$\theta^{(t+1)} \leftarrow \theta^{(t)} - \frac{\eta}{|\mathcal{S}_b|} \sum_{i \in \mathcal{S}_b} \partial_{\theta} \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)})$

$\lambda^{(t+1)} \leftarrow \lambda^{(t)} - \frac{\eta}{|\mathcal{S}_b|} \sum_{i \in \mathcal{S}_b} \partial_{\lambda} \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)})$

$\hat{\lambda}^{(t+1)} \leftarrow \hat{\lambda}^{(t)} + \frac{\eta}{|\mathcal{S}_b|} \sum_{i \in \mathcal{S}_b} \partial_{\hat{\lambda}} \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)})$

$t \leftarrow t + 1$

end

end

4 ANALYSIS

In this section, we prove that ORAT is a \mathcal{H} -consistent adversarial surrogate loss and study its generalization property. Subsequently, we focus on the case of binary classification ($\mathcal{Y} = \{\pm 1\}$) with margin-based loss function, i.e., $\ell(yf(\mathbf{x})) = \ell(f(\mathbf{x}), y) = \ell(f, \mathbf{x}, y)$, where we omit model parameter θ for simplicity. We first introduce several notations. Let \mathcal{H} be a Hypothesis function set from \mathbb{R}^d to \mathbb{R} . We say \mathcal{H} is symmetric, if for any $f \in \mathcal{H}$, $-f$ is also in \mathcal{H} . The 0/1 risk of a classifier $f \in \mathcal{H}$ is $\mathcal{R}_{\ell_0} = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\ell_0(f(\mathbf{x}), y)]$, where $\ell_0(f(\mathbf{x}), y) = \mathbb{1}_{yf(\mathbf{x}) \leq 0}$ is the 0/1 loss. Denote the adversarial 0/1 risk as $\mathcal{R}_{\tilde{\ell}_0}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\tilde{\ell}_0(f(\mathbf{x}), y)]$, where $\tilde{\ell}_0(f(\mathbf{x}), y) = \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_{\epsilon}(\mathbf{x})} \mathbb{1}_{yf(\tilde{\mathbf{x}}) \leq 0}$ is the adversarial 0/1 loss. We also define the ℓ_s -risk of f for a surrogate loss $\ell_s(f(\mathbf{x}), y)$ as $\mathcal{R}_{\ell_s}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\ell_s(f(\mathbf{x}), y)]$ and the minimal (ℓ_s, \mathcal{H})-risk, which is defined by $\mathcal{R}_{\ell_s, \mathcal{H}}^* = \inf_{f \in \mathcal{H}} \mathcal{R}_{\ell_s}(f)$.

4.1 CLASSIFICATION CALIBRATION AND CONSISTENCY

Designing a robust algorithm entails using an appropriate surrogate loss to the standard 0/1 loss since 0/1 loss is very hard to optimize for most hypothesis sets (Awasthi et al., 2021; Bao et al., 2020). We first provide a definition of \mathcal{H} -consistency.

Definition 1. (*\mathcal{H} -Consistency*). (Awasthi et al., 2021) Given a hypothesis set \mathcal{H} , we say that a loss function ℓ_1 is \mathcal{H} -consistent w.r.t. a loss function ℓ_2 , if, for all probability distributions and sequences of $\{f_n\}_{n \in \mathbb{N}} \subset \mathcal{H}$, there holds $\mathcal{R}_{\ell_1}(f_n) - \mathcal{R}_{\ell_1, \mathcal{H}}^* \xrightarrow{n \rightarrow +\infty} 0 \Rightarrow \mathcal{R}_{\ell_2}(f_n) - \mathcal{R}_{\ell_2, \mathcal{H}}^* \xrightarrow{n \rightarrow +\infty} 0$.

\mathcal{H} -consistency guarantees that the optimal minimizers of the surrogate adversarial loss are equal to or near the optimal minimizers of the 0/1 adversarial loss on a restricted hypothesis set \mathcal{H} .

For a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ with random variables X and Y . Let $\eta := \Pr(Y = 1 | X = \mathbf{x}) \in [0, 1]$, for any $\mathbf{x} \in \mathcal{X}$. Using conditional expectation, we can rewrite $\mathcal{R}_{\ell_s}(f)$ as $\mathcal{R}_{\ell_s}(f) = \mathbb{E}_X[\mathcal{C}_{\ell_s}(f, \mathbf{x}, \eta)]$, where $\mathcal{C}_{\ell_s}(f, \mathbf{x}, \eta) := \eta \ell_s(f, \mathbf{x}, 1) + (1 - \eta) \ell_s(f, \mathbf{x}, -1)$, $\forall \mathbf{x} \in \mathcal{X}$. Furthermore, the minimal inner ℓ_s -risk on \mathcal{H} is denoted by $\mathcal{C}_{\ell_s, \mathcal{H}}^*(\mathbf{x}, \eta) := \inf_{f \in \mathcal{H}} \mathcal{C}_{\ell_s}(f, \mathbf{x}, \eta)$. We now define \mathcal{H} -calibration.

Definition 2. (*\mathcal{H} -Calibration*). (Awasthi et al., 2021) Given a hypothesis set \mathcal{H} , we say that a loss function ℓ_1 is \mathcal{H} -calibrated with respect to a loss function ℓ_2 if, for any $\tau > 0$, $\eta \in [0, 1]$, and $\mathbf{x} \in \mathcal{X}$, there exists $\delta > 0$ such that, for all $f \in \mathcal{H}$, $\mathcal{C}_{\ell_1}(f, \mathbf{x}, \eta) < \mathcal{C}_{\ell_1, \mathcal{H}}^*(\mathbf{x}, \eta) + \delta \Rightarrow \mathcal{C}_{\ell_2}(f, \mathbf{x}, \eta) < \mathcal{C}_{\ell_2, \mathcal{H}}^*(\mathbf{x}, \eta) + \tau$.

As shown in Steinwart (2007), if ℓ_1 is \mathcal{H} -calibrated with respect to ℓ_2 , then \mathcal{H} -consistency holds for any probability distribution verifying the additional condition of minimizability.

Without considering the adversarial perturbations and omitting the parameter θ , the population version of the objective function of Eq.(4) can be formulated as follows:

$$\min_{\lambda} \max_{\hat{\lambda}} \frac{1}{k-m} \sum_{i=1}^n \hat{\mathcal{L}}(f, \lambda, \hat{\lambda}) \xrightarrow[\nu \rightarrow \infty]{\frac{k-m}{n} \rightarrow \nu, \frac{m}{n} \rightarrow \mu} \frac{1}{\nu} \min_{\lambda} \max_{\hat{\lambda}} \mathbb{E} \left[\hat{\lambda} - [\hat{\lambda} - [\ell(Yf(X)) - \lambda]_+]_+ \right] + \nu \lambda - \mu \hat{\lambda}.$$

Throughout the paper, we assume that $\mu > 0$ since if $\mu = 0$ then it will lead to $\hat{\lambda} = \infty$. As such, the population version of our ORAT loss is given by $(f_0^*, \lambda^*, \hat{\lambda}^*) = \arg \inf_{f, \lambda} \sup_{\hat{\lambda}} \left\{ \mathbb{E} \left[\hat{\lambda} - [\hat{\lambda} - [\ell(Yf(X)) - \lambda]_+]_+ \right] + \nu \lambda - \mu \hat{\lambda} \right\}$. It is difficult to directly work with the optima f_0^* since the above problem is a non-convex minmax problem and the standard minmax theorem does not apply here. Instead, we assume the existence of λ^* and $\hat{\lambda}^*$ and work with the minimizer $f^* = \arg \inf_f \mathcal{L}(f, \lambda^*, \hat{\lambda}^*)$ where $\mathcal{L}(f, \lambda^*, \hat{\lambda}^*) := \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf(X)) - \lambda^*]_+]_+ \right] + \nu \lambda^* - \mu \hat{\lambda}^*$. Since the term $\nu \lambda^* - \mu \hat{\lambda}^*$ does not depend on f , we have $f^* = \arg \inf_f \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf(X)) - \lambda^*]_+]_+ \right]$.

From the above observations, we denote by ϕ_{ORAT} and $\tilde{\phi}_{\text{ORAT}}$ for ORAT loss without and with the adversarial perturbation, respectively, as follows:

$$\phi_{\text{ORAT}}(t) = \hat{\lambda}^* - [\hat{\lambda}^* - [\ell(t) - \lambda^*]_+]_+, \quad \tilde{\phi}_{\text{ORAT}}(f, \mathbf{x}, y) = \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \phi_{\text{ORAT}}(yf(\tilde{\mathbf{x}})). \quad (5)$$

We can then obtain the following theorem. Its proof can be found in Appendix C.2.

Theorem 2. Let \mathcal{H} be a symmetric hypothesis set consisting of the family of all measurable functions \mathcal{H}_{all} , suppose $\nu > \min\{\hat{\lambda}^*, \mathcal{R}_{\ell, \mathcal{H}}^*\}$, $0 \leq \lambda^* < \hat{\lambda}^*$, λ^* and $\hat{\lambda}^*$ are bounded, and ℓ is a non-negative, continuous, and non-increasing margin-based loss.

(i) Then $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$.

(ii) Furthermore, $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -consistent with respect to $\tilde{\ell}_0$ for all distributions \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ that satisfy: $\mathcal{R}_{\tilde{\ell}_0, \mathcal{H}}^* = 0$ and there exists $f^* \in \mathcal{H}$ such that $\mathcal{R}_{\tilde{\phi}_{\text{ORAT}}}(f^*) = \mathcal{R}_{\phi_{\text{ORAT}}, \mathcal{H}_{\text{all}}}^* < +\infty$.

\mathcal{H} can be linear models or deep neural networks. The commonly used hinge loss, logistic loss, and cross-entropy loss all satisfy the conditions of ℓ (See Appendix C.3 for details). Furthermore, according to Theorem 1 and assuming $\ell \geq 0$, we have $0 \leq \lambda^* < \hat{\lambda}^*$. For ν , it should be larger than the smaller value among $\hat{\lambda}^*$ and minimal (ℓ, \mathcal{H}) -risk, which means k and m should be as far

away from each other as possible. We call $\mathcal{R}_{\tilde{\ell}_0, \mathcal{H}}^* = 0$ as the realizability condition. Therefore, we can conclude that ORAT loss satisfies \mathcal{H} -consistency. Our proof is mainly inspired by [Awasthi et al. \(2021\)](#). However, our analysis for the $\tilde{\phi}_{\text{ORAT}}$ adversarial surrogate loss is more involved since it is a composition function based on ℓ for which we need to consider the different conditions of λ^* and $\hat{\lambda}^*$.

4.2 GENERALIZATION ERROR

In this subsection, we present the generalization error bound for the proposed ORAT loss. Define the adversarial surrogate loss $\tilde{\ell}(yf(\mathbf{x})) = \max_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \ell(yf(\tilde{\mathbf{x}}))$ and the composite function class $\tilde{\ell}_{\mathcal{H}} := \{\tilde{\ell}(yf(\mathbf{x})) : f \in \mathcal{H}\}$. The generalization error studies the discrepancy between the empirical adversarial risk $\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S})$ defined on the training data and its population risk $\mathcal{R}_{\tilde{\ell}}(f)$ measuring the performance on the test data, where $\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) = \inf_{\lambda} \sup_{\hat{\lambda}} \frac{1}{k-m} \sum_{i=1}^n \hat{\mathcal{L}}(f, \lambda, \hat{\lambda})$. First, we need the Rademacher complexity ([Bartlett & Mendelson, 2002](#)) which is defined as follows.

Definition 3 (Rademacher complexity). *For any function class \mathcal{H} , given a dataset $\{\mathbf{x}_i\}_{i=1}^n$, the empirical Rademacher complexity is defined as $\mathfrak{R}_n(\mathcal{H}) = \frac{1}{n} \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{H}} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right]$, where $\sigma_1, \dots, \sigma_n$ are i.i.d. Rademacher random variables with $\mathbb{P}[\sigma_i = 1] = \mathbb{P}[\sigma_i = -1] = 1/2$.*

With these preparations, we can get the following theorem.

Theorem 3. *Suppose that the range of $\ell(f(\mathbf{x}), y)$ is $[0, M]$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of an i.i.d. training dataset of size n , the following holds for all $f \in \mathcal{H}$,*

$$\mathcal{R}_{\tilde{\ell}}(f) - \mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) \leq \frac{2}{\nu} \left(2\mathfrak{R}_n(\tilde{\ell}_{\mathcal{H}}) + \frac{M(2\sqrt{2} + 3\sqrt{\log(2/\delta)})}{\sqrt{2n}} \right).$$

The proof of Theorem 3 is by rewriting the empirical risk and population risk with optimal λ^* and $\hat{\lambda}^*$, and then analytically bound the Rademacher complexity of λ^* and $\hat{\lambda}^*$ by utilizing the boundness condition. See Appendix C.4 for details. Theorem 3 characterizes uniform convergence between the training and testing on ORAT given hypothesis set \mathcal{H} . The generalization error depends on the limit of ranked range $\frac{k-m}{n} \rightarrow \nu$ as well as the Rademacher complexity of the adversarial loss function class. It is worth noting the i.i.d. assumption over the sample is restrictive but required for applying the symmetrization trick in the proof, and how to relax this assumption to a more realistic assumption, for example, $n - m$ i.i.d. inliers with m outliers ([Laforgue et al., 2021](#)), is future work for us. Nevertheless, Theorem 3 still highlights the generalization ability of our ORAT objective with respect to the ranked range and adversarial loss in this ideal setting. We provide hypothesis set examples of linear classifiers and neural networks in Appendix C.5 for further characterizing the Rademacher complexity $\mathfrak{R}_n(\tilde{\ell}_{\mathcal{H}})$.

Remark 1. *Theorem 3 together with Theorem 2 indicates that learning with ORAT asymptotically converges to zero on the adversarial 0/1 risk. Let $f^{**} = \arg \inf_{f \in \mathcal{H}} \mathcal{R}_{\tilde{\ell}}(f)$ and $f_S^{**} = \arg \inf_{f \in \mathcal{H}} \mathcal{R}_{\tilde{\ell}}(f; \mathcal{S})$. By standard error decomposition, $\mathcal{R}_{\tilde{\ell}}(f_S^{**}) - \mathcal{R}_{\tilde{\ell}}(f^{**}) = (\mathcal{R}_{\tilde{\ell}}(f_S^{**}) - \mathcal{R}_{\tilde{\ell}}(f_S^{**}; \mathcal{S})) + (\mathcal{R}_{\tilde{\ell}}(f_S^{**}; \mathcal{S}) - \mathcal{R}_{\tilde{\ell}}(f^{**}; \mathcal{S})) + (\mathcal{R}_{\tilde{\ell}}(f^{**}; \mathcal{S}) - \mathcal{R}_{\tilde{\ell}}(f^{**}))$. The first term converges to 0 as n goes to ∞ and $\mathfrak{R}_n(\tilde{\ell}_{\mathcal{H}})$ is bounded. The second term is always non-positive as f_S^{**} minimizes $\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S})$. The last term is bounded by $\mathcal{O}(1/\sqrt{n})$ with high probability by Hoeffding's inequality ([Hoeffding, 1994](#)). Therefore, $\mathcal{R}_{\tilde{\ell}}(f_S^{**}) - \mathcal{R}_{\tilde{\ell}}(f^{**}) \rightarrow 0$ as $n \rightarrow \infty$. Combined with Theorem 2 (ii) and Definition 1, it shows that $\mathcal{R}_{\tilde{\ell}_0}(f_S^{**}) - \mathcal{R}_{\tilde{\ell}_0}(f^{**}) \rightarrow 0$ as $n \rightarrow \infty$.*

5 EXPERIMENTS

We evaluate the performance of the proposed ORAT with numerical experiments. Due to the limit of the space, we present the most significant information and results of our experiments. More detailed information, additional results, and code are in Appendix D, E, and supplementary files, respectively.

