

Robust Ensemble Model Training via Random Layer Sampling Against Adversarial Attack

Hakmin Lee*¹

zpqlam12@kaist.ac.kr

Hong Joo Lee*¹

dlghdwn008@kaist.ac.kr

Seong Tae Kim²

seongtae.kim@tum.de

Yong Man Ro¹

ymro@kaist.ac.kr

¹ Image and Video System Lab

Korea Advanced Institute of Science
and Technology

Daejeon, Republic of Korea

² Computer Aided Medical Procedures

Technical University of Munich

Munich, Germany

Abstract

Deep neural networks have achieved substantial achievements in several computer vision areas, but have vulnerabilities that are often fooled by adversarial examples that are not recognized by humans. This is an important issue for security or medical applications. In this paper, we propose an ensemble model training framework with random layer sampling to improve the robustness of deep neural networks. In the proposed training framework, we generate various sampled model through the random layer sampling and update the weight of the sampled model. After the ensemble models are trained, it can hide the gradient efficiently and avoid the gradient-based attack by the random layer sampling method. To evaluate our proposed method, comprehensive and comparative experiments have been conducted on three datasets. Experimental results show that the proposed method improves the adversarial robustness.

1 Introduction

Recently, deep learning models have shown exceptionally good performance in various computer vision tasks such as image classification [0, 24], object detection [22, 23], semantic segmentation [8, 16]. However, several studies have revealed that deep learning methods are vulnerable in case images are intervened by small perturbations that are not even perceptible for human-beings. These perturbed images are called adversarial examples, and the procedure to create such adversarial examples is called adversarial attack.

The adversarial attack algorithms could be categorized into two approaches. The one is white-box attack and the other one is black-box attack. In the white-box attacks, the attacker could access to the model's parameters, while in the black-box attacks, the attacker could not access to the model's parameters. In this paper, we focus on handling white box attacks. In the white-box attacks, interruptions of visually imperceptible perturbations to original images could lead to erroneous results. The perturbations of these approaches are

generated based on gradients. The gradient-based adversarial attack methods could simply and effectively cause malfunctions of target networks. Hence, the presence of adversarial examples has brought up challenges of great importance for security-critical computer vision applications.

To enhance the deep learning model against adversarial examples, several methods have been proposed. Previous adversarial defense methods have been proposed: Randomization [20] to make the attack ineffective, denoising [6], ensemble training [20], and training the model with adversarial examples [4, 17]. It has been demonstrated that training with adversarial examples is the most effective way to improve the model robustness [4]. However, it takes large training time and impractical on large-scale datasets [28].

To increase the robustness of the deep network, in recent studies, randomization and obfuscation are reported as practically effective strategies to improve the adversarial robustness. The example of randomization is noise addition at different levels of the system [29], randomized lossy compression [6], random projections [27], and random feature sampling [5]. The key idea of these approaches is to hide the gradient of the network by randomization. The adversarial attack is blunted in white box situations, where the defender can effectively perturb the weight through randomness and make a difference when the attacker makes adversarial examples based on a specific weight. In other words, increasing the diversity of the networks plays an important role in defense.

One of the simple ways to increase the model diversity and improve the robustness is to generate ensemble models. It has been empirically observed by Smith et al. [13] that ensemble networks which are trained with different random initialization can be robust to adversarial examples. Pang et al. [20] proposed an ensemble model training method to promote the diversity among the predictions. They proposed the adaptive diversity promoting regularizer to make individual network predict orthogonally. Most of these ensemble-based approaches showed prominent in terms of adversarial robustness. However, these ensemble models require a large number of parameters to improve the adversarial robustness.

