

Toward Robust Diagnosis: A Contour Attention Preserving Adversarial Defense for COVID-19 Detection

Kun Xiang^{1*}, Xing Zhang^{2*}, Jinwen She¹, Jinpeng Liu¹,
Haohan Wang³, Shiqi Deng¹, Shancheng Jiang^{14†}

¹Sun Yat-sen University

²Shuguang Hospital, Shanghai University of Traditional Chinese Medicine

³University of Illinois Urbana-Champaign

⁴Guangdong Provincial Key Laboratory of Fire Science and Technology

{xiangk, shejw3, dengshq5}@mail2.sysu.edu.cn, zhangxing2012@shutcm.edu.cn, liujp22@mails.tsinghua.edu.cn
haohanw@illinois.edu, jiangshch3@mail.sysu.edu.cn

Abstract

As the COVID-19 pandemic puts pressure on healthcare systems worldwide, the computed tomography image based AI diagnostic system has become a sustainable solution for early diagnosis. However, the model-wise vulnerability under adversarial perturbation hinders its deployment in practical situation. The existing adversarial training strategies are difficult to generalized into medical imaging field challenged by complex medical texture features. To overcome this challenge, we propose a Contour Attention Preserving (CAP) method based on lung cavity edge extraction. The contour prior features are injected to attention layer via a parameter regularization and we optimize the robust empirical risk with hybrid distance metric. We then introduce a new cross-nation CT scan dataset to evaluate the generalization capability of the adversarial robustness under distribution shift. Experimental results indicate that the proposed method achieves state-of-the-art performance in multiple adversarial defense and generalization tasks. The code and dataset are available at <https://github.com/Quinn777/CAP>.

Introduction

The coronavirus disease 2019 (COVID-19) pandemic has emerged as one of the pre-eminent global health crises of this century. After the earlier success in many fields (Ker et al. 2017; Anwar et al. 2018; Xiang et al. 2021), various deep neural network-based frameworks (e.g. vision transformer) are proposed as primary screening tools for CT image-based medical diagnosis. However, recent studies found that these systems have exposed vulnerabilities under carefully crafted adversarial perturbations which is usually imperceptible to the human eyes (Szegedy et al. 2013). It raises serious security concerns about their deployment, e.g., health insurance fraud and privacy disclosure, and even aggravating bias against AI-based Computer Aided Diagnosis system (CAD) (Mangaokar et al. 2020; Finlayson et al. 2019).

*These authors contributed equally.

†Corresponding author.

To alleviate this problem, plenty of defense methods have been proposed, including well designed models (Zhou et al. 2022), detection techniques (Lee et al. 2018) and regularization strategies (Chan et al. 2019), etc. Besides, adversarial training (AT) and its variants (Madry et al. 2017; Wong, Rice, and Kolter 2020; Rade and Moosavi-Dezfooli 2021) stand out and become the powerful weapons to defend adversarial attacks. They generally produce artificial noise bounded in ϵ -ball and formulate the key objective as a mini-max optimization problem. Despite their progress, due to the scarcity of medical images and the similarity of textures between classes, existing AT based methods are difficult to maintain attention to salient region or diagnostic points in medical imaging applications especially for diagnosis of COVID-19 based on CT scans. As shown in Figure 1, model under vanilla training can only provide rough and incomplete information about the infection localization. Despite TRADES (Zhang et al. 2019), as a representative AT method, improves the refined perception of local details, the model does not learn structured semantic features. For example, Figure 1c focuses too much attention on the central spine which is unlikely to be infected.

In addition, it has been observed that there is a significant robustness generalization gap under data distribution shift (Rice, Wong, and Kolter 2020; Raghunathan et al. 2019). Some studies provide a frequency-based explanation to reveal the unintuitive generalization behaviors between people and networks (Wang et al. 2020b). Other findings try to improve the model input-robustness with out-of-distribution (OoD) generalization, e.g., data augmentation (Cubuk et al. 2018), feature mixup (Xu et al. 2020) and adversarial domain adaptation (Song et al. 2018). However, existing adversarial robustness researches mainly focus on adapting the distribution discrepancy in natural image. In medical image data, the diverse distributions caused by different patients race, sampling equipment and post-processing methods, etc., usually result in a severer performance degeneration of models. Although the previous data driven methods introduce a larger data latent space, the excessive augmentation for lesion area could destroy the intrinsic structural information, such as lung outline and lesion margin.