5.1 EXPERIMENTAL SETTINGS

Datasets, Networks, and Baselines. Our experiments are based on three popular datasets, namely MNIST ([LeCun et al., 1998](#)), CIFAR-10, and CIFAR-100 ([Krizhevsky et al., 2009](#)). We follow the same training/testing splitting from the original datasets. The pixel values of all images are normalized into the range of $[0, 1]$. We adopt LeNet ([LeCun et al., 1998](#)), Small-CNN ([Wang et al., 2020](#)), and ResNet-18 ([He et al., 2016](#)) for MNIST, CIFAR-10, and CIFAR-100, respectively. Settings

of networks are in Appendix D.2. We compare ORAT with standard training (ST), adversarial training (AT). In addition, as there are no dedicated adversarial training methods for handling outliers, we compare three instance-reweighted AT methods GAIRAT (Zhang et al., 2021), MAIL (Liu et al., 2021), and WMMR (Zeng et al., 2021a), because they are the most recent works with the best performance against adversarial attacks. We also compare RAA (Zhu et al., 2021) since it considers label correction in training. To support our claims about the drawbacks of using a self-learning strategy to remove outliers before adversarial training in Section 1, we use the AoRR approach to remove examples with larger losses (potential outliers) from the original training set and obtain a cleaner set for adversarial training. We call this baseline AT w/o and compare it with our method. Appendix E.5 shows more details of this baseline. All methods are designed based on the cross-entropy loss.

Defense Settings and Robustness Evaluation. In training, following Zeng et al. (2021a); Liu et al. (2021), we consider robustness by setting $p = \infty$. For each dataset, we test two different perturbation bounds ϵ , i.e., $\epsilon \in \{0.1, 0.2\}$ for MNIST and $\epsilon \in \{2/255, 8/255\}$ for CIFAR-10 and CIFAR-100. The training attack is PGD¹⁰ with random start and step size $\epsilon/4$. We evaluate the robustness of our method and baselines using the standard accuracy on natural test data (Natural), FGSM, PGD²⁰, and the CW attack because they are frequently used in the literature of our compared methods. We also evaluate all methods on a very strong attack, AutoAttack (AA) (Croce & Hein, 2020), under specific settings. All of them are constrained by the same perturbation limit ϵ .

Hyperparameters (k and m) and Outliers Generation. Following Hu et al. (2020), we apply a grid search to select the values of k and m . To simulate the outliers in the training dataset, as in the work of Wang et al. (2019), we add the symmetric (uniform) and asymmetric (class-dependent) noises to the labels of the training data. For symmetric noise creation, we randomly choose training samples with a given probability γ and change each of their labels to another random label. For asymmetric noise creation, given γ , flipping labels only within a specific set of classes. More details of the simulation can be found in Appendix D.4. Note that we use 4 different $\gamma \in \{10\%, 20\%, 30\%, 40\%\}$.

5.2 RESULTS

Overall Performance. We report the overall results (accuracy on the testing sets) in Table 1. First, when there is no noise in the datasets, ORAT outperforms the AT method in three different attack settings on all datasets by a significant margin (0, 2.14%]. This is probably due to that the original datasets may contain outliers, which has been verified in a recent work Sanyal et al. (2021).

Second, under the symmetric noise setting, ORAT outperforms all compared methods in all attack scenarios in general. We observe the performance gap between noise and noise-free settings is small on MNIST because it is a relatively smaller dataset and LeNet achieves almost perfect performance. However, on a larger dataset such as CIFAR-10 and using a more complex neural network, the outliers have a stronger impact. The instance-reweighted AT methods such as GAIRAT, MAIL, and MAIL are inferior to AT and ORAT in 20%, 30%, 40% settings under $\epsilon = 2/255$, because they allocate the loss of the most outliers with higher weight during the training. For CIFAR-100, which is larger and more challenging than CIFAR-10 and the training network ResNet is also larger than Small-CNN. In this case, we find that outliers have a large influence on robust training. Compared to noise-free settings, the performance of AT is reduced by half under the noise settings. However, ORAT still outperforms all baselines in all settings. The performance gap between the baselines and ORAT is more than 2% and even can be achieved near 10% (in 30% symmetric noise, MAIL, FGSM, and $\epsilon=2/255$ settings). Furthermore, we find RAA is more robust than other compared methods but vulnerable to adversarial samples than ORAT. As we discussed before, their label correction approach may correct the wrong labels but inevitably introduces more extra noisy labels during training.

Third, for asymmetric noise settings, we can get similar observations. Note that ORAT can even outperform the baselines over 10% on the CIFAR-10 dataset. All adversarial methods outperform ST in all attack settings since ST does not have a mechanism to handle adversarial samples. In general, increasing the noise percentage, we can find there is a decreasing trend in the performance among all methods. On the other hand, increasing the value of ϵ will also decrease all model performance. Even a small amount of label noise causes classifiers to have significant adversarial errors.

We also report testing accuracy for AA on all datasets with 20% symmetric noise in Table 2. We can find our method outperforms all compared methods even if using a strong attack strategy. In addition, we compare AT w/o with our method and report performance in Table 3. Comparing Table 1 and Table 3 under the same setting, it is evident that simply removing outliers from training data

Noise	Defense	MNIST (LeNet)								CIFAR-10 (Small-CNN)								CIFAR-100 (ResNet-18)								
		$\epsilon=0.1$				$\epsilon=0.2$				$\epsilon=2/255$				$\epsilon=8/255$				$\epsilon=2/255$				$\epsilon=8/255$				
		Na	FG	PGD	CW	Na	FG	PGD	CW	Na	FG	PGD	CW	Na	FG	PGD	CW	Na	FG	PGD	CW	Na	FG	PGD	CW	
0	ST	99.51	92.43	85.68	86.05	99.41	50.19	16.76	18.65	92.48	44.27	24.87	25.06	92.63	15.03	6.82	7.15	73.67	19.64	10.66	11.13	73.86	3.58	1.61	1.66	
	AT	99.51	98.60	97.93	97.94	99.16	97.33	95.14	95.19	89.78	79.40	73.52	73.64	85.52	54.61	33.34	34.28	67.05	53.70	48.18	47.46	55.13	33.66	25.46	23.28	
	ORAT	99.53	98.70	98.03	98.04	99.30	97.79	96.43	96.38	89.54	79.53	73.96	73.86	85.91	54.85	33.46	34.56	68.00	55.60	50.06	49.50	57.82	35.51	27.22	25.26	
	Ours	99.53	98.70	98.03	98.04	99.30	97.79	96.43	96.38	89.54	79.53	73.96	73.86	85.91	54.85	33.46	34.56	68.00	55.60	50.06	49.50	57.82	35.51	27.22	25.26	
Symmetric	10%	ST	99.32	93.72	89.09	87.17	99.32	67.55	36.76	26.02	62.22	29.24	25.72	24.80	61.80	14.70	9.72	11.38	29.60	14.17	11.41	9.96	29.41	5.22	2.54	2.19
		AT	99.12	97.98	97.36	97.28	98.39	96.62	95.25	95.11	63.03	51.36	46.02	44.95	59.53	34.17	23.33	22.28	31.02	22.04	20.25	18.97	28.96	15.08	12.42	10.49
		GAI-RAT	99.11	97.87	97.33	97.20	98.35	96.57	95.22	95.08	60.41	49.23	44.72	42.02	57.12	35.48	24.22	21.52	31.05	21.31	19.76	18.08	27.73	15.13	12.28	10.26
		MAIL	99.09	98.06	97.50	97.40	98.48	96.69	95.47	95.20	55.14	47.44	44.75	39.64	53.54	33.61	26.18	20.23	29.06	18.90	16.62	14.79	27.65	13.21	10.03	7.26
	20%	ST	99.14	97.95	97.32	97.23	98.36	96.58	95.25	95.11	61.83	51.43	46.56	44.97	57.97	34.74	25.15	21.96	29.74	20.11	18.65	16.86	29.23	14.78	11.84	9.96
		AT	99.22	98.01	97.37	97.31	98.65	96.83	95.32	95.14	62.94	51.32	46.55	45.36	58.28	33.58	23.71	22.89	30.72	21.95	20.34	18.83	28.99	15.60	13.03	11.03
		GAI-RAT	99.52	98.45	97.78	97.79	99.20	97.61	96.07	96.11	63.83	52.04	48.77	46.65	60.42	36.13	26.86	23.48	35.76	25.72	22.27	21.28	35.35	20.37	15.50	14.02
		MAIL	99.12	93.35	89.05	86.87	99.14	70.18	41.67	29.30	58.89	31.23	25.97	24.13	57.28	15.88	12.8	13.26	26.16	14.74	11.98	10.29	26.22	5.09	2.45	1.92
	30%	ST	98.88	97.67	97.03	96.89	97.97	95.98	94.67	94.50	61.10	50.67	46.76	45.28	56.96	34.55	25.54	24.34	27.52	20.12	18.49	17.15	26.77	14.39	11.87	9.96
		AT	98.82	97.59	96.96	96.79	97.86	95.81	94.57	94.30	58.58	49.26	45.44	40.62	53.55	34.40	25.52	22.15	26.77	19.93	18.30	16.78	24.80	14.28	12.03	10.32
		GAI-RAT	98.81	97.82	97.29	97.12	97.98	96.14	94.81	94.58	54.04	47.05	44.78	41.05	52.93	34.28	27.74	21.90	25.13	17.01	16.06	14.32	24.54	12.85	10.02	7.70
		MAIL	98.88	97.59	97.00	96.89	97.85	95.89	94.50	94.36	59.01	49.64	45.96	43.71	56.04	34.65	26.22	23.70	26.72	19.43	18.04	16.53	26.61	13.81	11.94	9.82
40%	ST	99.07	97.70	97.07	97.05	98.25	96.28	94.41	94.20	59.53	49.56	45.73	45.60	56.71	33.24	24.51	24.12	28.16	20.73	19.15	17.74	25.34	14.79	12.46	10.25	
	AT	99.56	98.37	97.65	97.64	99.06	97.33	95.71	95.68	61.03	51.85	48.48	46.16	54.37	35.71	27.76	24.70	34.45	25.07	22.21	20.92	32.83	19.45	15.49	13.34	
	GAI-RAT	98.91	93.01	88.62	85.73	98.97	68.59	41.24	27.73	58.17	30.55	25.71	24.54	57.55	14.18	11.50	11.41	23.38	14.26	11.84	10.10	24.29	5.30	2.97	2.35	
	MAIL	98.33	96.86	96.27	96.11	96.56	94.37	92.55	92.18	37.16	43.45	41.54	40.13	40.65	31.59	26.74	24.98	22.20	18.28	17.33	15.92	23.41	13.75	11.25	9.26	
Asymmetric	10%	ST	99.39	93.05	86.09	89.09	99.39	65.48	22.17	35.02	65.20	31.18	22.97	22.95	64.90	13.41	9.65	10.10	31.81	15.10	11.23	10.03	31.71	4.34	1.84	1.71
		AT	99.46	98.14	97.36	97.61	99.08	97.29	95.21	95.64	64.71	51.75	45.88	45.70	60.82	33.98	20.71	21.30	33.21	22.38	20.53	19.09	30.83	15.80	12.49	10.55
		GAI-RAT	99.44	98.13	97.34	97.50	99.03	97.20	95.14	95.45	62.31	49.42	43.64	41.28	58.32	35.07	23.65	21.97	33.03	22.73	20.63	19.38	29.88	15.58	12.49	10.57
		MAIL	99.50	98.20	97.38	97.73	99.08	97.52	95.79	96.35	56.91	48.29	45.40	36.39	53.96	35.75	28.61	14.94	31.45	21.36	17.92	16.43	29.66	13.99	10.78	7.64
	20%	ST	99.45	98.11	97.38	97.59	99.08	97.28	95.30	95.67	63.93	51.58	46.73	43.92	60.01	34.71	23.37	20.70	32.52	22.08	19.63	17.81	30.89	14.86	11.91	10.01
		AT	99.42	98.11	97.34	97.51	99.08	97.24	95.17	95.45	62.05	50.39	45.61	45.52	58.56	32.75	22.70	22.13	33.01	23.07	21.07	19.69	30.35	16.09	12.79	10.70
		GAI-RAT	99.50	98.58	97.81	98.04	99.18	97.68	96.11	96.36	65.41	52.76	47.75	46.21	61.28	36.54	28.88	23.52	37.09	27.07	23.65	22.59	35.72	20.51	15.46	13.83
		MAIL	99.19	92.64	86.00	88.63	99.23	64.76	20.49	32.18	62.29	22.77	18.34	16.56	62.14	10.46	6.89	6.96	30.46	14.79	11.57	10.38	30.49	4.92	2.14	2.03
	30%	ST	99.44	98.12	97.29	97.59	99.07	97.11	95.20	95.84	37.29	46.54	42.36	41.63	35.48	32.05	24.20	23.33	31.67	21.15	19.44	18.19	29.96	15.47	12.46	10.51
		AT	99.41	98.10	97.20	97.54	99.01	96.99	95.16	95.66	37.36	43.98	40.30	38.73	35.12	31.98	24.05	23.06	31.64	21.64	19.40	18.42	29.15	15.45	12.58	10.70
		GAI-RAT	99.46	98.19	97.38	97.78	99.17	97.38	95.67	96.26	32.26	41.99	38.28	32.29	27.98	32.09	25.97	17.24	30.94	20.38	16.98	15.60	28.48	13.32	9.89	7.18
		MAIL	99.44	98.11	97.19	97.60	99.06	97.07	95.22	95.83	33.63	45.84	42.90	40.44	36.56	31.97	25.58	23.24	31.90	20.82	18.96	17.31	30.61	14.28	11.53	9.85
40%	ST	99.40	98.02	97.14	97.53	99.04	97.00	95.11	95.62	46.44	38.01	34.80	33.07	43.22	27.58	21.07	20.20	30.58	21.08	19.19	17.98	29.08	15.18	12.28	10.38	
	AT	99.47	98.57	97.82	98.12	99.24	97.45	95.98	96.63	39.98	46.51	42.98	41.30	35.71	33.71	26.04	24.06	34.58	24.18	21.05	20.11	33.80	19.35	14.84	13.70	
	GAI-RAT	97.43	81.71	73.74	75.11	97.58	50.69	16.99	22.82	58.87	19.16	16.29	16.04	59.07	10.52	7.01	8.81	26.99	12.97	10.29	9.56	27.11	4.32	1.91	1.67	
	MAIL	99.26	97.70	96.78	97.03	98.85	96.65	94.36	95.24	32.71	42.02	38.33	36.30	33.97	31.28	23.17	23.26	28.42	19.67	17.88	16.75	25.70	14.49	11.53	9.85	

Table 1: Testing accuracy (%) of seven methods on MNIST, CIFAR-10, and CIFAR-100 with different levels of symmetric and asymmetric noisy. The best results are shown in bold. We color the performance of all adversarial training methods on three different attacks. The performance gap between current method and ORAT are shown in green: $\leq 2\%$; yellow: (2%,5%]; orange: (5%,10%]; blue: $>10\%$. According to results, ORAT outperforms all adversarial training methods on all attack settings. ‘Na’ represents Natural. ‘FG’ represents FGSM.

cannot significantly improve AT performance, and our method still achieves the best performance. The stability evaluation results of each method with 10 random runs are shown in Table 4, where we use 40% symmetric noisy data as an example. Comparing Table 1 and Table 4 under the same setting, it is clear that the performance gap becomes larger when we report scores by using mean and standard deviation, and our method shows a stable and stronger ability in handling outliers and adversarial attacks.

