
Modeling Adversarial Noise for Adversarial Training

Dawei Zhou^{*1,2} Nannan Wang^{*1} Bo Han³ Tongliang Liu^{†2}

Abstract

Deep neural networks have been demonstrated to be vulnerable to adversarial noise, promoting the development of defense against adversarial attacks. Motivated by the fact that adversarial noise contains well-generalizing features and that the relationship between adversarial data and natural data can help infer natural data and make reliable predictions, in this paper, we study to model adversarial noise by learning the transition relationship between adversarial labels (i.e. the flipped labels used to generate adversarial data) and natural labels (i.e. the ground truth labels of the natural data). Specifically, we introduce an instance-dependent transition matrix to relate adversarial labels and natural labels, which can be seamlessly embedded with the target model (enabling us to model stronger adaptive adversarial noise). Empirical evaluations demonstrate that our method could effectively improve adversarial accuracy.

1. Introduction

Deep neural networks have been demonstrated to be vulnerable to adversarial noise (Goodfellow et al., 2015; Szegedy et al., 2014; Jin et al., 2019; Liao et al., 2018; Ma et al., 2018; Wu et al., 2020a). The vulnerability of deep neural networks seriously threatens many decision-critical deep learning applications (LeCun et al., 1998; He et al., 2016; Zagoruyko & Komodakis, 2016; Simonyan & Zisserman, 2015; Kaiming et al., 2017; Ma et al., 2021).

To alleviate the negative affects caused by adversarial noise, many adversarial defense methods have been proposed. A

^{*}Equal contribution ¹ISN Lab, School of Telecommunications Engineering, Xidian University, dwzhou.xidian@gmail.com, nnwang@xidian.edu.cn ²TML Lab, Sydney AI Centre, The University of Sydney ³Department of Computer Science, Hong Kong Baptist University, bhanml@comp.hkbu.edu.hk. Correspondence to: Tongliang Liu < tongliang.liu@sydney.edu.au >.

major class of adversarial defense methods focus on exploiting adversarial instances to help train the target model (Madry et al., 2018; Ding et al., 2019; Zhang et al., 2019; Wang et al., 2019), which achieve the state-of-the-art performance. However, these methods do not explicitly model the adversarial noise. The relationship between the adversarial data and natural data has not been well studied yet.

Studying (or modeling) the relationship between adversarial data and natural data is considered to be beneficial. If we can model the relationship, we can infer natural data information by exploiting the adversarial data and the relationship. Previous researches have made some explorations on this idea. Some data pre-processing based methods (Jin et al., 2019; Liao et al., 2018; Naseer et al., 2020; Zhou et al., 2021a;b), which try to recover the natural data by removing the adversarial noise, share the same philosophy. However, those methods suffer from the high dimensionality problem because both the adversarial data and natural data are high dimensional. The recovered data is likely to have *human-observable loss* (i.e., obvious inconsistency between processed instances and natural instances) (Xu et al., 2017) or contain residual adversarial noise (Liao et al., 2018).

To avoid the above problems, in this paper, we propose to model adversarial noise in the low-dimensional label space. Specifically, instead of directly modeling the relationship between the adversarial data and the natural data, we model the adversarial noise by learning the label transition from the *adversarial labels* (i.e., the flipped labels used to generate adversarial data) to the *natural labels* (i.e., the ground truth labels of the natural data). Note that the adversarial labels have not been well exploited in the community of adversarial learning, which guide the generation of adversarial noise and thus contain valuable information for modeling the well-generalizing features of the adversarial noise (Ilyas et al., 2019).

It is well-known that the adversarial noise depends on many factors, e.g., the data distribution, the adversarial attack strategy, and the target model. The proposed adversarial noise modeling method is capable to take into account of the factors because that the label transition is dependent of the adversarial instance (enabling us to model how the patterns of data distribution and adversarial noise affect label flipping) and that the transition can be seamlessly embed-

ded with the target model (enabling us to model stronger adaptive adversarial noise). Specifically, we employ a deep neural network to learn the complex instance-dependent label transition.

By using the transition relationship, we propose a defense method based on Modeling Adversarial Noise (called MAN). Specifically, we embed a label *transition network* (i.e., a matrix-valued transition function parameterized by a deep neural network) into the target model, which denotes the transition relationship from mixture labels (mixture of adversarial labels and natural labels) to natural labels (i.e., the ground truth labels of natural instances and adversarial instances). The transition matrix explicitly models adversarial noise and help us to infer natural labels. Considering that adversarial data can be adaptively generated, we conduct joint adversarial training on the target model and the transition network to achieve an optimal adversarial accuracy. We empirically show that the proposed MAN-based defense method could provide significant gains in the classification accuracy against adversarial attacks in comparison to state-of-the-art defense methods.

The rest of this paper is organized as follows. In Section 2, we introduce some preliminary information and briefly review related work on attacks and defenses. In Section 3, we discuss how to design an adversarial defense by modeling adversarial noise. Experimental results are provided in Section 4. Finally, we conclude this paper in Section 5.

2. Preliminaries

In this section, we first introduce some preliminary about notation and the problem setting. We then review the most relevant literature on adversarial attacks and adversarial defenses.

Notation. We use *capital* letters such as X and Y to represent random variables, and *lower-case* letters such as x and y to represent realizations of random variables X and Y , respectively. For norms, we denote by $\|x\|$ a generic norm. Specific examples of norms include $\|x\|_\infty$, the L_∞ -norm of x , and $\|x\|_2$, the L_2 -norm of x . Let $\mathbb{B}(x, \epsilon)$ represent the neighborhood of x : $\{\tilde{x} : \|\tilde{x} - x\| \leq \epsilon\}$, where ϵ is the perturbation budget. We define the *classification function* as $f : \mathcal{X} \rightarrow \{1, 2, \dots, C\}$, where \mathcal{X} is the feature space of X . It can be parametrized, e.g., by deep neural networks.

Problem setting. This paper focuses on a classification task under the adversarial environment, where adversarial attacks are utilized to craft adversarial noise to mislead the prediction of the target classification model. Let X and Y be the variables for natural instances and natural labels (i.e., the ground truth labels of natural instances) respectively. We sample natural data $\{(x_i, y_i)\}_{i=1}^n$ according to the distribution of the variables (X, Y) , where

$(X, Y) \in \mathcal{X} \times \{1, 2, \dots, C\}$. Given a deep learning model based classifier f and a pair of natural data (x, y) , the adversarial instance \tilde{x} satisfies the following constraint:

$$f(\tilde{x}) \neq y \quad \text{s.t.} \quad \|x - \tilde{x}\| \leq \epsilon. \quad (1)$$

The generation of the adversarial instance is guided by the adversarial label \tilde{y} which is different with the natural label y . Adversarial labels can be set by the attacker in target attacks, or be randomly initialized and then found by non-target attacks. Let \tilde{X} and \tilde{Y} be the variables for adversarial instances and adversarial labels respectively. We denote by $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$ the adversarial data drawn according to a distribution of the variables (\tilde{X}, \tilde{Y}) . Our aim is to design an adversarial defense to correct the adversarial label \tilde{y} into natural label y by modeling the relationship between adversarial data (i.e., \tilde{x} or \tilde{y}) and natural data (i.e., x or y).