In this paper, we focus on tackling the problem of network parameter increase in the ensemble models when improving the adversarial robustness. To this end, we propose a novel ensemble model training framework with random layer sampling and group optimization strategy. In the proposed ensemble models training framework, the ensemble model set is defined with M sub-models. Each sub-model has same structure with L layers and they have different weights. From the model set, we sample the layers categorically through the proposed random layer sampling method. Then, we generate sampled models which have the same structure with sub-model. Each sampled model predicts sample outputs. Through the proposed random layer sampling method, it is possible to increase possible recombination cases exponentially by simply adding linear parameters. Also we train the ensemble models with group optimization to promote the ensemble diversity. Through the group optimization, we could predict diverse sample outputs that are robust to adversarial examples. In summary, there are two advantages of the proposed method. Firstly, we can generate various ensemble models combination only with a few number of sub-models. Secondly, since our proposed layer sampling method generates a sampled model randomly, we can effectively take a gradient ambiguity and avoid reproducible for the adversarial attack. Then, the sample outputs guarantee the diversity. Therefore, we could predict robust prediction against adversarial examples. The contributions of our paper can be summarized below:

- We propose a novel ensemble model training framework with random layer sampling. Through the proposed framework, we can effectively construct M^L ensemble models

with M sub-models which has L layers. Compared to conventional ensemble methods, our method can generate a large number of predictions with a few number of sub-models. Then, the diversity of the predictions is increased with the random layer sampling training framework and group optimization strategy.

- Our method effectively hides the full gradient of network and improves the adversarial robustness. It is mainly due to the reason that our method predicts various sample outputs from different model combinations. Experimental results show that our method effectively improves the adversarial robustness compared with other ensemble-based defense methods.

2 Related Work

2.1 Adversarial Attack Method

The deep neural networks have been shown to be highly vulnerable to adversarial examples. It was first discovered by Szegedy et al. [26]. Then, Goodfellow et al. [8] proposed Fast Gradient Sign Method (FGSM). It is a fast and single-step adversarial attack version of [26]. It performs a single step update on the original sample x along the direction of the gradient of a loss function. After that, Moosavi et al. designed the DeepFool attack [13] starting from the assumption that models are fully linear. Under this assumption, there is a polyhedron that can separate individual classes. Recently, more powerful and effective attacks including C&W [4], PGD [17], EAD [9] are proposed to fool the networks. As stronger attack methods are reported, the need for developing better defense methods is increased.

2.2 Adversarial Defense Method

To improve the model's robustness against adversarial examples, several methods have been reported. There are many approaches including distillation-based approaches [21], adversarial training approaches [11, 12], and ensemble training approaches [10, 14, 15, 20]. In this section, we mainly describe the ensemble training approaches.

Pang et al. [20] improves the adversarial robustness by promoting ensemble diversity with Adaptive Diversity Promoting (ADP) regularization. They trained ensemble model through ADP regularization method. Through the ensemble diversity, the non-maximal predictions of each model are mutually orthogonal, then predicts robust output. Smith et al. [25] proposed ensemble models training method with predictive uncertainty estimation and detecting adversarial example. By quantifying the predictive uncertainty, they could optimize the ensemble models that are robust to adversarial examples. It could also detect adversarial examples with predictive uncertainty. Although these ensemble model training methods could improve the adversarial robustness, to improve the adversarial robustness, a large number of parameters are required. Another way to construct ensemble models is to add the noise in the layer or dropout the weight. The noise addition or dropout method can be interpreted as ensemble model combination [13] where the predictions are averaged over an ensemble of neurons. Liu et al. [15] reported a random self-ensemble (RSE) training for improving the robustness of deep neural models. They add a noise layer before each convolution layer in both training and prediction phases. They showed that the algorithm is equivalent to ensemble a huge amount of noisy models together, and ensure that the ensemble model can generalize well. They further prove the fact that the proposed method is equivalent to adding

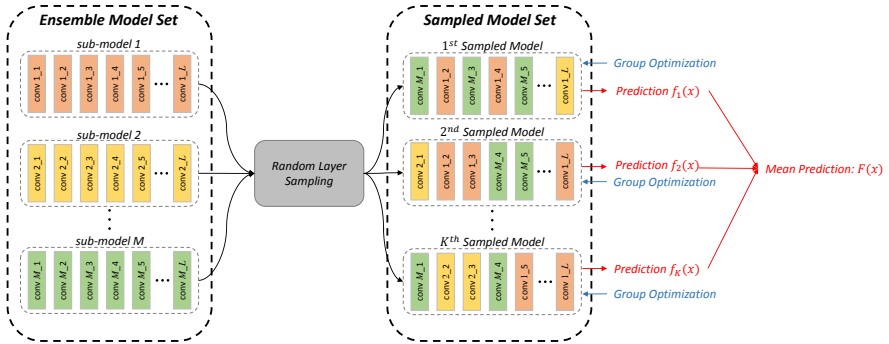


Figure 1: Overall procedure of the proposed ensemble model training framework with random layer sampling method. The ensemble model set consists of M sub-models. Each sub-model has L layers with different weights. From the ensemble model set, we sample each layer and construct sampled model. Then, we construct K different sampled models and update each sampled model parameter at once.