Convergence and Robustness Tendency. We show the tendency curves of the training loss on all datasets with different noise when using Algorithm 1 in the left panel of Figure 2. The sharp drops in the curves correspond to decreases in training learning rate. We can observe a steady decrease in the training loss on all datasets with the increase of training epochs when training against adversarial samples, which supports that Algorithm 1 can effectively optimize and solve Eq. (4) even if it is a non-smooth loss. We also compare the robust accuracy of AT and ORAT by using four attack strategies on all datasets with 20% symmetric noise in Figure 2 right panel. It is clear that our method

Defense	MNIST		CIFAR-10		CIFAR-100	
	$\epsilon=0.1$	$\epsilon=0.2$	$\epsilon=2/255$	$\epsilon=8/255$	$\epsilon=2/255$	$\epsilon=8/255$
ST	84.13	23.78	23.54	10.41	10.22	1.72
AT	96.73	93.23	44.56	23.06	17.11	9.89
GAIRAT	96.78	93.57	40.82	20.57	16.54	9.97
MAIL	97.01	93.80	40.39	18.95	13.67	7.29
WMMR	96.88	93.60	43.52	20.68	15.69	8.73
RAA	96.98	93.62	45.07	22.29	17.70	9.78
Ours	97.67	95.05	45.59	23.35	19.74	12.22

Table 2: Testing accuracy (%) on AutoAttack.

Noise	Defense	MNIST ($\epsilon = 0.1$)				CIFAR-100 ($\epsilon = 2/255$)			
		Na	FG	PGD	CW	Na	FG	PGD	CW
10%	AT w/o	98.91	98.08	97.61	97.55	29.81	21.09	19.59	18.23
	Ours	99.52	98.45	97.78	97.79	35.76	25.72	22.27	21.28
30%	AT w/o	97.82	96.97	96.47	96.35	24.18	19.17	17.31	16.71
	Ours	99.55	98.30	97.51	97.53	31.27	23.81	21.35	19.59

Table 3: Testing accuracy (%) of AT w/o and ORAT on symmetric noisy data.

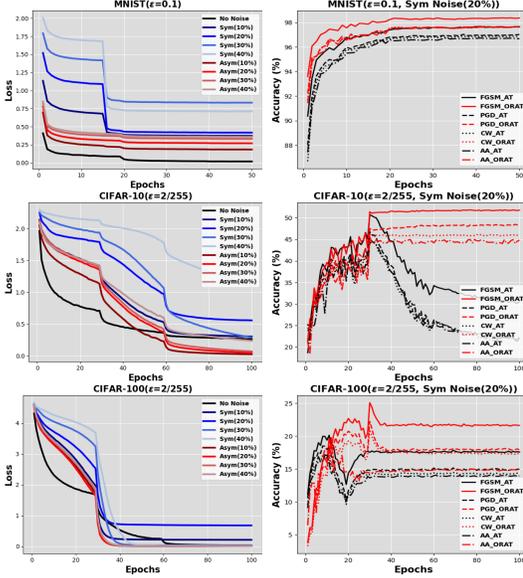


Figure 2: The tendency curves of training adversarial loss and test accuracy on three datasets.

Noise	Defense	MNIST ($\epsilon = 0.1$)				CIFAR-100 ($\epsilon = 2/255$)			
		Na	FG	PGD	CW	Na	FG	PGD	CW
40% Symmetric Noise	ST	98.38 (0.18)	77.47 (5.79)	59.31 (11.13)	51.47 (12.79)	21.53 (0.71)	13.95 (0.16)	12.12 (0.28)	10.12 (0.28)
	AT	96.60 (0.87)	95.11 (1.17)	94.36 (1.31)	94.16 (1.37)	21.82 (0.30)	18.14 (0.27)	17.02 (0.29)	15.62 (0.32)
	GAIRAT	98.04 (0.11)	96.60 (0.08)	95.98 (0.13)	95.79 (0.12)	21.58 (0.30)	17.94 (0.36)	16.88 (0.36)	15.50 (0.33)
	MAIL	98.06 (0.17)	96.88 (0.16)	96.31 (0.12)	96.11 (0.10)	17.52 (0.74)	14.31 (0.43)	13.46 (0.40)	12.09 (0.37)
	WMMR	98.07 (0.14)	96.69 (0.10)	96.04 (0.08)	95.85 (0.08)	20.64 (0.29)	17.14 (0.33)	16.10 (0.30)	14.76 (0.27)
	RAA	98.58 (0.11)	96.82 (0.08)	95.99 (0.10)	95.91 (0.09)	22.21 (0.32)	18.27 (0.30)	17.22 (0.30)	15.85 (0.33)
	Ours	99.34 (0.03)	98.02 (0.07)	97.19 (0.08)	97.18 (0.07)	30.12 (0.42)	23.62 (0.49)	21.48 (0.49)	20.02 (0.41)

Table 4: Mean and standard deviation (in parentheses) of testing accuracy (%) across 10 random runs.

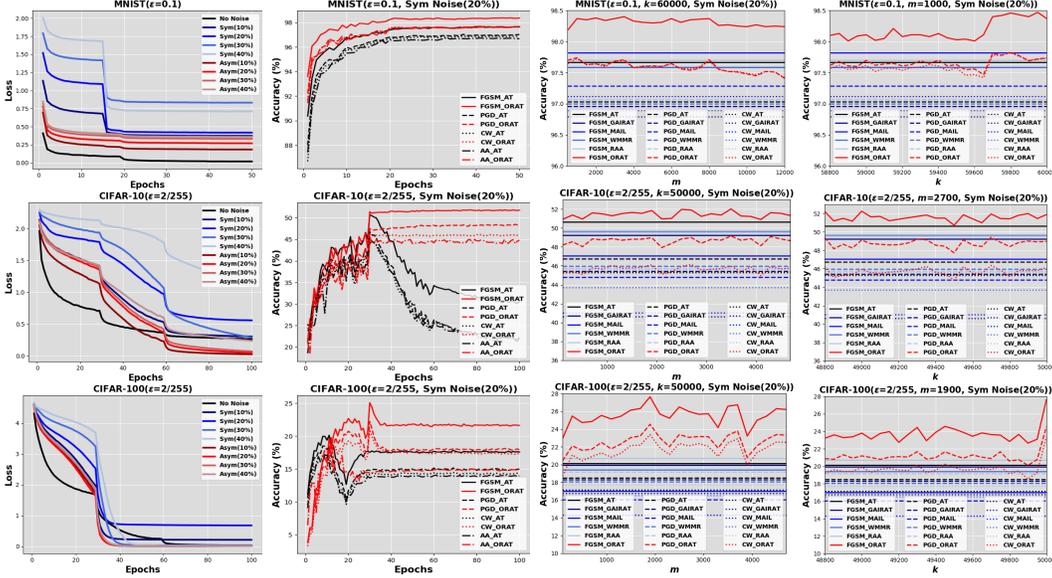


Figure 3: Effect of k and m on the test accuracy of ORAT on three datasets.

outperforms the original AT method in all settings. In general, these plots illustrate that we can consistently reduce the value of the objective function of ORAT, thus producing an increasingly robust classifier.

Effect of k and m . We study how the choices of k and m affect the performance of ORAT with two types of experiments on all datasets under 20% symmetric noise setting, together with those from other defense methods. In the first set of experiments, we fix k to the total number of training samples and run the algorithm with different values of m . The results are plotted in Figure 3 (left panel). We can see that there is a clear range of m with better performance than all compared methods. In the second set of experiments, we fix m with the best performance from the first set of experiments and run ORAT with different values of k . The results are shown in Figure 3 (right panel). Note that there is also a range of k with better performance, in particular, the optimal value of k is less than the number of total training samples. Similar trends are observed on all datasets.

6 CONCLUSION

In this work, we introduce the outlier robust adversarial training (ORAT), which considers both outliers in training data and adversarial attacks in the model training. We provide an optimizing algorithm and analyze the theoretical aspects of ORAT. Empirical results showed the effectiveness and robustness of ORAT on three benchmark datasets. In the future, we will study the optimization error or convergence rate of our proposed learning algorithm. We will also evaluate our method on large datasets and large deep neural networks. Although achieving fairness is not a goal of this work, we have found that our method can benefit the minority subgroup of data (see Figure 1 right panel). Studying fairness with ORAT is an interesting direction in the future. Furthermore, we plan to design an efficient method to automatically determine the hyperparameters k and m during the training.

REFERENCES

- Maximilian Augustin, Alexander Meinke, and Matthias Hein. Adversarial robustness on in-and out-distribution improves explainability. In *European Conference on Computer Vision*, pp. 228–245. Springer, 2020.
- Pranjal Awasthi, Natalie Frank, and Mehryar Mohri. Adversarial learning guarantees for linear hypotheses and neural networks. In *International Conference on Machine Learning*, pp. 431–441. PMLR, 2020.
- Pranjal Awasthi, Natalie Frank, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Calibration and consistency of adversarial surrogate losses. *Advances in Neural Information Processing Systems*, 2021.
- Han Bao, Clay Scott, and Masashi Sugiyama. Calibrated surrogate losses for adversarially robust classification. In *Conference on Learning Theory*, pp. 408–451. PMLR, 2020.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Julian Bitterwolf, Alexander Meinke, and Matthias Hein. Certifiably adversarially robust detection of out-of-distribution data. *Advances in Neural Information Processing Systems*, 33:16085–16095, 2020.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 39–57. IEEE, 2017.
- Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pp. 2206–2216. PMLR, 2020.
- Yinpeng Dong, Zhijie Deng, Tianyu Pang, Jun Zhu, and Hang Su. Adversarial distributional training for robust deep learning. *Advances in Neural Information Processing Systems*, 33:8270–8283, 2020.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2020.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *ICLR*, 2015.
- Stephen Gould, Basura Fernando, Anoop Cherian, Peter Anderson, Rodrigo Santa Cruz, and Edison Guo. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *arXiv preprint arXiv:1607.05447*, 2016.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018.
- Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5138–5147, 2019.
- Kai Han, Yunhe Wang, Yixing Xu, Chunjing Xu, Enhua Wu, and Chang Xu. Training binary neural networks through learning with noisy supervision. In *International Conference on Machine Learning*, pp. 4017–4026. PMLR, 2020.
- Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8340–8349, 2021.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*, pp. 409–426. Springer, 1994.
- Chester Holtz, Tsui-Wei Weng, and Gal Mishne. Learning sample reweighting for adversarial robustness. *ICLR Submission*, 2021.
- Shu Hu, Yiming Ying, Siwei Lyu, et al. Learning by minimizing the sum of ranked range. *Advances in Neural Information Processing Systems*, 33, 2020.
- Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pp. 2304–2313. PMLR, 2018.
- Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. *Advances in Neural Information Processing Systems*, 2008.
- Vladimir Koltchinskii and Dmitry Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Pierre Laforgue, Guillaume Staerman, and Stephan Cléménçon. Generalization bounds in the presence of outliers: a median-of-means study. In *International Conference on Machine Learning*, pp. 5937–5947. PMLR, 2021.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 1991.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.
- Feng Liu, Bo Han, Tongliang Liu, Chen Gong, Gang Niu, Mingyuan Zhou, Masashi Sugiyama, et al. Probabilistic margins for instance reweighting in adversarial training. *Advances in Neural Information Processing Systems*, 34, 2021.
- Yang Liu and Hongyi Guo. Peer loss functions: Learning from noisy labels without knowing noise rates. In *International Conference on Machine Learning*, pp. 6226–6236. PMLR, 2020.
- Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *International Conference on Machine Learning*, pp. 6448–6458. PMLR, 2020.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- Colin McDiarmid et al. On the method of bounded differences. *Surveys in combinatorics*, 141(1): 148–188, 1989.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.

- Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Icml*, 2010.
- Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? *Advances in Neural Information Processing Systems*, 33:3533–3545, 2020.
- Amartya Sanyal, Puneet K Dokania, Varun Kanade, and Philip Torr. How benign is benign overfitting? In *International Conference on Learning Representations*, 2021.
- Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. In *3rd International Conference on Learning Representations workshops, ICLR 2015*, 2015.
- Neeraj Varshney, Swaroop Mishra, and Chitta Baral. Investigating selective prediction approaches across several tasks in iid, ood, and adversarial settings. *arXiv preprint arXiv:2203.00211*, 2022.
- Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8688–8696, 2018.
- Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 322–330, 2019.
- Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2020.
- Mingyang Yi, Lu Hou, Jiacheng Sun, Lifeng Shang, Xin Jiang, Qun Liu, and Zhiming Ma. Improved ood generalization via adversarial training and pretraing. In *International Conference on Machine Learning*, pp. 11987–11997. PMLR, 2021.
- Dong Yin, Ramchandran Kannan, and Peter Bartlett. Rademacher complexity for adversarially robust generalization. In *International conference on machine learning*, pp. 7085–7094. PMLR, 2019.
- Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pp. 7164–7173. PMLR, 2019.
- Huimin Zeng, Chen Zhu, Tom Goldstein, and Furong Huang. Are adversarial examples created equal? a learnable weighted minimax risk for robustness under non-uniform attacks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 10815–10823, 2021a.
- Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu, and Weiran Xu. Adversarial self-supervised learning for out-of-domain detection. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5631–5639, 2021b.
- Runtian Zhai, Chen Dan, J Zico Kolter, and Pradeep Ravikumar. Doro: Distributional and outlier robust optimization. *ICML*, 2021.
- Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pp. 7472–7482. PMLR, 2019.
- Jingfeng Zhang, Jianing Zhu, Gang Niu, Bo Han, Masashi Sugiyama, and Mohan Kankanhalli. Geometry-aware instance-reweighted adversarial training. In *International Conference on Learning Representations*, 2021.

Yuchen Zhang and Percy Liang. Defending against whitebox adversarial attacks via randomized discretization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 684–693. PMLR, 2019.

Jianing Zhu, Jingfeng Zhang, Bo Han, Tongliang Liu, Gang Niu, Hongxia Yang, Mohan Kankanhalli, and Masashi Sugiyama. Understanding the interaction of adversarial training with noisy labels. *arXiv preprint arXiv:2102.03482*, 2021.

Appendix

A RELATED WORKS

Traditional Robust Learning. Training accurate machine learning models in the presence of noisy data is of great practical importance (Sukhbaatar et al., 2015). However, a degradation in the performance of classification models is inevitable when there are outliers in the training data. To combat outliers, the traditional robust learning methods are designed from four directions. 1) The *label correction methods* (Wang et al., 2018) improve the quality of the raw labels by correcting wrong labels into correct ones. However, it requires an extra clean dataset or potentially expensive detection process to estimate the outliers. 2) The *loss correction methods* (Han et al., 2020; Liu & Guo, 2020) improve the robustness by modifying the loss function based on an estimated noise transition matrix that defines the probability of mislabeling one class with another. However, these methods are sensitive to the noise transition matrix, which is also hard to be estimated. 3) The *refined training strategies* such as Co-teaching (Han et al., 2018; Yu et al., 2019), MentorNet (Jiang et al., 2018; Yu et al., 2019) are robust to outliers. These studies all rely on an auxiliary network for sample weighting or learning supervision, which is hard to adapt and tune. 4) Some simpler and arguably generic *robust loss functions* are also designed for robust learning. For example, a recent work Hu et al. (2020) proposed AoRR loss, which can mitigate the influence of the outliers if their proportion in training data is known. Furthermore, Some smoothing methods are proposed in (Lukasik et al., 2020; Chaudhari et al., 2019; Foret et al., 2020) and have been proved to be effective in solving the problems of label and data noise. However, none of these methods are related to adversarial robust learning.

Adversarial Robust Learning. The omnipotent DNN models are surprisingly vulnerable to adversarial examples (Goodfellow et al., 2015; Carlini & Wagner, 2017), which can easily mislead a DNN model to make erroneous predictions. To mitigate this issue, the adversarial training (AT) (Madry et al., 2018) is first proposed as one of the most effective robust learning methodologies against adversarial attacks. To improve adversarial robustness, instance-reweighted AT methods are studied by considering the unequal importance of the adversarial data in several recent works. Intuitively, the samples assigned a low weight to correspond to samples on which the classifier is already sufficiently robust. Specifically, the reweight mechanism in WMMR (Zeng et al., 2021a) and MAIL (Liu et al., 2021) is based on the multi-class probabilistic margin margin of the model outputs (Koltchinskii & Panchenko, 2002; Zhang & Liang, 2019). The reweighting method in work GAIKAT (Zhang et al., 2021) identifies non-robust (easily be-attacked) data by estimating how many steps the PGD method needs to attack natural data successfully. The most recent work BiLAW (Holtz et al., 2021) uses a validation set to learn weights based on bi-level optimization and meta-learning. The most significant assumption in these works is that the natural dataset is clean. However, the performance of the model based on these methods will be degraded if the training dataset contains outliers. In Sanyal et al. (2021), the authors identified label noise as one of the causes of adversarial vulnerability. However, no defense methods are proposed to solve this problem. The work Zhu et al. (2021) empirically studies the efficacy of AT for mitigating the effect of label noise in training data. However, their proposed annotator algorithm is based on the label correction strategy, which inevitably introduces more extra noisy labels due to the bottleneck of the classifier. In Dong et al. (2020), the authors proposed an adversarial distributional training. They focus on the distribution shift of adversarial samples but they do not consider the outliers problem. Several works Augustin et al. (2020); Bitterwolf et al. (2020); Lee et al. (2018) connect adversarial robustness to out-of-distribution (OOD) problems. However, they are in different settings from ours because the notion of outliers is different from OOD points.