Adversarial attacks. Adversarial instances are inputs maliciously designed by adversarial attacks to mislead deep neural networks (Szegedy et al., 2014). They are generated by adding imperceptible but adversarial noise to natural instances. Adversarial noise can be crafted by multi-step attacks following the direction of adversarial gradients, such as PGD (Madry et al., 2018) and AA (Croce & Hein, 2020). The adversarial noise is bounded by a small norm-ball $\|\cdot\|_p \leq \epsilon$, so that their adversarial instances can be perceptually similar to natural instances. Optimization-based attacks such as CW (Carlini & Wagner, 2017b) and DDN (Rony et al., 2019) jointly minimize the perturbation L_2 and a differentiable loss based on the logit output of a classifier. Some attacks such as FWA (Wu et al., 2020b) and STA (Xiao et al., 2018) focus on mimicking non-suspicious vandalism by exploiting the geometry and spatial information.

Adversarial defenses. The issue of adversarial instances promotes the development of adversarial defenses. A major class of adversarial defense methods is devoted to enhance the adversarial robustness in an adversarial training manner (Madry et al., 2018; Ding et al., 2019; Zhang et al., 2019; Wang et al., 2019). They augment training data with adversarial instances and use a min-max formulation to train the target model (Madry et al., 2018). However, these methods do not explicitly model the adversarial noise. The relationship between the adversarial data and natural data has not been well studied yet.

In addition, some data pre-processing based methods try to remove adversarial noise by learning a denoising function or a feature-squeezing function. For example, denoising based defenses (Liao et al., 2018; Naseer et al., 2020) transfer adversarial instances into clean instances, and feature squeezing based defenses (Guo et al., 2018) aim to reduce redundant but adversarial information. However, these methods suffer from the high dimensionality problem because both the adversarial data and natural data are high dimen-

sional. For examples, the recovered instances are likely to have significant inconsistency between the natural instances (Xu et al., 2017). Besides, the recovered instances may contain residual adversarial noise (Liao et al., 2018), which would be amplified in the high-level layers of the target model and mislead the final predictions. To avoid the above problems, we propose to model adversarial noise in the low-dimensional label space.

3. Modeling adversarial noise based defense

In this section, we present the *Modeling Adversarial Noise* (MAN) based defense method, to improve the adversarial accuracy against adversarial attacks. We first illustrate the motivation of the proposed defense method (Section 3.1). Next, we introduce how to model adversarial noise (Section 3.2). Finally, we present the training process of the proposed adversarial defense (Section 3.3). The code is available at <https://github.com/dwDavidxd/MAN>.

3.1. Motivation

Studying the relationship between adversarial data and natural data is considered to be beneficial for adversarial defense. If we can model the relationship, we can infer clean data information by exploiting the adversarial data and the relationship. Processing the input data to remove the adversarial noise is a representative strategy to estimate the relationship. However, data pre-processing based methods may suffer from high dimensionality problems. To avoid the problem, we would like to model adversarial noise in the low-dimensional label space.

Adversarial noise can be modeled because it has imperceptible but well-generalizing features. Specifically, classifiers are usually trained to solely maximize accuracy for natural data. They tend to use any available signal including those that are well-generalizing, yet brittle. Adversarial noise can arise as a result of perturbing these features (Ilyas et al., 2019), and it controls the flip from natural labels to adversarial labels. Thus, adversarial noise contains well-generalizing features which can be modeled by *learning the label transition from adversarial labels to natural labels*.

Note that the adversarial noise depends on many factors, such as data distribution, the adversarial attack strategy, and the target model. Modeling adversarial noise in label space is capable to take into account of these factors. Specifically, since that the label transition is dependent of the adversarial instance, we can model how the patterns of data distribution and adversarial noise affect label flipping, and exploit the modeling to correct the flipping of the adversarial label. In addition, we can model adversarial noise crafted by the stronger white-box adaptive attack, because the label transition can be seamlessly embedded with the target model. By

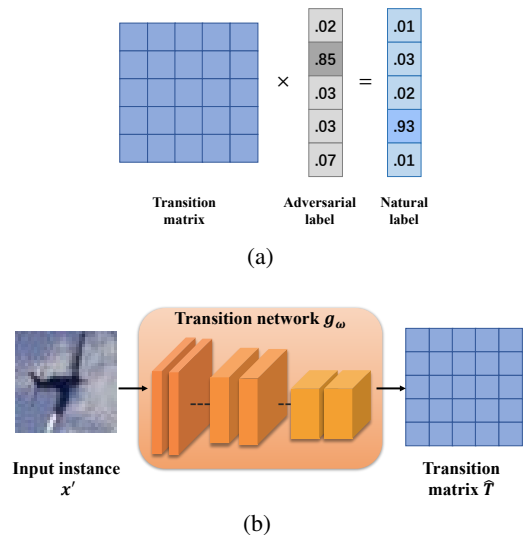


Figure 1. (a). Infer the natural label by utilizing the transition matrix and the adversarial label. (b). The transition network g_ω parameterized by ω takes the instance x' as input, and output the label transition matrix $\hat{T}(x'; \omega) = g_\omega(x')$.

jointly training the transition matrix with the target model, it can also be adaptive to adversarial attacks. We employ a deep neural network to learn the complex label transition from the well-generalizing and misleadingly predictive features of instance-dependent adversarial noise.

Motivated by the value of modeling adversarial noise for adversarial defense, we propose a defense method based on Modeling Adversarial Noise (MAN) by exploiting label transition relationship.

3.2. Modeling adversarial noise

We exploit a transition network to learn the label transition for modeling adversarial noise. The transition network can be regarded as a matrix-valued transition function parameterized by a deep neural network. The transition matrix, estimated by the label transition network, explicitly models adversarial noise and help us to infer natural labels from adversarial labels. The details are discussed below.

Transition matrix. To model the adversarial noise, we need to relate adversarial labels and natural labels in an explicit form (e.g., a matrix). Inspired by recent researches in label-noise learning (Xia et al., 2020; Liu & Tao, 2015; Xia et al., 2019; Yang et al., 2021; Wu et al., 2021; Xia et al., 2021), we design a label transition matrix, which can encode the probabilities that adversarial labels flip into natural labels. We then can infer natural labels by utilizing the transition matrix and adversarial data (see Figure 1(a)). Note that the transition matrix used in our work is different from that used in label-noise learning. The detailed illustration can be