a Lipchitz regularization and thus can improve the robustness of neural models. Dhillon et al. [10] proposed stochastic activation pruning for robust adversarial defense. During the forward pass, they prune a subset of the activations in each layer. Then, they scale up the remaining activations to normalize the dynamic range of the inputs to the subsequent layer.

3 Proposed Method

In this section, we describe our proposed ensemble model training framework. Figure 1 shows an overview of the proposed ensemble model training framework with random layer sampling and group optimization. As shown in the figure, we construct ensemble model set with M duplicate sub-models. The sub-model consists of L layers, and weight of each layer is differently initialized over sub-models. To train the ensemble model set, we sample the layer from the ensemble model set by the proposed random layer sampling method. Then, we construct K sampled models. The sampled model predicts sample output $f_k(x)$. From the sampled models, we optimize the sampled models at once with sample mean prediction to increase the diversity between sampled models. In the testing phase, we construct N different ensemble models by the proposed random layer sampling and decide a final decision by averaging predictions of N sampled models. The details of the proposed random layer sampling method and training framework will be described in the following subsections.

3.1 Random Layer Sampling

In this section, we describe how to sample the layer and generate a sampled model. As shown in Figure 1, we design an ensemble model set with M duplicate sub-models which consists of L layers. Each sub-model has the same structure while having different weights. From the ensemble model set, we sample the layers by our proposed Random Layer Sampling (RLS) method. In the proposed RLS method, we sample layers with $r_l(m)$ where l denotes the layer index, m denotes the sub-model index, and $r_l(m)$ denotes the categorical random

variable that indicates whether each layer is selected or not. In order to ensure that each layer is selected from one sub-model, we constrain the $r_l(m)$ as follows

$$\sum_{m=1}^M r_l(m) = 1, \quad (1)$$

$$r_l \sim \text{categorical}(x_1, x_2, x_3, \dots, x_M; \mu_1, \mu_2, \mu_3, \dots, \mu_M). \quad (2)$$

Therefore, the feed-forward operation can be described as

$$w_l = r_l \otimes \mathbf{W}_l, \quad (3)$$

$$z^{(l+1)} = w_l * y^{(l)}, \quad (4)$$

$$y^{(l+1)} = \sigma(z^{(l+1)}), \quad (5)$$

where μ_M denotes probability of each category separately specified, \mathbf{W}_l denotes a set of weights in l^{th} convolution layer, w_l denotes a sampled weight from categorical random variables, \otimes denotes element-wise multiplication, $*$ denotes a convolution operator, y^l denotes a feature vector of l^{th} layer, $z^{(l+1)}$ denotes the output vector of l^{th} layer, and σ is an activation function. Eq. 3 denotes that we sample only one weight of the layer in the M sub-models. By stacking these randomly sampled layers, we can construct sampled models which have the same structure but have different weights. Through the RLS, we can generate various ensemble models that have different weight effectively. Theoretically, we can generate M^L different ensemble models only with the M sub-models.

3.2 Training for Adversarial Robustness with Ensemble Diversity

It is widely known that ensemble of several individual models could improve the adversarial robustness [1, 2, 3, 4]. For the adversarial robustness, the diversity among individual sub-models should be sufficiently guaranteed. To guarantee the diversity, we trained the ensemble models by group optimization strategy. As shown in Figure 1, let $f_k(x) = p(y | x, \hat{w}^k)$ be a prediction of the sampled model where \hat{w}^k denotes the sampled weights, and $F(x) = \frac{1}{K} \sum f_k(x)$ be a mean prediction of the sampled models. Following the description of [5, 6], the ensemble model diversity can be defined as