B EXPLICIT FORMS OF $\partial_\theta \hat{\mathcal{L}}$, $\partial_\lambda \hat{\mathcal{L}}$, AND $\partial_{\hat{\lambda}} \hat{\mathcal{L}}$

From Eq.(4), we have $\hat{\mathcal{L}}(f_\theta, \lambda, \hat{\lambda}) := \frac{k-m}{n} \lambda + \frac{n-m}{n} \hat{\lambda} - [\hat{\lambda} - [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_+$. Denote $\mathbb{I}_{[a]}$ as an indicator function with $\mathbb{I}_{[a]} = 1$ if a is true and 0 otherwise. Then we can get

$$\begin{aligned} \partial_\theta \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)}) &= \partial \ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i) \cdot \mathbb{I}_{[\hat{\lambda}^{(t)} > [\ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i) - \lambda^{(t)}]_+]} \cdot \mathbb{I}_{[\ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i) > \lambda^{(t)}]}, \\ \partial_\lambda \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)}) &= \frac{k-m}{n} - \mathbb{I}_{[\hat{\lambda}^{(t)} > [\ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i) - \lambda^{(t)}]_+]} \cdot \mathbb{I}_{[\ell(f_{\theta^{(t)}}(\tilde{\mathbf{x}}_i), y_i) > \lambda^{(t)}]}, \end{aligned}$$

$$\partial_{\hat{\lambda}} \hat{\mathcal{L}}(f_{\theta^{(t)}}, \lambda^{(t)}, \hat{\lambda}^{(t)}) = \frac{n-m}{n} - \mathbb{I}_{[\hat{\lambda}^{(t)} > [\ell(f_{\theta^{(t)}}(\bar{\mathbf{x}}_i), y_i) - \lambda^{(t)}]_+]}$$

C PROOFS

C.1 PROOF OF THEOREM 1

Denote $[a]_+ = \max\{0, a\}$ as the hinge function. First, we introduce two Lemmas as follows,

Lemma 1. (Hu et al., 2020) For a set of real numbers $S = \{s_1, \dots, s_n\}$, $s_i \in \mathbb{R}$, and $s_{[i]}$ represents the i -th largest value after sorting the elements in S , we have

$$\sum_{i=1}^k s_{[i]} = \min_{\lambda \in \mathbb{R}} \left\{ k\lambda + \sum_{i=1}^n [s_i - \lambda]_+ \right\}.$$

Furthermore, $s_{[k]} \in \arg \min_{\lambda \in \mathbb{R}} \{k\lambda + \sum_{i=1}^n [s_i - \lambda]_+\}$.

Proof. We know $\sum_{i=1}^k s_{[i]}$ is the solution of

$$\max_{\mathbf{p}} \mathbf{p}^\top S, \text{ s.t. } \mathbf{p}^\top \mathbf{1} = k, \mathbf{0} \leq \mathbf{p} \leq \mathbf{1}.$$

We apply Lagrangian to this equation and get

$$L = -\mathbf{p}^\top S - \mathbf{v}^\top \mathbf{p} + \mathbf{u}^\top (\mathbf{p} - \mathbf{1}) + \lambda(\mathbf{p}^\top \mathbf{1} - k)$$

where $\mathbf{u} \geq \mathbf{0}$, $\mathbf{v} \geq \mathbf{0}$ and $\lambda \in \mathbb{R}$ are Lagrangian multipliers. Taking its derivative w.r.t. \mathbf{p} and set it to 0, we have $\mathbf{v} = \mathbf{u} - S + \lambda \mathbf{1}$. Substituting it back into the Lagrangian, we get

$$\min_{\mathbf{u}, \lambda} \mathbf{u}^\top \mathbf{1} + k\lambda, \text{ s.t. } \mathbf{u} \geq \mathbf{0}, \mathbf{u} + \lambda \mathbf{1} - S \geq \mathbf{0}.$$

This means

$$\sum_{i=1}^k s_{[i]} = \min_{\lambda} \left\{ k\lambda + \sum_{i=1}^n [s_i - \lambda]_+ \right\}. \quad (\text{C.1})$$

Furthermore, we can see that $\lambda = s_{[k]}$ is always one optimal solution for Eq.(C.1). So

$$s_{[k]} \in \arg \min_{\lambda} \left\{ k\lambda + \sum_{i=1}^n [s_i - \lambda]_+ \right\}.$$

□

Lemma 2. For a set of real numbers $S = \{s_1, \dots, s_n\}$, $s_i \in \mathbb{R}$, we have

$$\sum_{i=m+1}^n s_{[i]} = \max_{\lambda \in \mathbb{R}} \left\{ (n-m)\lambda - \sum_{i=1}^n [\lambda - s_i]_+ \right\}.$$

Furthermore, $s_{[m]} \in \arg \max_{\lambda \in \mathbb{R}} \{(n-m)\lambda - \sum_{i=1}^n [\lambda - s_i]_+\}$.

Proof.

$$\begin{aligned} \sum_{i=m+1}^n s_{[i]} &= \sum_{i=1}^n s_i - \sum_{i=1}^m s_{[i]} \\ &= \sum_{i=1}^n s_i - \min_{\lambda} \left\{ m\lambda + \sum_{i=1}^m [s_i - \lambda]_+ \right\} \\ &= -\min_{\lambda} \left\{ -\sum_{i=1}^m (s_i - \lambda) - (n-m)\lambda + \sum_{i=1}^n [s_i - \lambda]_+ \right\} \\ &= -\min_{\lambda} \left\{ -(n-m)\lambda + \sum_{i=1}^n [\lambda - s_i]_+ \right\} \\ &= \max_{\lambda} \left\{ (n-m)\lambda - \sum_{i=1}^n [\lambda - s_i]_+ \right\} \end{aligned}$$

The second equation holds because of Lemma 1. The fourth equation holds because the fact of $[a]_+ - a = [-a]_+$. Furthermore, we can see that $\lambda = s_{[m]}$ is always one optimal solution. So

$$s_{[m]} \in \arg \max_{\lambda \in \mathbb{R}} \left\{ (n-m)\lambda - \sum_{i=1}^n [\lambda - s_i]_+ \right\}.$$

□

Theorem C.1. (Theorem 1 restated) Suppose $\lambda \in \mathbb{R}$, $\hat{\lambda} \in \mathbb{R}$, then Eq.(3) is equivalent to

$$\begin{aligned} \min_{\theta, \lambda} \max_{\hat{\lambda}} & \frac{1}{k-m} \sum_{i=1}^n \left[\frac{k-m}{n} \lambda + \frac{n-m}{n} \hat{\lambda} - [\hat{\lambda} - [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_+ \right] \\ \text{s.t. } & \tilde{\mathbf{x}}_i = \arg \max_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x}_i)} \ell(f_\theta(\tilde{\mathbf{x}}), y_i) \end{aligned} \quad (\text{C.2})$$

Furthermore, $\hat{\lambda} > \lambda$, when the optimal solution is achieved.

Proof. To extract the sum of (m, k) -ranked range individual losses, we can first select a subset, which contains the bottom $n-m$ losses from the ranked list of $L(\{\mathbf{x}_j, y_j\}_{j=1}^n)$. Then we select top- $(k-m)$ individual losses from this subset as the finalized (m, k) -ranked range. Therefore, We sum the bottom $n-m$ individual losses as follows,

$$\sum_{i=m+1}^n \ell(f_\theta(\tilde{\mathbf{x}}_{[i]}, y_{[i]})) = \min_q \sum_{i=1}^n q_i \ell(f_\theta(\tilde{\mathbf{x}}_{[i]}, y_{[i]})) \quad \text{s.t. } q_i \in \{0, 1\}, \|q\|_0 = n-m,$$

where $q = \{q_1, \dots, q_n\} \in \{0, 1\}^n$, and q_i is an indicator. When $q_i = 0$, it indicates that the i -th individual loss is not included in the objective function. Otherwise, the objective function should include this individual loss. Next, we sum the top- $(k-m)$ individual losses from the bottom $n-m$ individual losses as follows,

$$\begin{aligned} & \min_q \sum_{i=1}^{k-m} (q \ell(f_\theta(\tilde{\mathbf{x}}), y))_{[i]} \quad \text{s.t. } q_i \in \{0, 1\}, \|q\|_0 = n-m \\ &= \min_{\lambda, q} (k-m)\lambda + \sum_{i=1}^n [q_i \ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \quad \text{s.t. } q_i \in \{0, 1\}, \|q\|_0 = n-m \\ &= \min_{\lambda, q} (k-m)\lambda + \sum_{i=1}^n q_i [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \quad \text{s.t. } q_i \in [0, 1], \|q\|_0 = n-m \quad (\text{C.3}) \\ &= \min_{\lambda} (k-m)\lambda + \sum_{i=m+1}^n [[\ell(f_\theta(\tilde{\mathbf{x}}), y) - \lambda]_+]_{[i]} \\ &= \min_{\lambda} (k-m)\lambda + \max_{\hat{\lambda}} \left\{ (n-m)\hat{\lambda} - \sum_{i=1}^n [\hat{\lambda} - [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_+ \right\}, \end{aligned}$$

where $q \ell(f_\theta(\tilde{\mathbf{x}}), y) = \{q_1 \ell(f_\theta(\tilde{\mathbf{x}}_1), y_1), \dots, q_n \ell(f_\theta(\tilde{\mathbf{x}}_n), y_n)\}$. The first equation holds because of Lemma 1. Since $q_i \ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) \geq 0$, we know the optimal $\lambda^* \geq 0$ from Lemma 1. If $q_i = 0$, $[q_i \ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda^*]_+ = 0 = q_i [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda^*]_+$. If $q_i = 1$, $[q_i \ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda^*]_+ = [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda^*]_+ = q_i [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda^*]_+$. Thus the second equation holds. It should be mentioned that the discrete indicator q_i can be replaced by a continue one, which means $q_i \in [0, 1]$. The third equation holds because we take the optimal q^* into the objective function and remove the constraints. The fourth equation can be obtained by applying Lemma 2.

Therefore,

$$\begin{aligned} & \min_{\theta} \frac{1}{k-m} \sum_{i=m+1}^k \ell(f_\theta(\tilde{x}_{[i]}, y_{[i]})) \\ &= \min_{\theta} \frac{1}{k-m} \left\{ \min_{\lambda} (k-m)\lambda + \max_{\hat{\lambda}} \left\{ (n-m)\hat{\lambda} - \sum_{i=1}^n [\hat{\lambda} - [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_+ \right\} \right\} \quad (\text{C.4}) \\ &= \min_{\theta, \lambda} \max_{\hat{\lambda}} \frac{1}{k-m} \sum_{i=1}^n \left[\frac{k-m}{n} \lambda + \frac{n-m}{n} \hat{\lambda} - [\hat{\lambda} - [\ell(f_\theta(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_+ \right]. \end{aligned}$$

Furthermore, according to Lemma 1 and 2, we know the optimal λ^* and $\hat{\lambda}^*$ can be obtained at the top- k and top- m values of loss ℓ , respectively. Since $m < k$, we have $\lambda^* < \hat{\lambda}^*$. Therefore, $\hat{\lambda} > \lambda$, when the optimal solution is achieved. \square

C.2 PROOF OF THEOREM 2

To prove Theorem 2, we first introduce the calibration function as follows,

Definition 4. (*Calibration function*). (Awasthi et al., 2021) Given a hypothesis set \mathcal{H} , we define the calibration function δ_{\max} for a pair of losses (ℓ_1, ℓ_2) as follows: for all $\mathbf{x} \in \mathcal{X}$, $\eta \in [0, 1]$ and $\tau > 0$,

$$\delta_{\max}(\tau, \mathbf{x}, \eta) = \inf_{f \in \mathcal{H}} \{ \mathcal{C}_{\ell_1}(f, \mathbf{x}, \eta) - \mathcal{C}_{\ell_1, \mathcal{H}}^*(\mathbf{x}, \eta) | \mathcal{C}_{\ell_2}(f, \mathbf{x}, \eta) - \mathcal{C}_{\ell_2, \mathcal{H}}^*(\mathbf{x}, \eta) \geq \tau \}. \quad (\text{C.5})$$

The calibration function gives the maximal δ satisfying the calibration condition (Definition 2). The following proposition is an important result from Steinwart (2007).

Proposition 1. (Steinwart, 2007). Given a hypothesis set \mathcal{H} , loss ℓ_1 is \mathcal{H} -calibrated with respect to ℓ_2 if and only if its calibration function δ_{\max} satisfies $\delta_{\max}(\tau, \mathbf{x}, \eta) > 0$ for all $\mathbf{x} \in \mathcal{X}$, $\eta \in [0, 1]$, and $\tau > 0$.

Next, we define the adversarial loss of $f \in \mathcal{H}$ at (\mathbf{x}, y) as

$$\tilde{\ell}_s(f, \mathbf{x}, y) = \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \ell_s(yf(\tilde{\mathbf{x}})). \quad (\text{C.6})$$

The above naturally motivates supremum-based surrogate losses that are commonly used to optimize the adversarial 0/1 loss (Goodfellow et al., 2015; Madry et al., 2018; Zhang et al., 2019). When ℓ_s is non-increasing, the following equality holds (Yin et al., 2019):

$$\sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \ell_s(yf(\tilde{\mathbf{x}})) = \ell_s\left(\inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} yf(\tilde{\mathbf{x}})\right). \quad (\text{C.7})$$

Therefore, the adversarial 0/1 loss $\tilde{\ell}_0$ has the equivalent form

$$\tilde{\ell}_0(f, \mathbf{x}, y) := \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \mathbb{1}_{yf(\tilde{\mathbf{x}}) \leq 0} = \mathbb{1}_{\inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} yf(\tilde{\mathbf{x}}) \leq 0}. \quad (\text{C.8})$$

In this paper, we aim to characterize surrogate losses ℓ_1 satisfying \mathcal{H} -calibration (Definition 2) with $\ell_2 = \tilde{\ell}_0$ and for the hypothesis sets \mathcal{H} which are regular for adversarial calibration.