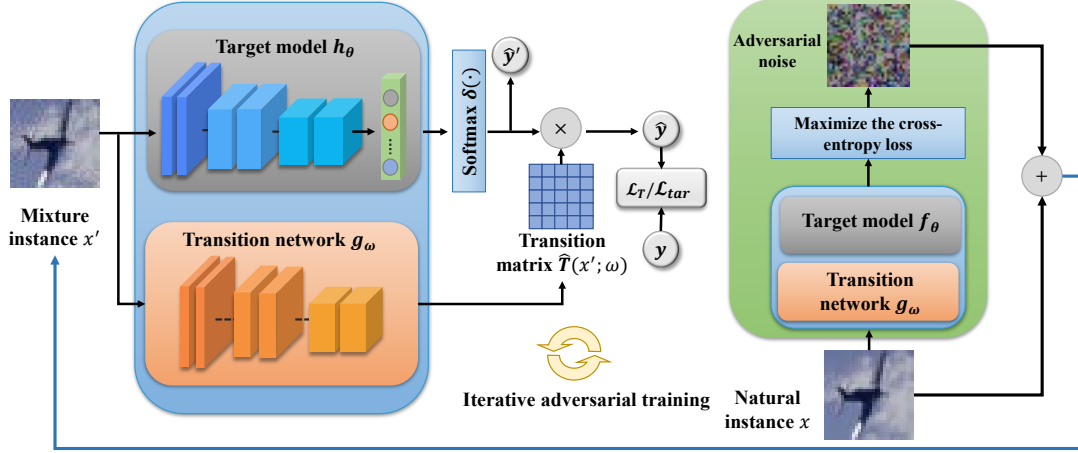


Figure 2. An overview of the training procedure for the proposed defense method. $\hat{\mathbf{y}}'$ and $\hat{\mathbf{y}}$ denote the probability of the estimated mixture label \hat{y}' and the probability of the inferred natural label \hat{y} respectively, i.e., $\hat{\mathbf{y}}' = \delta(h_\theta(x'))$ and $\hat{\mathbf{y}} = \hat{\mathbf{y}}' \cdot \hat{T}(x'; \omega)$. \mathbf{y} is y in the form of a vector.

found in Appendix A.

Considering that the adversarial noise is instance-dependent, the transition matrix should be designed to depend on input instances, that is, we should find a transition matrix specific to each input instance. In addition, we not only need to focus on the accuracy of adversarial instances, but also the accuracy of natural instances. An adversarially robust target model is often trained with both natural data and adversarial data. We therefore exploit the mixture of natural and adversarial instances (called mixture instances) as input instances, and design the transition matrix to model the relationship between the mixture of natural and adversarial labels (called mixture labels) to the ground-truth labels of mixture instances (called natural labels).

Specifically, let X' denote the variable for the mixture instances, Y' denote the variable for the mixture labels, and Y denote the variable for the natural labels. We combine the natural data and adversarial data as the mixture data, i.e., $\{(x'_i, y'_i)\}_{i=1}^{2n} = \{(x_i, y_i)\}_{i=1}^n \cup \{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$, where $\{(x'_i, y'_i)\}_{i=1}^{2n}$ is the data drawn according to a distribution of the random variables (X', Y') . Since we have the ground-truth labels of mixture instances (i.e., natural labels), we can extend the mixture data $\{(x'_i, y'_i)\}_{i=1}^{2n}$ into a triplet form $\{(x'_i, y'_i, y_i)\}_{i=1}^{2n}$. We utilize the transition matrix $T \in [0, 1]^{C \times C}$ to model the relationship between mixture labels Y' and natural labels Y . T is dependent of the mixture instances X' . It is defined as:

$$T_{i,j}(X' = x') = P(Y = j | Y' = i, X' = x'), \quad (2)$$

where $T_{i,j}$ denotes the (i, j) -th element of the matrix $T(X' = x')$. It indicates the probability of the mixture label i flipped to the natural label j for the input x' .

The basic idea is that given the *mixture class posterior probability* (i.e., the class posterior probability of the mixture labels Y') $P(\mathbf{Y}' | X' = x') = [P(Y' = 1 | X' = x'), \dots, P(Y' = C | X' = x')]^\top$, the *natural class posterior probability* (i.e., the class posterior probability of natural labels Y) $P(\mathbf{Y} | X' = x')$ could be inferred by exploiting $P(\mathbf{Y}' | X' = x')$ and $T(X' = x')$:

$$P(\mathbf{Y} | X' = x') = T(X' = x')^\top P(\mathbf{Y}' | X' = x'). \quad (3)$$

We then can obtain the natural labels by choose the class labels that maximize the robust class posterior probabilities. Note that the mixture class posterior probability $P(\mathbf{Y}' | X' = x')$ can be estimated by exploiting the mixture data.

Transition network. we employ a deep neural network (called *transition network*) to estimate the label transition matrix by exploiting the mixture data $\{(x'_i, y'_i, y_i)\}_{i=1}^{2n}$. Specifically, as shown in Figure 1(b), the transition network $g_\omega(\cdot)$ parameterized by ω takes the mixture instance x' as input, and output the label transition matrix $\hat{T}(x'; \omega) = g_\omega(x')$. We then can infer the natural class posterior probability $P(\mathbf{Y} | X' = x')$ according to Eq. 3, and thus infer the natural label. To optimize the parameter ω of the transition network, we minimize the difference between the inferred natural labels and the ground-truth natural labels. The loss function of the transition network is define as:

$$\mathcal{L}_T(\omega) = -\frac{1}{2n} \sum_{i=1}^{2n} \ell(\mathbf{y}'_i \cdot \hat{T}(x'_i; \omega), \mathbf{y}_i), \quad (4)$$

where $\mathbf{y}'_i, \mathbf{y}_i$ are y'_i, y_i in the form of vectors. $\ell(\cdot)$ is the cross-entropy loss between the inferred natural labels and the ground-truth natural labels, i.e., $\ell(\mathbf{y}'_i \cdot \hat{T}(x'_i; \omega), \mathbf{y}_i) = \mathbf{y}_i \cdot \log(\mathbf{y}'_i \cdot \hat{T}(x'_i; \omega))$.

3.3. Training

Considering that adversarial data can be adaptively generated, and the adversarial labels are also depended on the target model, we conduct joint adversarial training on the target model and the transition network to achieve the optimal adversarial accuracy. We provide the overview of the training procedure in Figure 2.

For the transition network, we optimize its model parameter ω according to Eq. 4. For the target model, considering that the final inferred natural class posterior probability is influenced by the target model, we also use the cross-entropy loss between the inferred natural labels and the ground-truth natural labels to optimize the parameter θ . The loss function for the target model h_θ is defined as:

$$\mathcal{L}_{tar}(\theta) = -\frac{1}{2n} \sum_{i=1}^{2n} [\mathbf{y}_i \cdot \log(\hat{\mathbf{y}}'_i \cdot \hat{T}(x'_i; \omega))], \quad (5)$$

$$\hat{\mathbf{y}}'_i = \delta(h_\theta(x'_i)),$$

where $\delta(\cdot)$ denote the softmax function. $\hat{\mathbf{y}}'_i$ denote the mixture label in the form of a vector predicted by the target model.