$$\alpha(f_k | x) = (f_k(x) - F(x))^2. \quad (6)$$

Therefore, the diversity of the ensemble model can be defined as the difference between individual sub-model prediction and mean prediction of sampled model. If we set the difference between ground-truth and sample output as mean square error, it can be represented as $MSE(f_k | x) = (y - f_k(x))^2$. The mean square error can be decomposed into

$$\mathbb{E}[MSE(F | x)] = \mathbb{E}[\overline{MSE}(f | x)] - \mathbb{E}[\bar{\alpha}(f | x)], \quad (7)$$

where

$$\overline{MSE}(f | x) = \frac{1}{K} \sum_{k=1}^K MSE(f_k | x), \text{ and } \bar{\alpha}(f | x) = \frac{1}{K} \sum_{k=1}^K \alpha(f_k | x). \quad (8)$$

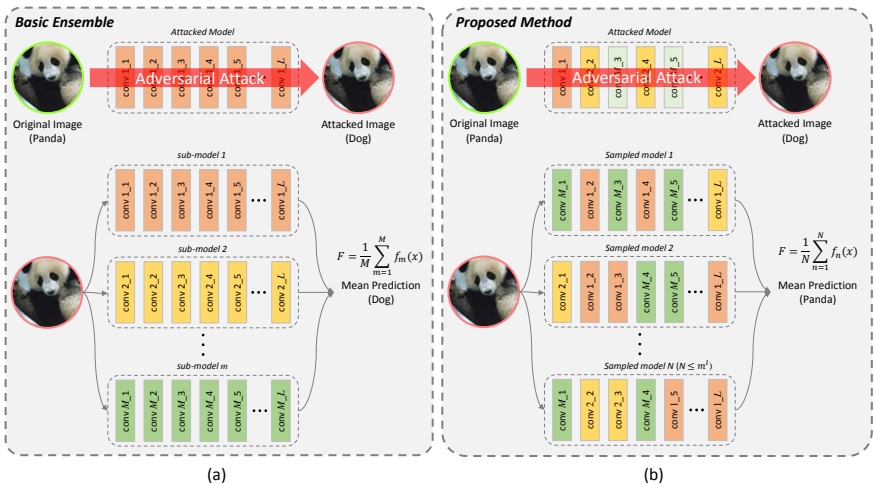


Figure 2: Comparison of defense scenarios between (a) basic ensemble method and (b) our method.

To minimize sample mean prediction, we group K different sampled models with the proposed random layer sampling method. By minimizing $\mathbb{E}[MSE(F | x)]$, the set of individual sampled model is optimized and the diversity is guaranteed.

3.3 Adversarial Defense Scenario

In this section, we describe the adversarial defense scenario. Figure 2 shows the comparison of defense scenario between basic ensemble method and our method. As shown in Figure 2 (a), in the basic ensemble model, if the sub-model is attacked, the prediction is corrupted by the adversarial attack. Note that the prediction score is calculated by averaging M predictions. To reduce the effect of adversarial attack, it is required to increase the number of sub-models but it is limited in real-world applications.

Compared to the basic ensemble approach, our method could generate M^L different ensemble models effectively only with M sub-models. As shown in Figure 2 (b), although a sampled model is attacked, at the test time, our method randomly samples N different models with the RLS method. It is very low probability to sample exactly same sub-model which is attacked ($1/M^L$). Since the weight of the attacked model are different from the ones of the sampled models, the gradients of the attacked model is different from sampled models. Therefore, the attack algorithm does not work properly in the sampled models. As a result, the robustness is improved in the proposed method.