For convenience, let $\underline{M}(f, \mathbf{x}, \epsilon) := \inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} f(\tilde{\mathbf{x}})$ and $\overline{M}(f, \mathbf{x}, \epsilon) := -\inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} -f(\tilde{\mathbf{x}}) = \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} f(\tilde{\mathbf{x}})$. Then we provide three useful Lemmas as follows,

Lemma 3. (Awasthi et al. (2021), Lemma 28). Let \mathcal{H} be a symmetric hypothesis set, ℓ be a surrogate loss function, and $\mathcal{X}_2 = \{\mathbf{x} \in \mathcal{X} : \text{there exists } f' \in \mathcal{H} \text{ such that } \underline{M}(f', \mathbf{x}, \epsilon) > 0\}$. If $\mathcal{X}_2 = \emptyset$, any loss ℓ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$. If $\mathcal{X}_2 \neq \emptyset$, then ℓ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$ if and only if for any $\mathbf{x} \in \mathcal{X}_2$,

$$\begin{aligned} \inf_{f \in \mathcal{H} : \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_\ell(f, \mathbf{x}, \frac{1}{2}) &> \inf_{f \in \mathcal{H}} \mathcal{C}_\ell(f, \mathbf{x}, \frac{1}{2}), \text{ and} \\ \inf_{f \in \mathcal{H} : \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \mathcal{C}_\ell(f, \mathbf{x}, \eta) &> \inf_{f \in \mathcal{H}} \mathcal{C}_\ell(f, \mathbf{x}, \eta), \quad \forall \eta \in (\frac{1}{2}, 1], \text{ and} \\ \inf_{f \in \mathcal{H} : 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_\ell(f, \mathbf{x}, \eta) &> \inf_{f \in \mathcal{H}} \mathcal{C}_\ell(f, \mathbf{x}, \eta), \quad \forall \eta \in [0, \frac{1}{2}). \end{aligned} \quad (\text{C.9})$$

Lemma 4. (Awasthi et al. (2021), Theorem 23 and Theorem 24). Let \mathcal{H} be a symmetric hypothesis set consisting of the family of all measurable functions \mathcal{H}_{all} , ϕ be a non-increasing margin-based loss, and $\tilde{\phi}(f, \mathbf{x}, y) = \sup_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \phi(yf(\tilde{\mathbf{x}}))$. If $\tilde{\phi}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$, then $\tilde{\phi}$ is \mathcal{H} -consistent with respect to $\tilde{\ell}_0$ for all distributions \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ that satisfy: $\mathcal{R}_{\tilde{\ell}_0, \mathcal{H}}^* = 0$ and there exists $f^* \in \mathcal{H}$ such that $\mathcal{R}_\phi(f^*) = \mathcal{R}_{\phi, \mathcal{H}_{\text{all}}}^* < +\infty$.

The proofs of the above two Lemmas can be found in Awasthi et al. (2021).

Lemma 5. Let \mathcal{H} be a symmetric hypothesis set and $f \in \mathcal{H}$. Suppose $0 \leq \lambda^* < \hat{\lambda}^*$, $\nu > \min\{\hat{\lambda}^*, \mathcal{R}_{\ell, \mathcal{H}}^*\}$, $\ell(yf(\mathbf{x})) \geq 0 \forall \mathbf{x}$, and λ^* is bounded, then $\lambda^* < \ell(0)$.

Proof. Based on the definition of $(f_0^*, \lambda^*, \hat{\lambda}^*) = \arg \inf_{f, \lambda} \sup_{\hat{\lambda}} \left\{ \mathbb{E} \left[\hat{\lambda} - [\hat{\lambda} - [\ell(Yf(X)) - \lambda]_{+}]_{+} \right] + \nu \lambda - \mu \hat{\lambda} \right\}$. We choose $f = 0$, $\lambda = \ell(0)$ and $\hat{\lambda} = \hat{\lambda}^*$ there holds

$$\begin{aligned} \nu \lambda^* - \mu \hat{\lambda}^* &\leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf_0^*(X)) - \lambda^*]_{+}]_{+} \right] + \nu \lambda^* - \mu \hat{\lambda}^* \\ &\leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(0) - \ell(0)]_{+}]_{+} \right] + \nu \ell(0) - \mu \hat{\lambda}^* \\ &= \nu \ell(0) - \mu \hat{\lambda}^* \end{aligned}$$

Thus $\nu \lambda^* \leq \nu \ell(0)$ which shows that $\lambda^* \leq \ell(0)$. Let $\beta = \ell(0) - \lambda$ which implies

$$(f_0^*, \beta^*, \hat{\lambda}^*) = \arg \inf_{f, \lambda} \sup_{\hat{\lambda}} \left\{ \mathbb{E} \left[\hat{\lambda} - [\hat{\lambda} - [\ell(Yf(X)) + \beta - \ell(0)]_{+}]_{+} \right] - \nu \beta - \mu \hat{\lambda} \right\}.$$

Let $(f_0^*, \beta^*, \hat{\lambda}^*)$ be the minimizer. we have, for any f and choosing $\beta = \ell(0)$, that

$$\begin{aligned} -\nu \beta^* - \mu \hat{\lambda}^* &\leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf(X)) + \beta^* - \ell(0)]_{+}]_{+} \right] - \nu \beta^* - \mu \hat{\lambda}^* \\ &\leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf(X)) + \ell(0) - \ell(0)]_{+}]_{+} \right] - \nu \ell(0) - \mu \hat{\lambda}^*. \end{aligned}$$

Therefore, we have

$$-\nu \beta^* \leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf(X))]_{+}]_{+} \right] - \nu \ell(0).$$

Since f is arbitrary, $\beta^* \geq \frac{\nu - \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf^*(X))]_{+}]_{+} \right]}{\nu}$. Since $\ell(Yf^*(X)) \geq 0$, we have

$$0 \leq \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf^*(X))]_{+}]_{+} \right] = \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - \ell(Yf^*(X))]_{+} \right] \leq \min \left\{ \hat{\lambda}^*, \inf_f \mathbb{E}[\ell(yf(\mathbf{x}))] \right\}.$$

By using the assumption $\nu > \min\{\hat{\lambda}^*, \mathcal{R}_{\ell, \mathcal{H}}^*\} = \min \left\{ \hat{\lambda}^*, \inf_f \mathbb{E}[\ell(yf(\mathbf{x}))] \right\}$, we get $\beta^* \geq \frac{\nu - \mathbb{E} \left[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(Yf^*(X))]_{+}]_{+} \right]}{\nu} > 0$. Consequently, the above arguments show that $0 \leq \lambda^* = \ell(0) - \beta^* < \ell(0)$ if $\nu > \min \left\{ \hat{\lambda}^*, \inf_f \mathbb{E}[\ell(yf(\mathbf{x}))] \right\}$. \square

Theorem C.2. (Theorem 2 restated) Let \mathcal{H} be a symmetric hypothesis set consisting of the family of all measurable functions \mathcal{H}_{all} , suppose $\nu > \min\{\hat{\lambda}^*, \mathcal{R}_{\ell, \mathcal{H}}^*\}$, $0 \leq \lambda^* < \hat{\lambda}^*$, λ^* and $\hat{\lambda}^*$ are bounded, and ℓ is a non-negative, continuous, and non-increasing margin-based loss.

(i) Then $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$.

(ii) Furthermore, $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -consistent with respect to $\tilde{\ell}_0$ for all distributions \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ that satisfy: $\mathcal{R}_{\tilde{\ell}_0, \mathcal{H}}^* = 0$ and there exists $f^* \in \mathcal{H}$ such that $\mathcal{R}_{\phi_{\text{ORAT}}}(f^*) = \mathcal{R}_{\phi_{\text{ORAT}}, \mathcal{H}_{\text{all}}}^* < +\infty$.

Proof. Below we will prove the theorem using Lemma 3 which is from Awasthi et al. (2021). Recall that, from the definition of $\phi_{\text{ORAT}}(t)$ in Eq.(5), $\ell(t)$ is a continuous and non-increasing function, and λ^* and $\hat{\lambda}^*$ are bounded, we can conclude $\phi_{\text{ORAT}}(t)$ is bounded, continuous, non-increasing.

By Lemma 3, if $\mathcal{X}_2 = \emptyset$, $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$. Consequently, it suffices to consider the case where $\mathcal{X}_2 \neq \emptyset$. In this case, in order to show $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$, from Lemma 3 we only need to show, $\forall \mathbf{x} \in \mathcal{X}_2$, that

$$\begin{aligned} \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}) &> \inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}), \text{ and} \\ \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \eta) &> \inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \eta), \quad \forall \eta \in (\frac{1}{2}, 1], \text{ and} \\ \inf_{f \in \mathcal{H}: 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \eta) &> \inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \eta), \quad \forall \eta \in [0, \frac{1}{2}). \end{aligned}$$

To this end, recall that, by the definition of inner ℓ_s -risk, the inner $\tilde{\phi}_{\text{ORAT}}$ -risk is given by

$$\begin{aligned} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, x, \eta) &= \eta \tilde{\phi}_{\text{ORAT}}(f, \mathbf{x}, +1) + (1 - \eta) \tilde{\phi}_{\text{ORAT}}(f, \mathbf{x}, -1) \\ &= \eta \phi_{\text{ORAT}}\left(\inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} f(\tilde{\mathbf{x}})\right) + (1 - \eta) \phi_{\text{ORAT}}\left(\inf_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} -f(\tilde{\mathbf{x}})\right) \\ &= \eta \phi_{\text{ORAT}}\left(\underline{M}(f, \mathbf{x}, \epsilon)\right) + (1 - \eta) \phi_{\text{ORAT}}\left(-\overline{M}(f, \mathbf{x}, \epsilon)\right). \end{aligned}$$

For any $\mathbf{x} \in \mathcal{X}_2$, let $M_{\mathbf{x}} = \sup_{f \in \mathcal{H}} \underline{M}(f, \mathbf{x}, \epsilon) > 0$. Since \mathcal{H} is symmetric consisting of all measurable functions, we have $-M_{\mathbf{x}} = \inf_{f \in \mathcal{H}} \overline{M}(f, \mathbf{x}, \epsilon) < 0$. Since $\phi_{\text{ORAT}}(\cdot)$ is continuous, for any $\mathbf{x} \in \mathcal{X}_2$ and $\tau > 0$, there exists $f_{\mathbf{x}}^\tau \in \mathcal{H}$ such that $\phi_{\text{ORAT}}(\underline{M}(f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon)) < \phi_{\text{ORAT}}(M_{\mathbf{x}}) + \tau$, $\phi_{\text{ORAT}}(-\overline{M}(f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon)) < \phi_{\text{ORAT}}(0) + \tau$, $\overline{M}(f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon) \geq \underline{M}(f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon) > 0$, $\underline{M}(-f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon) \leq \overline{M}(-f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon) = -\underline{M}(f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon) < 0$, and $\phi_{\text{ORAT}}(\underline{M}(-f_{\mathbf{x}}^\tau, \mathbf{x}, \epsilon)) < \phi_{\text{ORAT}}(0) + \tau$. Next we analyze three cases:

1. When $\eta = \frac{1}{2}$, since ϕ_{ORAT} is non-increasing,

$$\begin{aligned} &\inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}) \\ &= \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \frac{1}{2} \phi_{\text{ORAT}}\left(\underline{M}(f, \mathbf{x}, \epsilon)\right) + \frac{1}{2} \phi_{\text{ORAT}}\left(-\overline{M}(f, \mathbf{x}, \epsilon)\right) \\ &\geq \frac{1}{2} \phi_{\text{ORAT}}(0) + \frac{1}{2} \phi_{\text{ORAT}}(0) \\ &= \phi_{\text{ORAT}}(0) \\ &= \hat{\lambda}^* - [\hat{\lambda}^* - [\ell(0) - \lambda^*]_+]_+. \end{aligned}$$

For any $\mathbf{x} \in \mathcal{X}_2$, there exists $f' \in \mathcal{H}$ such that $\underline{M}(f', \mathbf{x}, \epsilon) > 0$ and $-\overline{M}(f', \mathbf{x}, \epsilon) \leq -\underline{M}(f', \mathbf{x}, \epsilon) < 0$, we obtain

$$\mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f', \mathbf{x}, \frac{1}{2}) = \frac{1}{2} \phi_{\text{ORAT}}\left(\underline{M}(f', \mathbf{x}, \epsilon)\right) + \frac{1}{2} \phi_{\text{ORAT}}\left(-\overline{M}(f', \mathbf{x}, \epsilon)\right)$$

According to Lemma 5, we have $0 \leq \lambda^* < \hat{\lambda}^*$ and $\lambda^* < \ell(0)$. Therefore, we also analyze two cases:

(a) If $0 < \lambda^* + \hat{\lambda}^* \leq \ell(0)$, then we have

$$\inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}) \geq \hat{\lambda}^* - [\hat{\lambda}^* - [\ell(0) - \lambda^*]_+]_+ = \hat{\lambda}^*.$$

On the other hand, since ϕ_{ORAT} is continuous, there exists $f' \in \mathcal{H}$ and $t = \underline{M}(f', \mathbf{x}, \epsilon)$, then $0 \leq \phi_{\text{ORAT}}\left(\underline{M}(f', \mathbf{x}, \epsilon)\right) < \hat{\lambda}^*$. Thus,

$$\begin{aligned} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f', \mathbf{x}, \frac{1}{2}) &= \frac{1}{2} \phi_{\text{ORAT}}\left(\underline{M}(f', \mathbf{x}, \epsilon)\right) + \frac{1}{2} \phi_{\text{ORAT}}\left(-\overline{M}(f', \mathbf{x}, \epsilon)\right) \\ &\leq \frac{1}{2} \phi_{\text{ORAT}}\left(\underline{M}(f', \mathbf{x}, \epsilon)\right) + \frac{1}{2} \hat{\lambda}^* \\ &< \frac{1}{2} \hat{\lambda}^* + \frac{1}{2} \hat{\lambda}^* = \hat{\lambda}^*. \end{aligned}$$

Therefore, for any $\mathbf{x} \in \mathcal{X}_2$,

$$\inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}) \leq \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f', \mathbf{x}, \frac{1}{2}) < \hat{\lambda}^* \leq \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}). \quad (\text{C.10})$$

(b) If $\lambda^* + \hat{\lambda}^* > \ell(0)$,

$$\inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}^-(f, \mathbf{x}, \frac{1}{2}) \geq \hat{\lambda}^* - [\hat{\lambda}^* - [\ell(0) - \lambda^*]_+]_+ = \ell(0) - \lambda^*.$$

On the other hand, recall both $\phi_{\text{ORAT}}(\cdot)$ and $\ell(\cdot)$ are continuous and non-increasing and $\ell(0) > \lambda^*$ from Lemma 5. Therefore, we can find $f' \in \mathcal{H}$ such that $\ell(0) > \ell(\underline{M}(f', \mathbf{x}, \epsilon)) > \lambda^*$, $\lambda^* + \hat{\lambda}^* > \ell(-\overline{M}(f', \mathbf{x}, \epsilon)) > \ell(0) > \lambda^*$, and $\ell(\underline{M}(f', \mathbf{x}, \epsilon)) + \ell(-\overline{M}(f', \mathbf{x}, \epsilon)) < 2\ell(0)$. Consequently, there holds

$$\begin{aligned}
& \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f', \mathbf{x}, \frac{1}{2}) \\
&= \frac{1}{2}\phi_{\text{ORAT}}(\underline{M}(f', \mathbf{x}, \epsilon)) + \frac{1}{2}\phi_{\text{ORAT}}(-\overline{M}(f', \mathbf{x}, \epsilon)) \\
&= \frac{1}{2}[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(\underline{M}(f', \mathbf{x}, \epsilon)) - \lambda^*]_+]_+] + \frac{1}{2}[\hat{\lambda}^* - [\hat{\lambda}^* - [\ell(-\overline{M}(f', \mathbf{x}, \epsilon)) - \lambda^*]_+]_+] \\
&= \frac{1}{2}[\ell(\underline{M}(f', \mathbf{x}, \epsilon)) - \lambda^*] + \frac{1}{2}[\ell(-\overline{M}(f', \mathbf{x}, \epsilon)) - \lambda^*] \\
&= \frac{1}{2}[\ell(\underline{M}(f', \mathbf{x}, \epsilon)) + \ell(-\overline{M}(f', \mathbf{x}, \epsilon))] - \lambda^* \\
&< \frac{1}{2} \times 2\ell(0) - \lambda^* = \ell(0) - \lambda^*.
\end{aligned}$$

Therefore, for any $\mathbf{x} \in \mathcal{X}_2$,

$$\inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \frac{1}{2}) \leq \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f', \mathbf{x}, \frac{1}{2}) < \ell(0) - \lambda^* \leq \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0 \leq \overline{M}(f, \mathbf{x}, \epsilon)} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \frac{1}{2}). \quad (\text{C.11})$$

2. When $\eta \in (\frac{1}{2}, 1]$, since ϕ_{ORAT} is non-increasing, for any $\mathbf{x} \in \mathcal{X}_2$,

$$\begin{aligned}
\inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) &= \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \eta\phi_{\text{ORAT}}(\underline{M}(f, \mathbf{x}, \epsilon)) + (1 - \eta)\phi_{\text{ORAT}}(-\overline{M}(f, \mathbf{x}, \epsilon)) \\
&\geq \eta\phi_{\text{ORAT}}(0) + (1 - \eta)\phi_{\text{ORAT}}(M_{\mathbf{x}}).
\end{aligned}$$