The details of the overall procedure are presented in Algorithm 1. Specifically, for each mini-batch $\mathcal{B} = \{x_i\}_{i=1}^m$ sampled from natural training set, we first generate adversarial instances $\tilde{\mathcal{B}} = \{\tilde{x}_i\}_{i=1}^m$ via a strong adversarial attack algorithm, and obtain the mixture mini-batch $\mathcal{B}' = \{x'_i\}_{i=1}^{2m}$ (i.e., mixture of \mathcal{B} and $\tilde{\mathcal{B}}$) with mixture labels $\{y'_i\}_{i=1}^{2m}$ (i.e., mixture of adversarial labels $\{\tilde{y}_i\}_{i=1}^m$ and natural labels $\{y_i\}_{i=1}^m$). Then, we input the mixture instances $\{x'_i\}_{i=1}^{2m}$ into the target model $h_\theta(\cdot)$ and output $\{\hat{\mathbf{y}}'_i\}_{i=1}^{2m}$. We also input the mixture instances $\{x'_i\}_{i=1}^{2m}$ into the transition network $g_\omega(\cdot)$ to estimate the label transition matrices $\{\hat{T}(x'_i; \omega)\}_{i=1}^{2m}$. We next infer the final prediction labels by exploiting $\{\hat{\mathbf{y}}'_i\}_{i=1}^{2m}$ and $\{\hat{T}(x'_i; \omega)\}_{i=1}^{2m}$. Finally, we compute the loss functions $\mathcal{L}_T(\omega)$ and $\mathcal{L}_{tar}(\theta)$, and update the parameters ω and θ . By iteratively conducting the procedures of adversarial instance generation and defense training, ω and θ are expected to be adversarially optimized.

To demonstrate the effectiveness of joint adversarial training, we conduct an ablation study by independently training the transition network with a fixed pre-trained target model. In addition, to prove that the improvement of our method is not mainly due to the introduction of more model parameters, we conduct an additional experiment by using an adversarially trained target model fused by two backbone networks as the baseline. The details could be found in Section 4.3.

4. Experiments

In this section, we first introduce the experiment setup in Section 4.1. Then, we evaluate the effectiveness of our de-

Algorithm 1 Training the defense model based on Modeling Adversarial Noise (MAN).

Input: Target model $h_\theta(\cdot)$ parameterized by θ , transition network $g_\omega(\cdot)$ parameterized by ω , batch size m , and the perturbation budget ϵ ;

- 1: **repeat**
- 2: Read mini-batch $\mathcal{B} = \{x_i\}_{i=1}^m$ from training set;
- 3: Craft adversarial instance $\{\tilde{x}_i\}_{i=1}^m$ at the given perturbation budget ϵ for each instance x_i in \mathcal{B} ;
- 4: Obtain the mixture mini-batch $\mathcal{B}' = \{x'_i\}_{i=1}^{2m}$ with mixture labels $\{y'_i\}_{i=1}^{2m}$;
- 5: **for** $i = 1$ to $2m$ (in parallel) **do**
- 6: Forward-pass x'_i through $h_\theta(\cdot)$ and obtain $\hat{\mathbf{y}}'_i$;
- 7: Forward-pass x'_i through $g_\omega(\cdot)$ to estimate $\hat{T}(x'_i; \omega)$;
- 8: Infer the final prediction label by exploiting $\hat{\mathbf{y}}'_i$ and $\hat{T}(x'_i; \omega)$;
- 9: **end for**
- 10: Calculate $\mathcal{L}_T(\omega)$ and $\mathcal{L}_{tar}(\theta)$ using Eq. 4 and Eq. 5;
- 11: Back-pass and update ω and θ ;
- 12: **until** training converged.

fense against representative and commonly used L_∞ norm and L_2 -norm adversarial attacks in Section 4.2. In addition, we conduct ablation studies in Section 4.3. Finally, we present that MAN is also suitable for detecting adversarial samples Section 4.4

4.1. Experiment setup

Datasets. We verify the effective of our defense method on two popular benchmark datasets, i.e., *CIFAR-10* (Krizhevsky et al., 2009) and *Tiny-ImageNet* (Wu et al., 2017). *CIFAR-10* has 10 classes of images including 50,000 training images and 10,000 test images. *Tiny-ImageNet* has 200 classes of images including 100,000 training images, 10,000 validation images and 10,000 test images. Images in the two datasets are all regarded as natural instances. All images are normalized into $[0,1]$, and are performed simple data augmentations in the training process, including random crop and random horizontal flip. For the target model, we mainly use ResNet-18 (He et al., 2016) for both *CIFAR-10* and *Tiny-ImageNet*.

Attack settings. Adversarial samples for evaluating defense models are crafted by applying state-of-the-art attacks. These attacks include L_∞ -norm PGD (Madry et al., 2018), L_∞ -norm AA (Croce & Hein, 2020), L_∞ -norm FWA (Wu et al., 2020b), L_2 -norm CW (Carlini & Wagner, 2017b) and L_2 -norm DDN (Rony et al., 2019). Among them, the AA attack algorithm integrates three non-target attacks and a target attack. Other attack algorithms belong to non-target attacks. The iteration number of PGD and FWA is set to

Table 1. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *CIFAR-10* and *Tiny-ImageNet*. The target model is ResNet-18. We show the most successful defense with **bold**.

Dataset	Defense	None	PGD-40	AA	FWA-40	CW ₂	DDN
CIFAR-10	AT	83.39	42.38	39.01	15.44	0.00	0.09
	MAN	82.72	44.83	39.43	29.53	43.17	10.63
	TRADES	80.70	46.29	42.71	20.54	0.00	0.06
	MAN_TRADES	80.34	48.65	44.40	29.13	1.46	0.31
	MART	78.21	50.23	43.96	25.56	0.02	0.07
	MAN_MART	77.83	50.95	44.42	31.23	1.53	0.47
Tiny-ImageNet	AT	48.40	17.35	11.27	10.29	0.00	0.29
	MAN	48.29	18.15	12.45	13.17	16.27	4.01
	TRADES	48.25	19.17	12.63	10.67	0.00	0.05
	MAN_TRADES	48.19	20.12	12.86	14.91	0.67	1.10
	MART	47.83	20.90	15.57	12.95	0.00	0.06
	MAN_MART	47.79	21.22	15.84	15.10	0.89	1.23

40 with step size 0.007. The iteration numbers of CW₂ and DDN are set to 200 and 40 respectively with step size 0.01. For *CIFAR-10* and *Tiny-ImageNet*, the perturbation budgets for L_2 -norm attacks and L_∞ -norm attacks are $\epsilon = 0.5$ and $8/255$ respectively.

Defense settings. We use three representative defense methods as the baselines: standard adversarial training method AT (Madry et al., 2018), optimized adversarial training methods TRADES (Zhang et al., 2019) and MART (Wang et al., 2019). For all baselines and our defense method, we use the L_∞ -norm non-target PGD-10 (i.e., PGD with iteration number of 10) with random start and step size $\epsilon/4$ to craft adversarial training data. The perturbation budget ϵ is set to $8/255$ for both *CIFAR-10* and *Tiny-ImageNet*. All the defense models are trained using SGD with momentum 0.9 and an initial learning rate of 0.1. For our defense method, we exploit the ResNet-18 as the transition network for both *CIFAR-10* and *Tiny-ImageNet*. Other detailed settings can be found in Appendix B.

4.2. Defense effectiveness

Defending against adaptive attacks. A powerful *adaptive attack* strategy has been proposed to break defense methods (Athalye et al., 2018; Carlini & Wagner, 2017a). In this case, the attacker can access the architecture and model parameters of both the target model and the defense model, and then can design specific attack algorithms. We study the following three adaptive attack scenarios for evaluating our defense method.