4 Experiments

4.1 Datasets and Implementation Details

To verify the effectiveness of the proposed method, we use three widely studied datasets-MNIST [24], CIFAR10 [10], and SVHN [19] dataset. MNIST dataset is a collection of handwritten digits in classes 0 to 9. It consists of 60,000 training images and 10,000 test

Attack Method	MNIST			CIFAR-10			SVHN		
	Para.	Baseline	Ours	Para.	Baseline	Ours	Para.	Baseline	Ours
No Attack	-	99.5	99.6	-	95.1	96.7	-	96.2	97.6
FGSM	$\epsilon=0.1$	89.5	98.6	$\epsilon=0.01$	45.4	65.1	$\epsilon=0.05$	40.4	56.2
	$\epsilon=0.2$	44.1	70.4	$\epsilon=0.02$	35.6	48.1	$\epsilon=0.1$	28.5	41.0
PGD	$\epsilon=0.05$	68.3	86.4	$\epsilon=0.01$	36.4	70.9	$\epsilon=0.01$	66.0	88.4
	$\epsilon=0.1$	30.8	62.5	$\epsilon=0.02$	12.0	39.2	$\epsilon=0.02$	35.3	68.5
Deepfool	-	3.7	91.4	-	0.4	84.9	-	1.1	84.1
C&W	$c=0.1$	43.7	83.9	$c=0.01$	41.7	91.7	$c=0.01$	38.8	92.3
	$c=1$	13.8	75.3	$c=0.1$	16.4	87.9	$c=0.1$	9.5	75.7
EAD	$c=0.1$	35.9	68.0	$c=0.01$	26.2	79.0	$c=0.01$	29.8	80.3
	$c=1$	12.8	48.3	$c=0.1$	10.8	47.6	$c=0.1$	5.9	69.2

Table 1: Classification accuracy (%) on adversarial examples. Models consist of Resnet-50.

Attack Method	MNIST			CIFAR-10			SVHN		
	Para.	Baseline	RSL	Para.	Baseline	RSL	Para.	Baseline	RSL
No Attack	-	98.2	99.1	-	92.1	93.7	-	96.5	97.0
FGSM	$\epsilon=0.4$	67.3	84.3	$\epsilon=0.1$	45.4	53.3	$\epsilon=0.1$	27.8	40.1
	$\epsilon=0.8$	16.7	28.6	$\epsilon=0.2$	22.0	28.1	$\epsilon=0.2$	17.4	26.8
PGD	$\epsilon=0.1$	66.55	83.9	$\epsilon=0.01$	54.7	78.8	$\epsilon=0.02$	42.5	74.0
	$\epsilon=0.2$	23.57	49.1	$\epsilon=0.02$	37.4	58.9	$\epsilon=0.04$	18.4	45.0
Deepfool	-	1.7	82.2	-	0.9	67.9	-	1.3	68.4
C&W	$c=0.1$	58.5	89.5	$c=0.01$	61.0	91.5	$c=0.01$	47.1	88.6
	$c=1$	15.2	65.8	$c=0.1$	18.2	74.3	$c=0.1$	13.9	68.5
EAD	$c=0.1$	48.1	65.7	$c=0.01$	51.4	84.3	$c=0.01$	18.0	67.4
	$c=1$	22.2	45.6	$c=0.1$	28.5	44.9	$c=0.1$	3.8	61.2

Table 2: Classification accuracy (%) on adversarial examples. Models consist of VGG-16.

images. SVHN dataset is similar to MNIST dataset but colorful street view house numbers. It has 10 classes and consists of 63,257 training images and 26,032 test images. CIFAR10 dataset consist of 60,000 images with 10 classes. It consists of 50,000 training images and 10,000 test images. Each class has 6,000 images. The pixel value of images is scaled to be in an interval $[-1, 1]$. We evaluate our method with two publicly available networks (ResNet-50 [9] and VGG-16 [24]). To verify that the proposed method is compact while robust to adversarial attack, we only use two sub-models and 3 sampled models for training ($M=2, K=3$). All networks are trained from the scratch by using SGD optimizer. The initial learning rate is set as 0.01 and divided by 10 every 50 epochs. We use a weight decay of 5×10^{-4} and a momentum of 0.9. We run the training process for 50 epochs on MNIST, 150 epochs on CIFAR10 and SVHN.

4.2 Performance Comparison under Adversarial Attacks

We evaluate the performance against well-known white-box attacks. We apply five adversarial attack methods (FGSM [8], PGD [17], Deepfool [18], C&W [2], and EAD [4]) on three datasets. For the adversarial attack setting, the iteration step is set to be 10 for PGD with step size of $\epsilon/10$ where ϵ denotes a magnitude of noise. For the C&W and EAD, we perform with constant c . With a selected c , we then run 1000 iterations of gradient decent with Adam optimizer with the learning rate of 0.01.