On the other hand, for any $\mathbf{x} \in \mathcal{X}_2$ and $\tau > 0$,

$$\begin{aligned}
\mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) &= \eta\phi_{\text{ORAT}}(\underline{M}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) + (1 - \eta)\phi_{\text{ORAT}}(-\overline{M}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) \\
&< \eta[\phi_{\text{ORAT}}(M_{\mathbf{x}}) + \tau] + (1 - \eta)[\phi_{\text{ORAT}}(0) + \tau] \\
&= \eta\phi_{\text{ORAT}}(M_{\mathbf{x}}) + (1 - \eta)\phi_{\text{ORAT}}(0) + \tau.
\end{aligned}$$

Since $\eta > \frac{1}{2}$ and $M_{\mathbf{x}} > 0$, we have

$$\begin{aligned}
& \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) - \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) \\
&> [\eta\phi_{\text{ORAT}}(0) + (1 - \eta)\phi_{\text{ORAT}}(M_{\mathbf{x}})] - [\eta\phi_{\text{ORAT}}(M_{\mathbf{x}}) + (1 - \eta)\phi_{\text{ORAT}}(0) + \tau] \\
&= [2\eta - 1][\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})] - \tau \\
&> 0,
\end{aligned}$$

where we take $0 < \tau < [2\eta - 1][\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})]$. Therefore, for any $\eta \in (\frac{1}{2}, 1]$ and $\mathbf{x} \in \mathcal{X}_2$, there exists $0 < \tau < [2\eta - 1][\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})]$ such that

$$\inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) \leq \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) < \inf_{f \in \mathcal{H}: \underline{M}(f, \mathbf{x}, \epsilon) \leq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta). \quad (\text{C.12})$$

3. When $\eta \in [0, \frac{1}{2})$, since ϕ_{ORAT} is non-increasing, for any $\mathbf{x} \in \mathcal{X}_2$,

$$\begin{aligned}
\inf_{f \in \mathcal{H}: \overline{M}(f, \mathbf{x}, \epsilon) \geq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) &= \inf_{f \in \mathcal{H}: \overline{M}(f, \mathbf{x}, \epsilon) \geq 0} \eta\phi_{\text{ORAT}}(\underline{M}(f, \mathbf{x}, \epsilon)) + (1 - \eta)\phi_{\text{ORAT}}(-\overline{M}(f, \mathbf{x}, \epsilon)) \\
&\geq (1 - \eta)\phi_{\text{ORAT}}(0) + \inf_{f \in \mathcal{H}: \overline{M}(f, \mathbf{x}, \epsilon) \geq 0} \eta\phi_{\text{ORAT}}(\underline{M}(f, \mathbf{x}, \epsilon)) \\
&\geq (1 - \eta)\phi_{\text{ORAT}}(0) + \eta\phi_{\text{ORAT}}(M_{\mathbf{x}}).
\end{aligned}$$

On the other hand, for any $\mathbf{x} \in \mathcal{X}_2$ and $\tau > 0$,

$$\begin{aligned} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) &= \eta \phi_{\text{ORAT}}(\underline{M}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) + (1 - \eta) \phi_{\text{ORAT}}(-\overline{M}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) \\ &= \eta \phi_{\text{ORAT}}(\underline{M}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) + (1 - \eta) \phi_{\text{ORAT}}(\underline{M}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) \\ &< \eta[\phi_{\text{ORAT}}(0) + \tau] + (1 - \eta) \phi_{\text{ORAT}}(\underline{M}(f_{\mathbf{x}}^{\tau}, \mathbf{x}, \epsilon)) \\ &< \eta[\phi_{\text{ORAT}}(0) + \tau] + (1 - \eta)[\phi_{\text{ORAT}}(M_{\mathbf{x}}) + \tau] \\ &= \eta \phi_{\text{ORAT}}(0) + (1 - \eta) \phi_{\text{ORAT}}(M_{\mathbf{x}}) + \tau. \end{aligned}$$

Since $\eta < \frac{1}{2}$ and $M_{\mathbf{x}} > 0$, we have

$$\begin{aligned} &\inf_{f \in \mathcal{H}: \overline{M}(f, \mathbf{x}, \epsilon) \geq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) - \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) \\ &> [(1 - \eta) \phi_{\text{ORAT}}(0) + \eta \phi_{\text{ORAT}}(M_{\mathbf{x}})] - [\eta \phi_{\text{ORAT}}(0) + (1 - \eta) \phi_{\text{ORAT}}(M_{\mathbf{x}}) + \tau] \\ &= (1 - 2\eta)[\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})] - \tau, \end{aligned}$$

where we take $0 < \tau < (1 - 2\eta)[\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})]$. Therefore for any $\eta \in [0, \frac{1}{2})$ and $\mathbf{x} \in \mathcal{X}_2$, there exists $0 < \tau < (1 - 2\eta)[\phi_{\text{ORAT}}(0) - \phi_{\text{ORAT}}(M_{\mathbf{x}})]$ such that

$$\inf_{f \in \mathcal{H}} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) \leq \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(-f_{\mathbf{x}}^{\tau}, \mathbf{x}, \eta) < \inf_{f \in \mathcal{H}: \overline{M}(f, \mathbf{x}, \epsilon) \geq 0} \mathcal{C}_{\tilde{\phi}_{\text{ORAT}}}(f, \mathbf{x}, \eta) \quad (\text{C.13})$$

From (C.10), (C.11), (C.12), (C.13), we conclude that $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -calibrated with respect to $\tilde{\ell}_0$. Thus, (i) holds.

According to Lemma 4, we can conclude that the $\tilde{\phi}_{\text{ORAT}}$ is \mathcal{H} -consistent with respect to $\tilde{\ell}_0$ for all distributions \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ that satisfy: $\mathcal{R}_{\tilde{\ell}_0, \mathcal{H}}^* = 0$ and there exists $f^* \in \mathcal{H}$ such that $\mathcal{R}_{\phi_{\text{ORAT}}}(f^*) = \mathcal{R}_{\phi_{\text{ORAT}}, \mathcal{H}_{\text{all}}}^* < +\infty$. Therefore, (ii) holds. \square

C.3 CROSS-ENTROPY AS A MARGIN-BASED LOSS

The cross-entropy loss can be rewritten as a margin-based loss. For example, in binary classification, the conventional binary cross-entropy (bce) loss is given by $bce = -(y \log(\sigma(f(\mathbf{x}))) + (1 - y) \log(1 - \sigma(f(\mathbf{x}))))$ when $y = \{0, 1\}$. Here σ is the sigmoid function. It is clear that this conventional bce loss is not a margin-based loss. However, we can transfer the negative label 0 to -1. In this case, by the property of the sigmoid function $1 - \sigma(\mathbf{x}) = \sigma(-\mathbf{x})$, the original bce loss can be rewritten as $bce = -\log(\sigma(yf(\mathbf{x})))$ when $y = \{-1, 1\}$. This is in fact a non-negative, continuous, and non-increasing margin-based loss.

C.4 PROOF OF THEOREM 3

To get the generalization error bound, we need an equivalent formulation of equation 4 which is stated in the following lemma.

Lemma 6. Suppose $\lambda \in \mathbb{R}, \hat{\lambda} \in \mathbb{R}$, then the empirical risk $\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S})$ defined by Eq. equation 4 is equivalent to

$$\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) = \frac{1}{k - m} \left(\min_{\lambda \in \mathbb{R}} \left\{ k\lambda + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \right\} - \min_{\hat{\lambda} \in \mathbb{R}} \left\{ m\hat{\lambda} + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \hat{\lambda}]_+ \right\} \right) \quad (\text{C.14})$$

Proof. According to Eq.(C.3), we have

$$\begin{aligned} &\min_{\lambda} (k - m)\lambda + \max_{\hat{\lambda}} \left\{ (n - m)\hat{\lambda} - \sum_{i=1}^n [\hat{\lambda} - [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+] \right\} \\ &= \min_{\lambda} (k - m)\lambda + \sum_{i=m+1}^n [[\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+]_{[i]} \\ &= \min_{\lambda, q} (k - m)\lambda + \sum_{i=1}^n q_i [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \quad \text{s.t. } q_i \in [0, 1], \|q\|_0 = n - m. \end{aligned}$$

Under the constraints, we can rewrite the last formula as

$$\begin{aligned}
& (k-m)\lambda + \sum_{i=1}^n q_i [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \\
&= (k-m)\lambda + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ - \sum_{i=1}^n (1-q_i) [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ \\
&= k\lambda + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ - \sum_{i=1}^n (1-q_i) \{[\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ + \lambda\}.
\end{aligned}$$

The last equality holds because $\sum_{i=1}^n (1-q_i) = n - (n-m) = m$.

For the term $\sum_{i=1}^n (1-q_i) \{[\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ + \lambda\}$, we assume $\ell(f^*(\tilde{\mathbf{x}}_i), y_i), \forall i$, are sorted in descending order when getting the optimal model f^* . For example, $\ell(f^*(\tilde{\mathbf{x}}_1), y_1) \geq \ell(f^*(\tilde{\mathbf{x}}_2), y_2) \geq \dots \geq \ell(f^*(\tilde{\mathbf{x}}_n), y_n)$. Since $\lambda^* \geq 0$, the optimal q^* should be $q_1^* = \dots = q_m^* = 0, q_{m+1}^* = \dots = q_n^* = 1$. Note that λ^* must be an optimal solution of the problem

$$\min_{\lambda} (k-m)\lambda + \sum_{i=m+1}^n q_i^* [\ell(f^*(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+.$$

From Lemma 1, we know $\ell(f^*(\tilde{\mathbf{x}}_{m+1}), y_{m+1}) \geq \lambda^*$, which implies that $\ell(f^*(\tilde{\mathbf{x}}_i), y_i) - \lambda^* \geq 0$ holds for $q_i < 1$. Therefore, $\sum_{i=1}^n (1-q_i) \{[\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+ + \lambda\} = \sum_{i=1}^n (1-q_i) \ell(f(\tilde{\mathbf{x}}_i), y_i)$. Furthermore, we know

$$\min_{\hat{\lambda}} \left\{ m\hat{\lambda} + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \hat{\lambda}]_+ \right\} = \max_q \left\{ \sum_{i=1}^n (1-q_i) \ell(f(\tilde{\mathbf{x}}_i), y_i) \mid q_i \in [0, 1], \|q\|_0 = n-m \right\}.$$

Then we get

$$\begin{aligned}
& \frac{1}{k-m} \left(\min_{\lambda} (k-m)\lambda + \max_{\hat{\lambda}} \left\{ (n-m)\hat{\lambda} - \sum_{i=1}^n [\hat{\lambda} - [\ell(f(\tilde{\mathbf{x}}_i), y_i) - \lambda]_+] \right\} \right) \\
&= \frac{1}{k-m} \left(\min_{\lambda} \left\{ k\lambda + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_{[i]}), y_{[i]}) - \lambda]_+ \right\} - \min_{\hat{\lambda}} \left\{ m\hat{\lambda} + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_{[i]}), y_{[i]}) - \hat{\lambda}]_+ \right\} \right).
\end{aligned}$$

The proof is complete. \square

Considering the limit case of equation C.14, the population risk $\mathcal{R}_{\tilde{\ell}}(f)$ can be written as

$$\begin{aligned}
& \frac{1}{k-m} \left(\min_{\lambda \in \mathbb{R}} \left\{ k\lambda + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_{[i]}), y_{[i]}) - \lambda]_+ \right\} - \min_{\hat{\lambda} \in \mathbb{R}} \left\{ m\hat{\lambda} + \sum_{i=1}^n [\ell(f(\tilde{\mathbf{x}}_{[i]}), y_{[i]}) - \hat{\lambda}]_+ \right\} \right) \\
& \xrightarrow[n \rightarrow \infty]{\frac{k-m}{n} \rightarrow \nu, \frac{m}{n} \rightarrow \mu} \frac{1}{\nu} \left(\min_{\lambda \in \mathbb{R}} \left\{ (\nu+\mu)\lambda + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - \lambda]_+ \right\} - \min_{\hat{\lambda} \in \mathbb{R}} \left\{ \mu\hat{\lambda} + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - \hat{\lambda}]_+ \right\} \right) = \mathcal{R}_{\tilde{\ell}}(f).
\end{aligned}$$

The next Lemma tells us if the loss function is bounded, we can constrain the problem of $\mathcal{R}_{\tilde{\ell}}(f)$ and $\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S})$ in the bounded range as well.

Lemma 7. *Suppose that the range of ℓ is $[0, M]$. Then we have*

$$\mathcal{R}_{\tilde{\ell}}(f) = \frac{1}{\nu} \left(\min_{\lambda \in [0, M]} \left\{ (\nu+\mu)\lambda + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - \lambda]_+ \right\} - \min_{\hat{\lambda} \in [0, M]} \left\{ \mu\hat{\lambda} + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - \hat{\lambda}]_+ \right\} \right), \quad (\text{C.15})$$

and so does the empirical risk

$$\mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) = \frac{1}{k-m} \left(\min_{\lambda \in [0, M]} \left\{ k\lambda + \sum_{i=1}^n [\tilde{\ell}(f(\mathbf{x}_i), y) - \lambda]_+ \right\} - \min_{\hat{\lambda} \in [0, M]} \left\{ m\hat{\lambda} + \sum_{i=1}^n [\tilde{\ell}(f(\mathbf{x}_i), y) - \hat{\lambda}]_+ \right\} \right). \quad (\text{C.16})$$

Proof. The proof of equation C.16 is straight forward. By Lemma 1 and Lemma 2, we know $\lambda_{\mathcal{S}}^* = \tilde{\ell}(f(\mathbf{x}_{[k]}), y_{[k]})$ and $\hat{\lambda}_{\mathcal{S}}^* = \tilde{\ell}(f(\mathbf{x}_{[m]}), y_{[m]})$ are a pair of solution of equation C.16. Since $\tilde{\ell}(f(\mathbf{x}), y) = \max_{\tilde{\mathbf{x}} \in \mathcal{B}_\varepsilon(\mathbf{x})} \ell(f(\tilde{\mathbf{x}}), y) \in [0, M]$ for any \mathbf{x}, y , we have $\lambda_{\mathcal{S}}^*, \hat{\lambda}_{\mathcal{S}}^* \in [0, M]$.

Next we move on to equation C.15. Let λ^* and $\hat{\lambda}^*$ be a pair of solution of equation C.15. Let $\lambda = M$, then we have

$$\begin{aligned} (\nu + \mu)\lambda^* &\leq (\nu + \mu)\lambda^* + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - \lambda^*]_+ \\ &\leq (\nu + \mu)M + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) - M]_+ \\ &\leq (\nu + \mu)M + \mathbb{E}[M - M]_+ = (\nu + \mu)M, \end{aligned}$$

which implies $\lambda^* \leq M$. On the other hand, assume $\lambda^* = -\varepsilon$ for some $\varepsilon > 0$. Let $\lambda = 0$, then we have

$$-(\nu + \mu)\varepsilon + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y) + \varepsilon]_+ = (1 - (\nu + \mu))\varepsilon + \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y)]_+ \leq \mathbb{E}[\tilde{\ell}(f(\mathbf{x}), y)]_+,$$

which contradicts with $(1 - (\nu + \mu))\varepsilon > 0$. Therefore we have $\lambda^* \geq 0$. Similarly, we can show that $\hat{\lambda}^* \in [0, M]$. The proof is complete. \square

The next lemma shows the uniform convergence of learning with ORAT without using perturbation.