Scenario (i): disturb the final output. Considering that the models in baselines (i.e., the target models) are completely leaked to the attacker in the white-box setting, for fair comparison, we utilize white-box adversarial attacks against the combination of the target model and the transition matrix. Similar to attacks against baselines, the goal

of the non-target attack in this scenario is to maximize the distances between final predictions of our defense and the ground-truth natural labels. The adversarial instance \tilde{x} is crafted by solving the following optimization problem:

$$\begin{aligned} \max_{\tilde{x}} \mathcal{L}(\tilde{\mathbf{y}} \cdot \hat{T}(\tilde{x}; \omega), \mathbf{y}), \\ \text{subject to: } \|x - \tilde{x}\| \leq \epsilon, \end{aligned} \quad (6)$$

where $\tilde{\mathbf{y}} = \delta(h_\theta(\tilde{x}))$ and $\mathcal{L}(\cdot)$ denote the specific loss function used by each attack. Similarly, we can generate adversarial instances via the target attack. The details of the target attack are presented in Appendix C.1.

We combine this attack strategy with five representative adversarial attacks introduced in Section *Attack settings* to evaluate defenses. The average natural accuracy (i.e., the results in the third column) and the average adversarial accuracy of defenses are shown in Table 1.

The results show that our defense (i.e., MAN) achieves superior adversarial accuracy compared with AT. This presents that our defense is effective. Although our method has a slight drop in the natural accuracy (0.80%), it provides more gains for adversarial robustness (e.g., 5.78% against PGD-40 and 1.08% against stronger AA). In addition, our method achieves significant improvements against some more destructive attacks (e.g., the adversarial accuracy is increased from 15.44% to 29.53% against FWA-40 and from 0.09% to 10.63% against DDN). The standard deviation is shown in Appendix C.1. Besides, we evaluate the effectiveness of our defense method on other model architecture by using the VggNet-19 as the target model and the transition network. In addition, we evaluate the robustness performance of our defense method at a small batch-size (e.g., 128). These detailed results are also shown in Appendix C.1.

Note that the training mechanism in our method can be regarded as the standard adversarial training on the combination of the target network and the transition matrix, our

method thus is applicable to different adversarial training methods. To avoid the bias caused by different adversarial training methods, we apply the optimized adversarial training methods TRADES and MART to our method respectively (i.e., MAN_TRADES and MAN_MART). As shown in Table 1, the results show that our method can improve the adversarial accuracy. Although our method has a slight drop in the natural accuracy (e.g., 0.49% on *CIFAR-10* and 0.08% on *Tiny-ImageNet* for MART), it provides more gains for adversarial robustness (e.g., 1.43% and 1.53% against PGD-40 on *CIFAR-10* and *Tiny-ImageNet* respectively). Besides, we find that using TRADES and MART affects the improvement of defense effectiveness against L_2 -norm based attacks (e.g., CW₂ and DDN). We will further study and address this issue in the future work.

Scenario (ii): attack the transition matrix. In this scenario, we design an adversarial attack to destroy the crucial transition matrix of our defense method. If the transition matrix is destroyed, the defense would immediately become ineffective. Since the ground-truth transition matrix is not given, we use the target attack strategy to craft adversarial samples. We choose an anti-diagonal identity matrix (see Figure 5 in Appendix C.2) as an example of the target transition matrix in the target attack. The optimization goal is designed as:

$$\begin{aligned} \max_x & -\mathcal{L}_{mse}(\hat{T}(\tilde{x}; \omega), T^*), \\ \text{subject to: } & \|x - \tilde{x}\| \leq \epsilon, \end{aligned} \quad (7)$$

where \mathcal{L}_{mse} denote the *mean square error* loss.

We use L_∞ -norm PGD-40 with this target attack strategy to evaluate the MAN-based defense trained in scenario (i). The adversarial accuracy is 70.18% on *CIFAR-10*, which represents that our defense is effective against such type of adaptive target attack. This may be because that the attack in scenario (i) also tries to craft adversarial data by destroying the transition matrix to reduce the adversarial accuracy. The transition network adversarially trained with such adversarial data thus have good robustness against the attack designed for the transition matrix.

Scenario (iii): dual attack. We explore another adaptive attack scenario. In this scenario, the attack not only disturbs the final prediction labels, but also disturbs the output of the target model. We call such attack *dual adaptive attack*. The optimization goal of the non-target dual attack can be designed as:

$$\begin{aligned} \max_x & [\mathcal{L}(\tilde{\mathbf{y}} \cdot T(\tilde{x}; \omega), \mathbf{y}) + \mathcal{L}_{ce}(\tilde{\mathbf{y}}, \mathbf{y})], \\ \text{subject to: } & \|x - \tilde{x}\| \leq \epsilon, \end{aligned} \quad (8)$$

where $\tilde{\mathbf{y}} = \delta(h_\theta(\tilde{x}))$, $\mathcal{L}(\cdot)$ denote the specific loss function used by each adversarial attack and $\mathcal{L}_{ce}(\cdot)$ is the cross-entropy loss of the target model, i.e. $\mathcal{L}_{ce}(\tilde{\mathbf{y}}, \mathbf{y}) =$

Table 2. Adversarial accuracy (percentage) of our MAN-based defense against dual white-box adaptive attacks on *CIFAR-10*. The target model is ResNet-18.

Defense	None	PGD-40	AA	FWA-40	CW	DDN
MAN	82.72	70.63	45.13	64.09	45.21	38.16
Model _T	8.67	2.36	1.43	1.65	2.10	0.73

Table 3. Adversarial accuracy (percentage) of our MAN-based defense against general attacks on *CIFAR-10*. The target model is ResNet-18.

Defense	None	PGD-40	AA	FWA-40	CW	DDN
MAN	89.01	81.07	79.90	80.02	77.89	77.82
Model _T	88.98	0.00	0.00	0.00	0.00	0.00

$-\frac{1}{n} \sum_{i=1}^n \mathbf{y}_i \cdot \log(\tilde{\mathbf{y}})$. Similarly, we can generate adversarial instances via the target attack. The details of the target attack are presented in Appendix C.3.

We combine this dual attack strategy with five attacks to evaluate our defense model trained in *Scenario (i)*. As shown in Table 2, the results in the second row are the adversarial accuracy of our MAN-based defense. The results in the third row (Model_T) are the accuracy of the target model for adversarial instances. The results demonstrate that our defense method can help infer natural labels by using the transition matrix and adversarial labels, and can provide effective protection against multiple attacks in this scenario.

We note that the adversarial accuracy of MAN in Table 2 is higher than that in Table 1. This may be because, for some input instances, the dual attack mainly uses the gradient information of the attack loss against the target model, to generate adversarial noise for breaking the target model. Our MAN-based defense can model the adversarial noise (i.e., learn the transition relationship for the adversarial label predicted by the target model) and infer the final natural label. The final adversarial accuracy thus can be improved.

We just showed three examples for non-targeted and targeted attacks designed to our defense model in the above three scenarios. Note that more different attacks can be designed, which however is beyond the scope of our work in this paper.