Table 1 and 2 show the classification accuracy on each adversarial example on ResNet-50 and VGG-16, respectively. We conduct experiment with various settings. ‘‘No Attack’’ denotes the normal setting when testing with normal data. In the case of baseline, we train

Defense	CIFAR-10				
	FGSM	PGD	Deepfool	C&W	EAD
No defense	45.4	36.4	0.4	41.7	26.2
ADP [20]	70.2	61.7	20.3	52.0	61.9
RSE [15]	51.8	45.8	82.1	91.2	75.4
Stochastic Dropout [2]	58.5	59.0	56.4	53.4	38.1
Ours	65.1	70.9	84.9	91.7	79.0

Table 3: Classification accuracy comparison with other defense methods on CIFAR-10.

Model	# of Attacked Sampled Model	FGSM	PGD	C&W
Resnet-50	1	65.1	70.9	91.7
	5	60.5	65.3	84.5
	10	58.7	60.9	82.3
	15	63.1	60.2	81.7

Table 4: Classification accuracy comparison on CIFAR-10 dataset when the attacker attacks multiple sampled models.

a sub-model and attacking that sub-model. In the case of our method, we sample a model from model set by proposed random layer sampling and attack the sampled network. Then, we evaluate accuracy by mean of 10 sample output ($M=2, N=10$). As shown in the tables, our method significantly improves adversarial robustness compared to baseline. Especially, in the case of recently proposed powerful attack methods C&W and EAD, the accuracy of the baseline is significantly dropped on three datasets. In the case of our method, although the adversarial attack is powerful, the accuracy does not drop significantly. Although the attacker knows the full structure and the weight of the sampled model, it attacks different sampled model. With the proposed random layer selection method, we can effectively defend adversarial examples with various sampled model.

4.3 Performance Comparison with Other Defense Methods

We compare our method with other defense methods. We use Resnet-50 network as backbone and test on CIFAR-10 dataset. Table 3 shows classification accuracy comparison with other ensemble based defense methods. We set the attack parameters (ϵ , c , and *number of iteration*) same as section 4.2. As shown in the table, our method outperforms other ensemble methods. Compared with other defense methods, our proposed method improves ensemble diversity through group optimization strategy. Also the sampled models can hide the gradient through random layer sampling. Through the experiment, we prove that the proposed method effectively improves the robustness against to adversarial attack.

4.4 Multiple Attack and Defense

If the attacker knows that the model consists of more than two models, the attacker could attack multiple models. In the multiple attacker scenario, our method could still operate robustly. To verify that our method is also robust to multiple attacks, we conduct multiple attack and defense scenario. Table 4 shows the results of the proposed method when the attacker attacks sampled models. As shown in the table, there are only few accuracy drops

Dataset	Stochastic Dropout	RSE	ADP	Ours
MNIST	0.53 ± 0.005	0.55 ± 0.015	0.52	0.43 ± 0.021
CIFAR-10	0.58 ± 0.004	0.60 ± 0.031	0.54	0.51 ± 0.018
SVHN	0.62 ± 0.003	0.66 ± 0.024	0.62	0.58 ± 0.023

Table 5: Interrater Agreement (IA) score comparison with other ensemble methods.

Method	# of Samples	FGSM	PGD	C&W
Base ensemble	3	50.2	53.7	46.7
	5	51.1	55.2	47.6
Our ensemble	5	64.8	70.3	91.5
	10	65.2	70.8	91.7
	15	65.1	70.9	91.7

Table 6: Classification accuracy on adversarial examples according to the number of sample.

even the attacker attacks multiple attacks. In the case of FGSM, since it is hard to generate adversarial example that could attack more than 10 sampled model, the attacks do not work properly. Also, it is impossible to attack all sampled network. Therefore, it can be interpreted that our method is also robust to multiple attacks.