Lemma 8. *Suppose that the range of $\ell(f(\mathbf{x}), y)$ is $[0, M]$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of an i.i.d. training dataset of size n , the following holds for all $\ell_f \in \ell_{\mathcal{H}}$,*

$$\mathcal{R}_\ell(f) - \mathcal{R}_\ell(f; \mathcal{S}) \leq \frac{2}{\nu} \left(2\mathfrak{R}_n(\ell_{\mathcal{H}}) + \frac{M(2\sqrt{2} + 3\sqrt{\log(2/\delta)})}{\sqrt{2n}} \right).$$

Proof. By the subadditivity of max operator, for any $\ell_f \in \ell_{\mathcal{H}}$, we have

$$\begin{aligned} &\mathcal{R}_\ell(f) - \mathcal{R}_\ell(f; \mathcal{S}) \\ &= \frac{1}{\nu} \min_{\lambda \in [0, M]} \left\{ (\nu + \mu)\lambda + \mathbb{E}[\ell(f(\mathbf{x}), y) - \lambda]_+ \right\} - \frac{1}{\nu} \min_{\lambda \in [0, M]} \left\{ (\nu + \mu)\lambda + \frac{1}{n} \sum_{i=1}^n (\ell(f(\mathbf{x}_i), y_i) - \lambda)_+ \right\} \\ &\quad + \frac{1}{\nu} \min_{\hat{\lambda} \in [0, M]} \left\{ \mu\hat{\lambda} + \frac{1}{n} \sum_{i=1}^n (\ell(f(\mathbf{x}_i), y_i) - \hat{\lambda})_+ \right\} - \frac{1}{\nu} \min_{\hat{\lambda} \in [0, M]} \left\{ \mu\hat{\lambda} + \mathbb{E}[\ell(f(\mathbf{x}), y) - \hat{\lambda}]_+ \right\} \\ &\leq \max_{\lambda \in [0, M]} \left\{ \frac{1}{\nu} \mathbb{E}[\ell(f(\mathbf{x}), y) - \lambda]_+ - \frac{1}{n\nu} \sum_{i=1}^n (\ell(f(\mathbf{x}_i), y_i) - \lambda)_+ \right\} := L_1(f, \ell) \tag{C.17} \end{aligned}$$

$$+ \max_{\hat{\lambda} \in [0, M]} \left\{ \frac{1}{n\nu} \sum_{i=1}^n (\ell(f(\mathbf{x}_i), y_i) - \hat{\lambda})_+ - \frac{1}{\nu} \mathbb{E}[\ell(f(\mathbf{x}), y) - \hat{\lambda}]_+ \right\} := L_2(f, \ell). \tag{C.18}$$

Without loss of generality, we consider equation C.17, the bound for equation C.18 can be derived in a similar manner. Taking supremum on both sides, we have

$$\sup_{\ell_f \in \ell_{\mathcal{H}}} L_1(f, \ell) \leq \sup_{\ell_f \in \ell_{\mathcal{H}}, \lambda \in [0, M]} \left\{ \frac{1}{\nu} \mathbb{E}[\ell(f(\mathbf{x}), y) - \lambda]_+ - \frac{1}{n\nu} \sum_{i=1}^n (\ell(f(\mathbf{x}_i), y_i) - \lambda)_+ \right\} := \Phi(\mathcal{S}).$$

It is standard to verify that $\Phi(\mathcal{S})$ satisfies the bounded differences condition with parameter $\frac{M}{\nu}$ and one can apply McDiarmid's inequality (McDiarmid et al., 1989) so that with probability at least $1 - \delta/4$, there holds

$$\Phi(\mathcal{S}) \leq \mathbb{E}[\Phi(\mathcal{S})] + \frac{M}{\nu} \sqrt{\frac{\log(4/\delta)}{2n}}.$$

By further standard reduction from the expectation to Rademacher complexity (Theorem 3.3 Mohri et al. (2018)), with probability at least $1 - \delta/2$, there holds

$$\Phi(\mathcal{S}) \leq 2\mathfrak{R}_n\left(\frac{1}{\nu}(\mathcal{G})_+\right) + \frac{3M}{\nu} \sqrt{\frac{\log(4/\delta)}{2n}}, \tag{C.19}$$

where $\mathcal{G} = \{\ell_f - \lambda | \ell_f \in \ell_{\mathcal{H}}, \lambda \in [0, M]\}$ and $(\cdot)_+ = \max(\cdot, 0)$. Since the ramp function $(\cdot)_+$ is 1-Lipschitz and $(0)_+ = 0$, by Ledoux-Talagrand contraction inequality (Ledoux & Talagrand, 1991) we have

$$\begin{aligned} \mathfrak{R}_n\left(\frac{1}{\nu}(\mathcal{G})_+\right) &\leq \frac{1}{\nu}\mathfrak{R}_n(\mathcal{G}) = \frac{1}{\nu}\mathbb{E}_\sigma\left[\sup_{\ell_f \in \ell_{\mathcal{H}}, \lambda \in [0, M]} \left(\frac{1}{n}\sum_{i=1}^n \sigma_i \ell(f(\mathbf{x}_i), y_i) - \frac{1}{n}\sum_{i=1}^n \sigma_i \lambda\right)\right] \\ &\leq \frac{1}{\nu}\left(\mathbb{E}_\sigma\left[\sup_{\ell_f \in \ell_{\mathcal{H}}} \frac{1}{n}\sum_{i=1}^n \sigma_i \ell(f(\mathbf{x}_i), y_i)\right] + \mathbb{E}_\sigma\left[\sup_{\lambda \in [0, M]} \frac{1}{n}\sum_{i=1}^n \sigma_i \lambda\right]\right) \\ &\leq \frac{1}{\nu}\left(\mathfrak{R}_n(\ell_{\mathcal{H}}) + \frac{M}{n}\mathbb{E}_\sigma\left|\sum_{i=1}^n \sigma_i\right|\right) \\ &\leq \frac{1}{\nu}\left(\mathfrak{R}_n(\ell_{\mathcal{H}}) + \frac{M}{\sqrt{n}}\right), \end{aligned} \quad (\text{C.20})$$

where the last inequality follows by $\left(\mathbb{E}_\sigma\left[\sum_{i=1}^n \sigma_i\right]\right)^2 \leq \mathbb{E}_\sigma\left(\sum_{i=1}^n \sigma_i\right)^2 = n$. By putting equation C.20 into equation C.19, we have

$$\sup_{\ell_f \in \ell_{\mathcal{H}}} L_1(f, \ell) \leq \frac{1}{\nu}\left(2\mathfrak{R}_n(\ell_{\mathcal{H}}) + \frac{M(2\sqrt{2} + 3\sqrt{\log(4/\delta)})}{\sqrt{2n}}\right)$$

with probability at least $1 - \delta/2$. The lemma holds by noting $\sup_{\ell_f \in \ell_{\mathcal{H}}} \{R(f, \ell) - R_n(f, \ell)\} \leq \sup_{\ell_f \in \ell_{\mathcal{H}}} L_1(f, \ell) + \sup_{\ell_f \in \ell_{\mathcal{H}}} L_2(f, \ell)$. \square

The next corollary is straight-forward from Lemma 8 by replacing ℓ with $\tilde{\ell}$.

Corollary 1 (Theorem 3 restated). *Suppose that the range of $\ell(f(\mathbf{x}), y)$ is $[0, M]$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the draw of an i.i.d. training dataset of size n , the following holds for all $\ell_f \in \ell_{\mathcal{H}}$,*

$$R_{\tilde{\ell}}(f) - R_{\tilde{\ell}}(f; \mathcal{S}) \leq \frac{2}{\nu}\left(2\mathfrak{R}_n(\tilde{\ell}_{\mathcal{H}}) + \frac{M(2\sqrt{2} + 3\sqrt{\log(2/\delta)})}{\sqrt{2n}}\right).$$

C.5 EXAMPLES OF HYPOTHESIS SETS

We give two examples of hypothesis sets: linear classifiers and nonlinear neural networks, that satisfy the condition in Theorem 2 and 3. Suppose $\ell : \mathbb{R} \rightarrow [0, M]$ is monotonically non-increasing and L -Lipschitz continuous. In this case, the adversarial loss can be written (Yin et al., 2019) as $\tilde{\ell}(f_\theta(\mathbf{x}), y) := \max_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} \ell(f_\theta(\tilde{\mathbf{x}}), y) = \ell(\min_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} y f_\theta(\mathbf{x}))$. Therefore, given any function class \mathcal{H} , we can define the function class $\tilde{\mathcal{H}} \subseteq \mathbb{R}^{\mathcal{X} \times \{\pm 1\}}$, such that $\tilde{\mathcal{H}} = \{\min_{\tilde{\mathbf{x}} \in \mathcal{B}_\epsilon(\mathbf{x})} y f_\theta(\mathbf{x}) : f_\theta \in \mathcal{H}\}$. By the Ledoux-Talagrand contraction inequality (Ledoux & Talagrand, 1991), we have $\mathfrak{R}_n(\tilde{\ell}_{\mathcal{H}}) \leq L\mathfrak{R}_n(\tilde{\mathcal{H}})$. Hence we only need to characterize the Rademacher complexity of $\tilde{\mathcal{H}}$ given \mathcal{H} .

Linear Classifiers. Let the hypothesis set $\mathcal{H}_{lin} \subseteq \mathbb{R}^{\mathcal{X}}$ be a set of linear functions of $\mathbf{x} \in \mathcal{X}$. More specifically, we consider prediction vector θ with l_p ($p \geq 1$) norm constraint, i.e. $\mathcal{H}_{lin} = \{f_\theta(\mathbf{x}) = \theta^\top \mathbf{x} : \|\theta\|_p \leq r\}$. Then it is straight-forward to check \mathcal{H}_{lin} is a symmetric hypothesis set. Furthermore, let $\tilde{d} = d^{1-1/p}$. Awasthi et al. (2020) showed that $\max\left\{\mathfrak{R}_n(\mathcal{H}_{lin}), \epsilon r \frac{\max\{\tilde{d}, 1\}}{2\sqrt{2n}}\right\} \leq \mathfrak{R}_n(\tilde{\mathcal{H}}_{lin}) \leq \mathfrak{R}_n(\mathcal{H}_{lin}) + \epsilon r \frac{\max\{\tilde{d}, 1\}}{2\sqrt{n}}$. Combined with Theorem 3 with probability at least $1 - \delta$, we have $\mathcal{R}_{\tilde{\ell}}(f) - \mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) \leq \frac{2}{\nu}\left(2L\mathfrak{R}_n(\mathcal{H}_{lin}) + L\epsilon r \frac{\max\{\tilde{d}, 1\}}{2\sqrt{n}} + \frac{M(2\sqrt{2} + 3\sqrt{\log(2/\delta)})}{\sqrt{2n}}\right)$, where $\mathfrak{R}_n(\mathcal{H}_{lin})$ is given by classical result in Kakade et al. (2008).

Neural Networks. We consider feedforward neural networks with ReLU activation function ρ , i.e. $\rho(t) = \max\{0, t\}$. In particular, if the hypothesis set is one-layer neural networks defined as $\mathcal{H}_{one} = \{f_\theta(\mathbf{x}) = \theta_0^\top \rho(\Theta \mathbf{x}) : \|\theta_0\|_1 \leq r_0, \|\Theta_i\|_p \leq r\}$ where $\Theta_i \in \mathbb{R}^d$ is the i -th row of $\Theta \in \mathbb{R}^{d' \times d}$. This is a symmetric hypothesis set. Furthermore, the Rademacher complexity can be upper bounded

(Awasthi et al., 2020) as $\mathfrak{R}_n(\tilde{\mathcal{H}}_{one}) \leq \frac{rr_0 \max\{1, \frac{\tilde{d}(\|\mathbf{X}\|_\infty + \epsilon)\}}{\sqrt{n}}\}}{\sqrt{n}} (1 + \sqrt{d(d'+1)\log(36)})$, where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$. Combined with Theorem 3 with probability at least $1 - \delta$, we have $\mathcal{R}_{\tilde{\ell}}(f) - \mathcal{R}_{\tilde{\ell}}(f; \mathcal{S}) \leq \frac{2}{\nu} \left(\frac{2Lrr_0 \max\{1, \frac{\tilde{d}(\|\mathbf{X}\|_\infty + \epsilon)\}}{\sqrt{n}}\}}{\sqrt{n}} \times (1 + \sqrt{d(d'+1)\log(36)}) + \frac{M(2\sqrt{2}+3\sqrt{\log(2/\delta)})}{\sqrt{2n}} \right)$. Such bound implies the generalization error depends on the perturbation size ϵ , which demonstrates the intrinsic complexity of adversarial training.

D ADDITIONAL EXPERIMENTAL DETAILS

D.1 SOURCE CODE

For the purpose of review, the source code is accessible in the supplementary file.

D.2 SETTINGS OF NETWORKS AND COMPUTING INFRASTRUCTURE DESCRIPTION

For all networks, we training them by using (mini-batch) stochastic gradient descent with momentum 0.9, weight decay $2e-4$, batch size 128, epochs 50 (for LeNet) / 100 (for Small-CNN) / 100 (for ResNet-18), and initial learning rate 0.03 (for LeNet) / 0.1 (for Small-CNN) / 0.1 (for ResNet-18) which is divided by the 10 at 20-th and 40-th epoch for LeNet / 30-th and 60-th epoch for Small-CNN and ResNet-18.

All algorithms are implemented in Python 3.6 and trained and tested on an Intel(R) Xeon(R) CPU W5590 @3.33GHz with 48GB of RAM and an NVIDIA Quadro RTX 6000 GPU with 24GB memory.

D.3 TRAINING SETTINGS ON TOY EXAMPLES

In this section, we provide more details about how to generate synthetic datasets in Figure 1.

We generate two sets of 2D synthetic data (Figure 1). Each dataset contains 200 samples from Gaussian distributions with different means and variances. We consider both the case of the balanced (Figure 1 left) and the imbalanced (Figure 1 right) data distributions, in the former, the training data for the two classes are approximately equal while in the latter one class has a dominating number of samples in comparison to the other. In the balanced dataset (Figure 1 left), we create two outliers. One is in the blue class (shown as red \times), the other is in the red class (shown as blue \circ). In the imbalanced dataset, we create one outlier in the blue class (shown as red \times). For both datasets, the yellow squares around data samples represent the samples are perturbed within a ℓ_∞ ball.

For the balanced dataset (Figure 1 left), we build a simple network contains two linear layers and one ReLU layer (Nair & Hinton, 2010). The number of hidden units is three. For the imbalanced dataset (Figure 1 right), the network contains four linear layers and three ReLU layers. The number of hidden units is 20. We train these networks using SGD with 0.9 momentum for 3000 (balanced dataset) / 100,000 (imbalanced dataset) iterations with the learning rate of 0.02. We set $k = 20$ and $m = 2$ for balanced dataset, and $k = 20$ and $m = 1$ for imbalanced dataset when run our ORAT algorithm. In AT and ORAT, the training attack is PGD¹⁰ and we set the perturbation bound $\epsilon = 0.01$ and the PGD step size $\epsilon/4$.

D.4 DETAILS OF OUTLIERS GENERATION BY USING ASYMMETRIC NOISE

In asymmetric noise generation procedure, for MNIST, flipping $2 \rightarrow 7$, $3 \rightarrow 8$, $5 \leftrightarrow 6$ and $7 \rightarrow 1$; for CIFAR-10, flipping TRUCK \rightarrow AUTOMOBILE, BIRD \rightarrow AIRPLANE, DEER \rightarrow HORSE, CAT \leftrightarrow DOG; for CIFAR-100, the 100 classes are grouped into 20 super-classes with each having 5 sub-classes, then flipping between two randomly selected sub-classes within each super-class.

D.5 k AND m SETTINGS ON REAL DATASETS

We provide a reference for setting k and m to reproduce our ORAT experimental results (Table 1) on real datasets in Table D.1.