Defending against general attacks. We also evaluate the effectiveness of the proposed defense method against general adversarial attacks. The general adversarial attacks usually only focus on disrupting the performance of the target model. We utilize five adversarial attacks to evaluate the proposed MAN-based defense method. To train the defense model, we use the adversarial instances crafted by non-target L_∞ -norm PGD-10 attack against the target model as the adversarial training data. The adversarial accuracy of our MAN-based defense and the adversarial accuracy of the target model (Model_T) are shown in Table 3. It can be seen that our defense method achieve a great defense effect

Table 4. Adversarial accuracy (percentage) of our defense method for different target models on *CIFAR-10*. MAN* denote the fine-tuned defense model on VGG.

Defense	PGD-40	AA	FWA-40	CW ₂	DDN
General adversarial attacks					
ResNet-MAN	81.07	79.90	80.02	77.89	77.82
VGG-MAN	71.93	68.62	68.19	70.91	70.76
WRN-MAN	72.08	69.13	68.83	71.12	70.93
Adaptive adversarial attacks					
ResNet-MAN	44.83	39.43	29.53	43.17	10.63
VGG-MAN	11.71	7.36	6.98	6.31	2.29
VGG-MAN*	33.26	30.30	13.42	10.07	3.83

against general adversarial attacks. Note that we only train the defense model in the proposed method (i.e., the transition network) and fix the model parameters of the target model in this experiment. In addition, to illustrate that our method does not utilize gradient masking, we conduct additional experiments listed in *Athalye et al. (2018)*. Detailed results can be found in the Appendix D

Defense transferability. The work in *Ilyas et al. (2019)* shows that any two target models are likely to learn similar mixture pattern. Therefore, modeling the adversarial noise which manipulate such features would apply to both target models for improving adversarial accuracy. We apply the MAN-based defense which is trained for the ResNet-18 (ResNet), to other naturally pre-trained target models, such as VGG-19 (VGG) (*Simonyan & Zisserman, 2015*) and Wide-ResNet 28×10 (WRN) (*Zagoruyko & Komodakis, 2016*) for evaluating the transferability of our defense method. Against general adversarial attacks, we deploy the defense model trained in Section *Defending against general attacks* on the VGG and WRN target models. In addition, against adaptive adversarial attacks, we deploy our defense model trained in Section *Scenario (i)* on VGG.

The performances on *CIFAR-10* are reported in Table 4. It can be seen that our defense method has a certain degree of transferability for providing cross-model protections against adversarial attacks. Moreover, for black-box target models whose model parameters and training manners cannot be accessed (e.g., the VGG target model), we can fine-tune the transition network in our defense method to further improve the defense effect against adaptive adversarial attacks. As shown by MAN* in Table 4, the fine-tuned defense model achieves a higher adversarial accuracy.

4.3. Ablation study

To demonstrate that compared with only training the transition network, the joint adversarial training on the target model and the transition network could achieve better defense effectiveness, we conduct an ablation study. Specifically, we use a naturally pre-trained ResNet-18 target model, and train the transition network independently (called MAN-). The model parameters of the target model are fixed in the




Figure 3. Ablation study by independently training the transition network.

training procedure. We use the adversarial instances crafted by the adaptive L_∞ -norm PGD-10 as adversarial training data. Compared with jointly trained MAN, the results shown in Section 4.3 demonstrated the joint adversarial training manner can provide positive gains.

In addition, to demonstrate that the improvement of our method has little relationship with the fact that the transition network introduces more model parameters, we conduct an experiment by using two parallel ResNet-18 as the target model of the baseline. We use the AT method to train the new target model. The result demonstrates that the gain of our method is not due to the increase of model parameters. The details and results can be found in Appendix E.

4.4. Detecting adversarial samples

Besides improving adversarial robustness, we find that the proposed MAN can be utilized to detect adversarial samples. By observing the distribution of element values on the diagonal of the generated transition matrix, we can discriminate whether the input is natural or adversarial. Specifically, we use the label predicted by the target model as the index to obtain the element value p from the diagonal of the transition matrix. If the value of the obtained element is larger than the values of other elements on the diagonal, the input sample is a natural sample, otherwise, it is an adversarial sample. Therefore, we take $1 - p$ as the probability that the input is adversarial. We use 10,000 test samples on *CIFAR-10* to compute the AUROC against L_∞ -norm PGD-40 and L_2 -norm DDN (see Figure 4). These results further validate the roles of the transition network and the transition matrix in defending against adversarial attacks.




Figure 4. AUROC of detecting adversarial samples.

5. Conclusion

Traditional adversarial defense methods typically focus on directly exploiting adversarial instances to remove adversarial noise or train an adversarially robust model. In this paper, motivated by that the relationship between adversarial data and natural data can help infer clean data from adversarial data, we study to model adversarial noise by learning the label transition relationship for improving adversarial accuracy. We propose a defense method based on Modeling Adversarial Noise (called MAN). Specifically, we embed a label transition matrix into the target model, which denote the transition relationship from adversarial labels to natural labels. The transition matrix explicitly models adversarial noise and help us infer natural labels. We design a transition network to generate the instance-independent transition matrix. Considering that adversarial data can be adaptively generated, we conduct joint adversarial training on the target model and the transition network to achieve an optimal adversarial accuracy. The empirical results demonstrate that our defense method can provide effective protection against white-box general attacks and adaptive attacks. Our work provides a new adversarial defense strategy for the community of adversarial learning. In future, we will further optimize the MAN-based defense method to improve its transferability and its performance when applied to other adversarial training methods. In addition, how to effectively learn the transition matrix for the dataset with more classes is also our future focus.

6. Acknowledgements

This work was supported in part by the National Key Research and Development Program of China under Grant 2018AAA0103202, in part by the National Natural Science Foundation of China under Grant 61922066, 61876142, 62036007 and 62006202, in part by the Technology Innovation Leading Program of Shaanxi under Grant 2022QFY01-15, in part by Open Research Projects of Zhejiang Lab under Grant 2021KG0AB01, in part by the RGC Early Career Scheme No. 22200720, in part by Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652, in part by Australian Research Council Projects DE-190101473, IC-190100031, and DP-220102121, in part by the Fundamental Research Funds for the Central Universities, and in part by the Innovation Fund of Xidian University. The authors thank the reviewers and the meta-reviewer for their helpful and constructive comments on this work. Thanks to Chaojian Yu for his important advice on Section *attack the transition matrix*.

References

- Athalye, A., Carlini, N., and Wagner, D. A. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Carlini, N. and Wagner, D. Magnet and "efficient defenses against adversarial attacks" are not robust to adversarial examples. *arXiv preprint arXiv:1711.08478*, 2017a.
- Carlini, N. and Wagner, D. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 39–57. IEEE, 2017b.
- Carmon, Y., Raghuathan, A., Schmidt, L., Duchi, J. C., and Liang, P. Unlabeled data improves adversarial robustness. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11190–11201, 2019.
- Croce, F. and Hein, M. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Ding, G. W., Lui, K. Y. C., Jin, X., Wang, L., and Huang, R. On the sensitivity of adversarial robustness to input data distributions. In *ICLR (Poster)*, 2019.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- Guo, C., Rana, M., Cissé, M., and van der Maaten, L. Countering adversarial images using input transformations. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., and Madry, A. Adversarial examples are not bugs, they are features. *arXiv preprint arXiv:1905.02175*, 2019.
- Jin, G., Shen, S., Zhang, D., Dai, F., and Zhang, Y. APE-GAN: adversarial perturbation elimination with GAN. In *International Conference on Acoustics, Speech and Signal Processing*, pp. 3842–3846, 2019.