4.5 Ensemble Diversity Comparison

One of the main contributions of our method is to guarantee the ensemble model diversity. To verify this, we measure the diversity of the proposed method by using Interrater Agreement (IA) score [14, 10]. This score explicitly quantifies the diversity of ensemble models. The lower the IA, the more diverse the predictions of the models. In the case of Stochastic Dropout, RSE, and our method, we use 10 sample and repeat 10 times. Since the ADP method use fixed model, we implement only one time. Table 5 shows the IA scores comparison with other ensemble methods. As shown in the table, our method shows lower IA score than other methods. It means that our proposed method guarantees the diversity.

4.6 Effect of Number of Sample for Adversarial Robustness

We analyze the effect of the number of sample for adversarial robustness. Table 6 shows the classification accuracy according to the number of sample. In the case of Base ensemble, we construct 3 and 5 ResNet-50 sub-models and train individually. Then, we select one sub-model and generate adversarial examples. As shown in the table, as the number of samples increase, the adversarial robustness is also increased. However, to improve the robustness, a large number of parameter is required. On the contrary, in the proposed method, with only two sub-models, it is possible to generate various sampled models with random layer sampling. As a result, our method can hide the weight of attacked sampled-model effectively. Therefore, our method effectively defense the adversarial examples with small number of sub-models.

5 Conclusion

This paper presents an ensemble model training framework with a random layer sampling method for adversarial robustness. In the proposed method, we design a model set that consists of multiple sub-models and construct sampled models with the random layer sampling method. The sampled models are trained by a group optimization strategy to guarantee diversity. After the training, our method predicts various sample outputs by recombination of layers. Therefore, our method effectively hides the full gradient of the models and improves the adversarial robustness. Comprehensive and comparative experiments show that the proposed method could defend adversarial attack effectively.

Acknowledgements

This work was conducted by Center for Applied Research in Artificial Intelligence (CARAI) grant funded by DAPA and ADD (UD190031RD).

References

- [1] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *International Conference on Machine Learning (ICML)*, 2018.
- [2] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [4] Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Ead: elastic-net attacks to deep neural networks via adversarial examples. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [5] Zhipeng Chen, Benedetta Tondi, Xiaolong Li, Rongrong Ni, Yao Zhao, and Mauro Barni. Secure detection of image manipulation by means of random feature selection. *IEEE Transactions on Information Forensics and Security*, 14(9):2454–2469, 2019.
- [6] Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E Kounavis, and Duen Horng Chau. Shield: Fast, practical defense and vaccination for deep learning using jpeg compression. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 196–204, 2018.
- [7] Guneet S Dhillon, Kamyar Aizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense. *International Conference on Learning Representations (ICLR)*, 2018.

- [8] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015.
- [9] 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*, pages 770–778, 2016.
- [10] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical Report*, 2009.
- [11] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. In *Advances in neural information processing systems*, pages 231–238, 1995.
- [12] Ludmila I Kuncheva and Christopher J Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2): 181–207, 2003.
- [13] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, pages 6402–6413, 2017.
- [14] 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.
- [15] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 369–385, 2018.
- [16] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [17] 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.
- [18] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016.
- [19] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [20] Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. *International Conference on Machine Learning (ICML)*, 2019.
- [21] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016.

- [22] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, Xiangyu Zhang, and Jian Sun. Object detection networks on convolutional feature maps. *IEEE transactions on pattern analysis and machine intelligence*, 39(7):1476–1481, 2016.
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations (ICLR)*, 2015.
- [25] Lewis Smith and Yarin Gal. Understanding measures of uncertainty for adversarial example detection. *arXiv preprint arXiv:1803.08533*, 2018.
- [26] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [27] Nguyen Xuan Vinh, Sarah Erfani, Sakrapee Paisitkriangkrai, James Bailey, Christopher Leckie, and Kotagiri Ramamohanarao. Training robust models using random projection. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 531–536. IEEE, 2016.
- [28] 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.
- [29] Zhonghui You, Jinmian Ye, Kunming Li, Zenglin Xu, and Ping Wang. Adversarial noise layer: Regularize neural network by adding noise. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 909–913. IEEE, 2019.
- [30] Zhilu Zhang, Adrian V Dalca, and Mert R Sabuncu. Confidence calibration for convolutional neural networks using structured dropout. *arXiv preprint arXiv:1906.09551*, 2019.