Noise	MNIST				CIFAR-10				CIFAR-100				
	$\epsilon=0.1$		$\epsilon=0.2$		$\epsilon=2/255$		$\epsilon=8/255$		$\epsilon=2/255$		$\epsilon=8/255$		
	k	m	k	m	k	m	k	m	k	m	k	m	
0	60000	1	60000	1	45000	1	45000	1	49950	1	49950	1	
Symmetric Noise	10%	59950	2000	60000	2000	50000	300	50000	500	50000	300	50000	500
	20%	59950	6000	60000	3000	50000	300	50000	300	50000	500	49800	500
	30%	59950	5000	60000	5000	50000	10	50000	200	49800	100	50000	500
	40%	59950	11000	60000	11000	50000	100	50000	50	49950	100	49950	500
Asymmetric Noise	10%	60000	100	60000	10	49950	300	50000	500	49900	100	49950	500
	20%	59950	100	59950	100	49950	500	50000	300	50000	100	50000	500
	30%	59950	10	60000	100	49950	200	50000	300	50000	100	50000	500
	40%	59950	10	60000	10	50000	450	50000	500	50000	100	49800	500

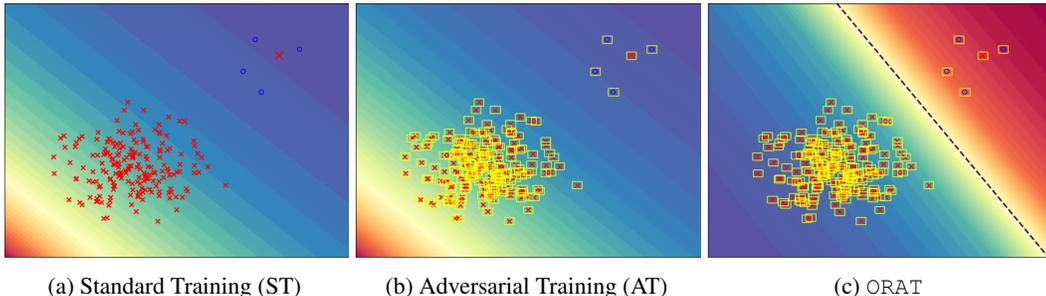
Table D.1: The k and m settings of ORAT on real datasets in different noise.

Figure E.1: An additional illustrative example of standard training (ST), adversarial training (AT), and ATRR for binary classification on an imbalanced synthetic dataset with one outlier (shown as red \times) in the blue class. The yellow squares around data samples represent the samples are perturbed within a ℓ_∞ ball. The dashed line is the decision boundary. The figure is better viewed in color.

E ADDITIONAL EXPERIMENTAL RESULTS

E.1 MORE EXPERIMENTS ON TOY EXAMPLE

We generate additional 2D synthetic data as shown in Figure E.1 to demonstrate the performance of our ORAT method. This imbalanced dataset contains 200 samples from Gaussian distribution with different means and variances. For this dataset, we create one outlier in the blue class (shown as red \times). In order to train this dataset, we build a network, which contains two linear layers and one ReLU layer. The number of hidden units is 64. We train this network using SGD with 0.9 momentum for 100,000 iterations with a learning rate of 0.1. We set $k = 5$ and $m = 1$ for ORAT. Similarly, in AT and ORAT, the training attack is PGD¹⁰, the perturbation bound $\epsilon = 0.01$, and the PGD step size is $\epsilon/4$.

From Figure E.1, we can find the classifiers are trained from ST (a) and AT (b) cannot separate two classes in the training data. However, we find the classifier is training by using our proposed ORAT can separate these two classes. The results demonstrate that our ORAT can eliminate the influence of the outliers when doing the adversarial training.

E.2 MORE EXPERIMENTS ON REAL DATASETS

In the main paper, we only show the tendency curves for MNIST when $\epsilon=0.1$ and CIFAR-10 and CIFAR-100 when $\epsilon=2/255$. In this section, we show more results on three datasets with 20% symmetric noise by setting a big value of ϵ in Figure E.2. Similar to the observations in Figure 2, we can find the losses are dramatically decreased in the first row of Figure E.2, which means Algorithm 1 can be successfully applied to solve ORAT optimization problem. From the second row of Figure E.2, it is obvious that the performance of our method is higher than the original AT approach on all attacks.

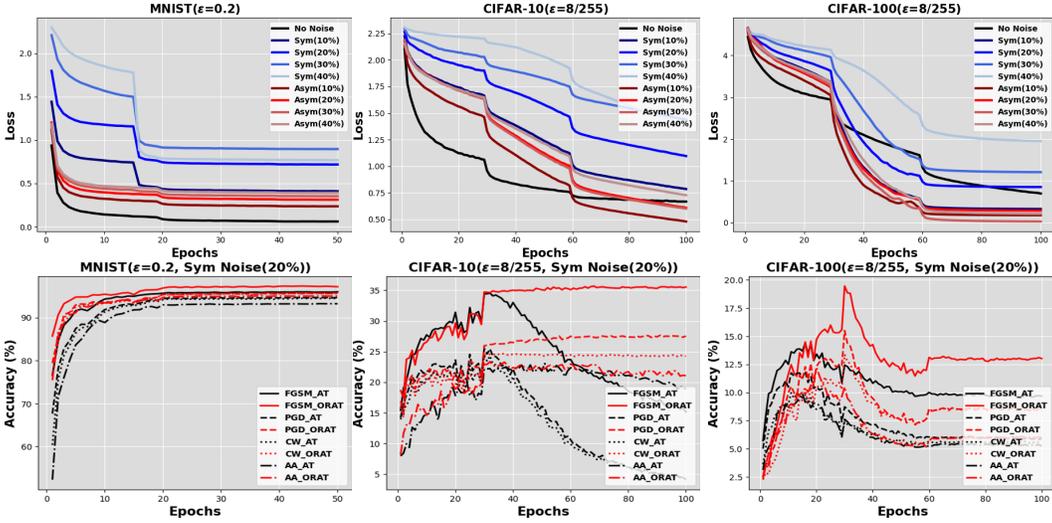


Figure E.2: The tendency curves of training adversarial loss and test accuracy on three datasets. The sharp drops in the curves correspond to decreases in training learning rate.

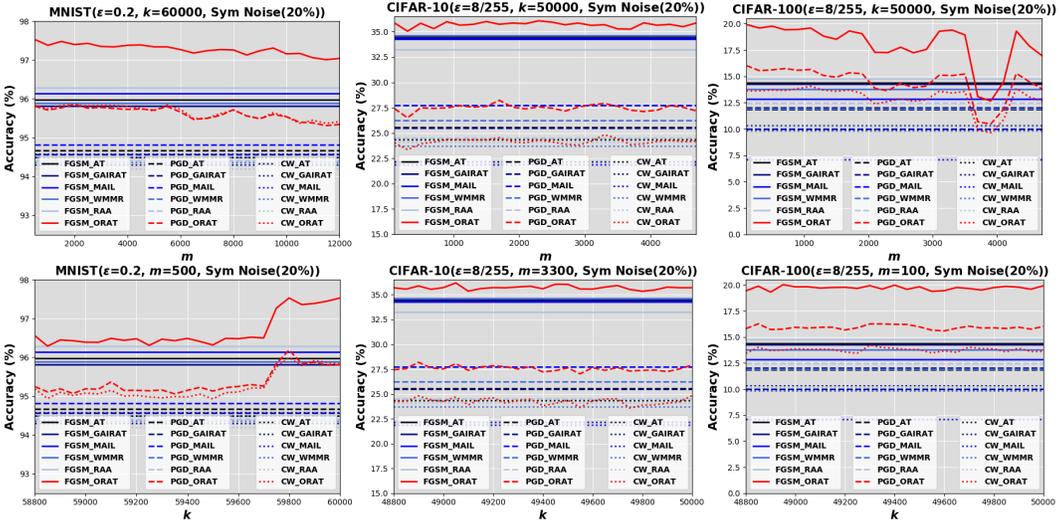


Figure E.3: Effect of k and m on the test accuracy of ORAT on three datasets.

E.3 MORE EXPERIMENTS ON THE EFFECT OF k AND m

We conduct more experiments to study the effect of hyperparameters k and m with using 20% symmetric noise on all datasets by setting a big value of ϵ . The results are shown in Figure E.3. Similar to the results that we get in Figure 3, we can see that there is a clear range of m with better performance than all compared methods. Fix m and test various k , we can find the performance can be improved by using some specific k values.

E.4 CONNECTION WITH ADVERSARIAL TRAINING ON OUT-OF-DISTRIBUTION PROBLEMS

Out-of-Distribution (OOD) problem exists due to the training and test data distributions mismatching (Hendrycks et al., 2021). Although the OOD problem setting is different from our outliers problem setting, some similarities exist between OOD data and outliers. For example, both of them are not from the data generating distribution. Therefore, whether the OOD methods can directly apply to solving our outlier problem in adversarial training is a question. Some works such as Zeng et al. (2021b); Varshney et al. (2022); Yi et al. (2021) connect adversarial robustness to out-of-distribution (OOD) problems.

Noise		Defense	CIFAR-100($\epsilon = 2/255$)			
			Na	FG	PGD	CW
Symmetric Noise	10%	pre-trained AT	24.20	17.09	14.83	14.17
		Ours	35.76	25.72	22.27	21.28
	20%	pre-trained AT	19.53	16.64	14.01	13.26
		Ours	34.45	25.07	22.21	20.92
	30%	pre-trained AT	19.41	16.22	13.06	13.04
		Ours	31.27	23.81	21.35	19.59
	40%	pre-trained AT	18.86	15.63	12.78	11.91
		Ours	29.38	22.99	20.85	19.20
Asymmetric Noise	10%	pre-trained AT	20.15	19.66	17.78	16.9
		Ours	37.09	27.07	23.65	22.59
	20%	pre-trained AT	24.60	17.43	16.67	15.15
		Ours	36.05	25.76	22.83	21.47
	30%	pre-trained AT	22.86	16.92	15.94	14.89
		Ours	34.58	24.18	21.05	20.11
	40%	pre-trained AT	21.58	16.18	15.32	14.48
		Ours	33.65	23.35	20.76	19.46

Table E.1: Testing accuracy (%) of pre-trained AT and our method (ORAT) on CIFAR-100 ($\epsilon = 2/255$) with different levels of symmetric and asymmetric noise. The best results are shown in bold.

Specifically, Zeng et al. (2021b) focuses on OOD detection. The problem in Zeng et al. (2021b) is that not enough labeled OOD samples can be used for training the OOD detection model. To improve the diversity of the unlabeled data augmentation, they apply an adversarial attack technique on unlabeled data to generate pseudo-positive samples. Then use these pseudo-positive samples with labeled data to improve the performance of the OOD detection model. However, their approach cannot directly apply to our setting since we only focus on supervised learning. All training data points are labeled in our setting, and the adversarial training works on labeled data. The authors in Varshney et al. (2022) test different selective prediction approaches for Natural Language Processing systems in in-domain, OOD, and adversarial settings. They regard several existing datasets as adversarial datasets for testing. However, no adversarial training approach is proposed and involved in Varshney et al. (2022). For Yi et al. (2021), the authors theoretically and experimentally show that a model (original AT Madry et al. (2018) or pre-trained AT (Salman et al., 2020)) robust to input perturbation generalizes well on OOD data.

Therefore, we test whether the pre-trained AT method (Salman et al., 2020) can solve outlier problems in adversarial training. Following the experimental setting from Yi et al. (2021), we download the ImageNet-based adversarially pre-trained robust ResNet-18 model in the setting of L_∞ and $\epsilon = 2/255$ from the public repository¹. Then fine-tune it on our noisy training datasets. We report pre-trained AT testing accuracy (%) on CIFAR-100 in Table E.1. To make the comparison explicit, we also attach our method performance. From Table E.1, we can find our method outperforms pre-trained AT under all settings. Most of the performance gaps between pre-trained AT and our method in Table E.1 are more than 5%. One reason is that the pre-trained AT is not designed to handle outliers. According to these results, it is clear that pre-trained AT cannot directly apply to solving our problem even if it has a good performance on OOD data.

E.5 EXTENSION OF TABLE 3

Self-learning (Han et al., 2019) is a useful strategy for learning model on noise data. For example, we can use AoRR to filter examples with larger loss (potential outliers), then conducting adversarial training on the cleaner set. We call this method AT w/o. However, it is not an end-to-end training approach. In contrast, our method is an end-to-end method, which means it is very easy to be conducted. To compare the effectiveness of ORAT and this AT w/o approach, we conduct experiments on MNIST with symmetric noise and CIFAR-100 with symmetric noise as follows.

In the first stage, for each dataset, we apply a grid search to select the values of k and m for training the model using the AoRR approach that can return a good testing accuracy. Then we use the trained

¹<https://github.com/microsoft/robust-models-transfer>

Noise	Defense	MNIST ($\epsilon = 0.1$)				CIFAR-100 ($\epsilon = 2/255$)				
		Na	FG	PGD	CW	Na	FG	PGD	CW	
Symmetric Noise	10%	AT w/o	98.91	98.08	97.61	97.55	29.81	21.09	19.59	18.23
		Ours	99.52	98.45	97.78	97.79	35.76	25.72	22.27	21.28
	20%	AT w/o	98.77	97.76	97.20	97.12	27.81	20.92	19.15	17.95
		Ours	99.56	98.37	97.65	97.64	34.45	25.07	22.21	20.92
	30%	AT w/o	97.82	96.97	96.47	96.35	24.18	19.17	17.31	16.71
		Ours	99.55	98.30	97.51	97.53	31.27	23.81	21.35	19.59
	40%	AT w/o	97.03	95.85	95.22	95.05	21.17	17.81	16.85	15.48
		Ours	99.36	98.00	97.22	97.20	29.38	22.99	20.85	19.20

Table E.2: Testing accuracy (%) of self-learning based method (Self-learning) and our method (ORAT) on MNIST ($\epsilon = 0.1$) and CIFAR-100 ($\epsilon = 2/255$) with different levels of symmetric noise. The best results are shown in bold.

model to test the loss for each sample from the training set. Therefore, we can obtain a training sample loss list. Next, we delete data points for the m largest losses in the training set to construct a clean set. This is because the AoRR uses m to determine how many examples (potential outliers) with the largest losses are ignored during each training epoch.

In the second stage, after we get a clean set, we use the conventional AT approach to train the model on the clean set and test the trained model on the testing set.

We report the testing accuracy (%) of the AT w/o approach on MNIST (symmetric noise, $\epsilon = 0.1$) and CIFAR-100 (symmetric noise, $\epsilon = 2/255$) in Table E.2. To make the comparison explicit, we also attach our method performance. From Table E.2, we can find our method outperforms the AT w/o approach under all settings. For example, the performance gap between the AT w/o approach and our method (ORAT) on MNIST can achieve more than 2% under the 40% symmetric noise setting. Most of the performance gaps on CIFAR-100 can achieve more than 4%.

One reason for low performance from the self-learning approach is that the training data points ignored by AoRR may contain clean data points. In this case, the constructed clean set is smaller than the original dataset. This may hurt the final model performance. Moreover, removing the examples with the largest losses before the adversarial training may lose the important feature information from the original training dataset. In other words, this compromises the richness and representational power of the data. In contrast, our ORAT method considers all examples during adversarial training. According to these results, it is clear that our approach (ORAT) gives a better solution than the self-learning approach for solving outlier problems in adversarial training either in the algorithm efficiency or effectiveness.

E.6 MORE ANALYSIS ON STABILITY OF ORAT

To evaluate the stability of each method, we report the the mean and standard deviation of testing accuracy (%) of all methods on MNIST (40% symmetric noise, $\epsilon = 0.1$) and CIFAR-100 (40% symmetric noise, $\epsilon = 2/255$) in Table 4. For each method, the reported performance is obtained by averaging the testing accuracy according to 10 random seeds. From Table 4, we can find our method still outperforms the compared methods in both datasets. For MNIST, our method can even outperform AT by more than 2%. Most importantly, we can find that the standard deviation in our method is less than or equal to that of other compared methods. For CIFAR-100, we can find the mean value of our method ORAT even higher than the reported performance in our submission. The standard deviation of the performance of our method differs from the comparison methods by at most 0.26% (compared to ST on FGSM attack). Comparing Table 4 and Table 1, it is clear that the performance gap becomes larger when we report scores by using mean and standard deviation, and our method shows a stable and stronger ability in handling outliers and adversarial attacks.