- Kaiming, H., Georgia, G., Piotr, D., and Ross, G. Mask r-cnn. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, PP:1–1, 2017.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liao, F., Liang, M., Dong, Y., Pang, T., Hu, X., and Zhu, J. Defense against adversarial attacks using high-level representation guided denoiser. In *Conference on Computer Vision and Pattern Recognition*, pp. 1778–1787, 2018.
- Liu, T. and Tao, D. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2015.
- Ma, X., Li, B., Wang, Y., Erfani, S. M., Wijewickrema, S. N. R., Schoenebeck, G., Song, D., Houle, M. E., and Bailey, J. Characterizing adversarial subspaces using local intrinsic dimensionality. In *International Conference on Learning Representations*, 2018.
- Ma, X., Niu, Y., Gu, L., Wang, Y., Zhao, Y., Bailey, J., and Lu, F. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*, 110:107332, 2021.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations*, 2018.
- Naseer, M., Khan, S., Hayat, M., Khan, F. S., and Porikli, F. A self-supervised approach for adversarial robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 262–271, 2020.
- Rony, J., Hafemann, L. G., Oliveira, L. S., Ayed, I. B., Sabourin, R., and Granger, E. Decoupling direction and norm for efficient gradient-based L2 adversarial attacks and defenses. In *Conference on Computer Vision and Pattern Recognition*, pp. 4322–4330, 2019.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations*, 2015.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019.
- Wu, D., Xia, S.-T., and Wang, Y. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Wu, J., Zhang, Q., and Xu, G. Tiny imagenet challenge. *Technical Report*, 2017.
- Wu, K., Wang, A. H., and Yu, Y. Stronger and faster wasserstein adversarial attacks. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 10377–10387, 2020b.
- Wu, S., Xia, X., Liu, T., Han, B., Gong, M., Wang, N., Liu, H., and Niu, G. Class2simi: A noise reduction perspective on learning with noisy labels. In *International Conference on Machine Learning*, pp. 11285–11295. PMLR, 2021.
- Xia, X., Liu, T., Wang, N., Han, B., Gong, C., Niu, G., and Sugiyama, M. Are anchor points really indispensable in label-noise learning? *arXiv preprint arXiv:1906.00189*, 2019.
- Xia, X., Liu, T., Han, B., Wang, N., Gong, M., Liu, H., Niu, G., Tao, D., and Sugiyama, M. Part-dependent label noise: Towards instance-dependent label noise. *Advances in Neural Information Processing Systems*, 33, 2020.
- Xia, X., Liu, T., Han, B., Gong, C., Wang, N., Ge, Z., and Chang, Y. Robust early-learning: Hindering the memorization of noisy labels. In *International Conference on Learning Representations*, 2021.
- Xiao, C., Zhu, J., Li, B., He, W., Liu, M., and Song, D. Spatially transformed adversarial examples. In *6th International Conference on Learning Representations*, 2018.
- Xu, W., Evans, D., and Qi, Y. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.
- Yang, S., Yang, E., Han, B., Liu, Y., Xu, M., Niu, G., and Liu, T. Estimating instance-dependent label-noise transition matrix using dnns. *arXiv preprint arXiv:2105.13001*, 2021.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In Wilson, R. C., Hancock, E. R., and Smith, W. A. P. (eds.), *Proceedings of the British Machine Vision Conference 2016*, 2016.
- Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., and Jordan, M. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pp. 7472–7482. PMLR, 2019.

Zhou, D., Liu, T., Han, B., Wang, N., Peng, C., and Gao, X. Towards defending against adversarial examples via attack-invariant features. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 12835–12845, 2021a.

Zhou, D., Wang, N., Peng, C., Gao, X., Wang, X., Yu, J., and Liu, T. Removing adversarial noise in class activation feature space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7878–7887, 2021b.

Appendices

A. The difference of the transition matrix between our method and label-noise learning

In label-noise learning, the transition matrix is used to infer clean labels from given noisy labels. The transition matrix denotes the probabilities that **clean** labels flip into **noisy** labels (Xia et al., 2020; Liu & Tao, 2015; Yang et al., 2021). The label-noise learning methods utilize the transition matrix to correct the training loss on noisy data (i.e., the clean instance with the **noisy** label). Given the noisy class posterior probability, the clean class posterior probability can be obtained, i.e., $p(\mathbf{y}|x) = (T(x)^\top)^{-1}p(\mathbf{y}^*|x)$, where \mathbf{y}^* denotes the noisy label in the form of a vector. Differently, in our method, the transition matrix is used to infer natural labels from observed adversarial labels. The transition matrix denotes the probabilities that **adversarial** labels flip into **natural** labels. We utilize the transition matrix to help compute the training loss on the adversarial instance with the **natural** label. For the observed adversarial class posterior probability, the natural class posterior probability can be obtained, i.e., $p(\mathbf{y}|\tilde{x}) = T(\tilde{x})^\top p(\tilde{\mathbf{y}}|\tilde{x})$. In addition, to the best of our knowledge, our method for the first time utilizes the transition matrix to explicitly model adversarial noise for improving adversarial robustness.

B. Training settings

For all baselines and our defense method, we use the L_∞ -norm non-target PGD-10 (i.e., PGD with iteration number of 10) with random start and step size $\epsilon/4$ to craft adversarial training data. The perturbation budget ϵ is set to $8/255$ for both *CIFAR-10* and *Tiny-ImageNet*. All the defense models are trained using SGD with momentum 0.9 and an initial learning rate of 0.1. The weight decay is 2×10^{-4} for *CIFAR-10*, and is 5×10^{-4} for *Tiny-ImageNet*. The batch-size is set as 1024 to reduce time cost. For a fair comparison, we adjust the hyperparameter settings of the defense methods so that the natural accuracy is not severely compromised and then compare the robustness. The epoch number is set to 100. The learning rate is divided by 10 at the 75-th and 90-th epoch. We report the evaluation results of the *last* checkpoint rather than those of the *best* checkpoint.

C. Defending against adaptive attacks**C.1. Scenario (i): disturb the final output**

Table 5. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *CIFAR-10*. The target model is ResNet-18. We show the most successful defense with **bold**.

Defense	None	PGD-40	AA	FWA-40	CW ₂	DDN
AT	83.39±0.95	42.38±0.56	39.01±0.51	15.44±0.32	0.00±0.00	0.09±0.03
MAN	82.72±0.53	44.83±0.47	39.43±0.73	29.53±0.47	43.17±0.68	10.63±0.49
TRADES	80.70±0.63	46.29±0.59	42.71±0.49	20.54±0.47	0.00±0.00	0.06±0.01
MAN_TRADES	80.34±0.61	48.65±0.41	44.40±0.56	29.13±0.60	1.46±0.21	0.31±0.05
MART	78.21±0.65	50.23±0.70	43.96±0.67	25.56±0.61	0.02±0.00	0.07±0.01
MAN_MART	77.83±0.67	50.95±0.61	44.42±0.69	31.23±0.58	1.53±0.27	0.47±0.07

Table 6. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *Tiny-ImageNet*. The target model is ResNet-18. We show the most successful defense with **bold**.

Defense	None	PGD-40	AA	FWA-40	CW ₂	DDN
AT	48.40±0.68	17.35±0.59	11.27±0.53	10.29±0.47	0.00±0.00	0.29±0.03
MAN	48.29±0.57	18.15±0.51	12.45±0.67	13.17±0.69	16.27±0.71	4.01±0.19
TRADES	48.25±0.71	19.17±0.58	12.63±0.51	10.67±0.68	0.00±0.00	0.05±0.01
MAN_TRADES	48.19±0.62	20.12±0.49	12.86±0.59	14.91±0.47	0.67±0.12	1.10±0.15
MART	47.83±0.65	20.90±0.59	15.57±0.52	12.95±0.49	0.00±0.00	0.06±0.01
MAN_MART	47.79±0.59	21.22±0.60	15.84±0.47	15.10±0.39	0.89±0.16	1.23±0.21

In scenario (i), the target attack solve the following optimization problem:

$$\max_{\tilde{x}} -\mathcal{L}(\tilde{\mathbf{y}} \cdot \hat{T}(\tilde{x}; \omega), \mathbf{y}^*), \text{ subject to: } \|x - \tilde{x}\| \leq \epsilon, \tag{9}$$

where \mathbf{y}^* is the target label in the form of a vector set by the attacker. The adversarial accuracy of defense methods against white-box adaptive attacks on *CIFAR-10* and *Tiny-ImageNet* are shown in Table 5 and Table 6 respectively.

In addition, we evaluate the adversarial accuracy of defense methods against white-box adaptive attacks on *CIFAR-10* by using the VggNet-19 as the target model and the transition network. As shown in Table 7, our method still achieves better performance.

Moreover, we evaluate the robustness performance of our defense method at a small batch-size, such as 128. The initial learning rate is still 0.1. The results are shown in Table 8.

Table 7. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *CIFAR-10*. The target model is VggNet-19.

Defense	None	PGD-40	AA	FWA-40	CW ₂	DDN
AT	80.91±0.61	29.83±0.43	26.00±0.31	7.55±0.19	0.10±0.01	0.15±0.02
MAN	80.25±0.50	37.13±0.61	33.19±0.43	17.88±0.44	38.13±0.37	6.10±0.09

Table 8. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *CIFAR-10*. The target model is ResNet-18. The batch-size is 128.

Defense	None	PGD-40	AA
AT	84.92±0.49	46.47±0.47	43.55±0.43
MAN	84.70±0.41	48.06±0.50	44.67±0.51

C.2. Scenario (ii): attack the transition matrix

In scenario (ii), we design an adversarial attack to destroy the crucial transition matrix of our defense. Since the ground-truth transition matrix is not given, we use the target attack strategy to craft adversarial examples. We choose an anti-diagonal identity matrix as an example of the target transition matrix in the target attack. The target transition matrix T^* is shown in Figure 5. Note that there may be other attacks that can be designed, but this is beyond the scope of our work, and we would not explore further in this paper.

0	0			...				0	1	
0	0							0	1	0
		∴					0	1	0	
						0	1	0		
∴				0	1	0				∴
			0	1	0					
		0	1	0						
	0	1	0				∴			
0	1	0						0	0	
1	0			...				0	0	

Target transition matrix

Figure 5. Target transition matrix T^* for the target adversarial attack.

C.3. Scenario (iii): dual attack

In scenario (iii), the optimization goal of the target attack can be designed as:

$$\max_{\tilde{x}} [-\mathcal{L}(\tilde{y} \cdot T(\tilde{x}; \omega), \mathbf{y}^*) + \mathcal{L}_{ce}(\tilde{y}, \mathbf{y})], \text{ subject to: } \|x - \tilde{x}\| \leq \epsilon, \quad (10)$$

where \mathbf{y}^* is the target label in the form of a vector set by the attacker.

Note that we report the **last** adversarial accuracy (i.e., the adversarial accuracy of the last epoch models obtained during training) following the discussion in Carmon et al. (2019), instead of the **best** results.

D. The possibility of gradient masking

We consider five behaviors listed in Athalye et al. (2018) to identify the gradient masking. **(i) One-step attacks do not perform better than iterative attacks.** The accuracy against PGD-1 is 76.81% (vs 44.83% against PGD-40). **(ii) Black-box attacks are not better (on attack success rate) than white-box attacks.** We use ResNet-18 with standard AT as the target model to craft adversarial data. The accuracy against PGD-40/AA is 70.13%/67.30% (vs 44.83%/39.43% in white-box setting). **(iii) Unbounded attacks reach 100% success.** The accuracy against PGD-40 with $\epsilon = 255/255$ is 0%. **(iv) Random sampling does not find adversarial examples.** For samples that is not successfully attacked by PGD, we randomly sample 10^5 points within the ϵ -ball, and do not find adversarial data. **(v) Increasing distortion bound increases success.** The accuracy against PGD-40 with increasing ϵ (8/255, 16/255, 32/255 and 64/255) is 44.83%, 26.59%, 14.66% and 8.12%. These results show that our method does not use gradient masking.

E. The influence of more model parameters

We think that the improvement of the adversarial robustness has little relationship with the fact that the transition network introduces more model parameters. The transition network is only used to learn the transition matrix, and it does not directly learn the logit output for the instance to predict the class label. The main reason why our method can improve the robustness is that we explicitly model adversarial noise in the form of the transition matrix. We can use this transition matrix to infer the natural label. In addition, considering that the model parameters of our method are indeed larger than those of AT, we conduct the following experiment to compare AT and our method in the case of similar number of model parameters.

Considering that both the target model and the transition network in our method are mainly based on ResNet-18 on *CIFAR-10*, we use two parallel ResNet-18 as the target model to classify the input. We use the average of their logit outputs as the final logit output. We use the training mechanism of AT to train the new target model. The target model for our MAN-based defense method is still ResNet-18. In this way, The number of model parameters of our method is similar to that of AT. As shown in Table 9, the results show that our method still has higher adversarial accuracy. This demonstrates that the improvement of the adversarial robustness is not due to the increase of model parameters.

Table 9. Adversarial accuracy (percentage) of defense methods against white-box adaptive attacks on *CIFAR-10*. The defense models have a similar number of model parameters.

Defense	None	PGD-40	AA	FWA-40	CW ₂	DDN
AT	83.53±0.61	42.40±0.59	39.09±0.60	15.46±0.47	0.00±0.00	0.08±0.01
MAN	82.72±0.53	44.83±0.47	39.43±0.73	29.53±0.47	43.17±0.68	10.63±0.49