Figure 1: Saliency map (Simonyan, Vedaldi, and Zisserman 2013) of three models under different training strategies. Top: Original COVID-19 positive CT scan with expert annotations. Middle: Adversarial noise generated by projected gradient descent (PGD). Bottom: Image under attack. All the models using (b) vanilla training, (c) TRADES and (d) CAP have correct prediction in clean samples, while only the last two models stay robust in adversarial case.

In this paper, we use the latest vision transformers as backbones and take the joint consideration of adversarial robustness and medical data distribution shift. A new medical benchmark data, including 874 COVID-19 and 873 NonCOVID-19 CT images across ten countries, is first established for validation. We then introduce a Contour Attention Preserving (CAP) framework for transferable adversarial training. In order to overcome the attention bias of models to infected areas, we intend to provide a spatial prior of lung cavity contour based on edge features extracted by self-guided filter (He, Sun, and Tang 2010) transformation. We notice that the diagnostic features (such as diffuse patchy shadows and ground glass) are partially destroyed during filtering, which means the enhanced images should be separated from the original images and handled via different branches during the training process. Hence via surrogate model to indirectly minimize the prediction discrepancy between original and augmented counterparts, we generate a regularization term containing implicit prior information from the view of parameters. Without any labels, it is injected into the attention map of the training model in the form of parameter noises. Considering that a visual robust model should not rely on different metrics of distribution diversity, we then propose an optimization framework with hybrid distance metric. The framework uses different distance metrics in internal and external empirical risk calculations to correct the worst-case uncertainty. Figure 1d shows the remarkable resolution of our method for scattered pulmonary nodules and provides lung contour information close to pixel level segmentation. In Section 4, several experimental evidences prove the effectiveness of our method. The main contributions of this paper are summarized as:

- We propose an attention parameter regularization technique based on self-guided filtering. The regularization injects the prior knowledge of lung contour into vision

transformers and introduce a soft consistency between natural images and augmentations.

- A transferable adversarial training framework, CAP, is indicated to defend the imperceptible perturbations in COVID-19 CT data. We raise a hybrid distance metric to further optimize the min-max problem.
- We propose a new CT scan benchmark, COVID-19-C, to evaluate the generalization capacity. Extensive experiments demonstrate that we consistently improve the adversarial robustness and interpretability of state-of-the-art method even under complex data distribution shift.

Related Work

Robustness under Distribution Shift

A large number of works demonstrate the fragility of neural network in artificially corruptions and unseen data distributions (Geirhos et al. 2017; Hendrycks and Dietterich 2018). (Wang et al. 2019) theoretically show global representation can activate more domain-specific signals in the process of gradient propagation. (Huang et al. 2020) applies a feature-critic training heuristic to enforce model’s tendency in predicting with none-predictive features. Another group of methodologies to tackle generalization problem are based on data augmentation. (Zhang et al. 2017) creates mixture samples via different inputs with soft labels, while CutMix (Yun et al. 2019) randomly cuts out samples before mixing. Several works provide a new Fourier-based perspective to fit the data distribution bias (Chen et al. 2021; Xu et al. 2021). Furthermore, regularization loss is also widely used for keeping the consistency of predictions between original and augmented inputs (Wang et al. 2022; Abuduweili et al. 2021). But for texture-sensitive medical imaging task, these strong regularization constraints will inject inductive bias especially with a excessive augmentation method.

Robustness under Adversarial Attacks

Robust training method has experienced explosive development to address the problem of adversarial deception (Szegedy et al. 2013; Goodfellow, Shlens, and Szegedy 2014). Powerful defense ability is obtained by minimizing the training loss, e.g., gradient based methods (Athalye, Carlini, and Wagner 2018; Kurakin, Goodfellow, and Bengio 2018), carefully designed training strategies (Xie et al. 2020; Deng et al. 2021) and self-attention based networks (Benz et al. 2021; Paul and Chen 2022). In medical vision tasks, it is proved that AI systems show less input-robustness than what they have achieved in natural image data (Ma et al. 2021; Kaviani, Han, and Sohn 2022). (Hirano, Minagi, and Takemoto 2021) find CNNs are vulnerable to nontargeted universal perturbation and provide simple detectors to search the feature difference. (Lal et al. 2021) develops a feature fusion technique to defend speckle noise attacks in riabetic retinopathy recognition task while other work improve robustness for 3D whole heart segmentation (Liao et al. 2020). However, there are few works study the adversarial robustness on COVID-19 CT scans, which present less obvious characteristic patterns. Besides, the adversarial robustness

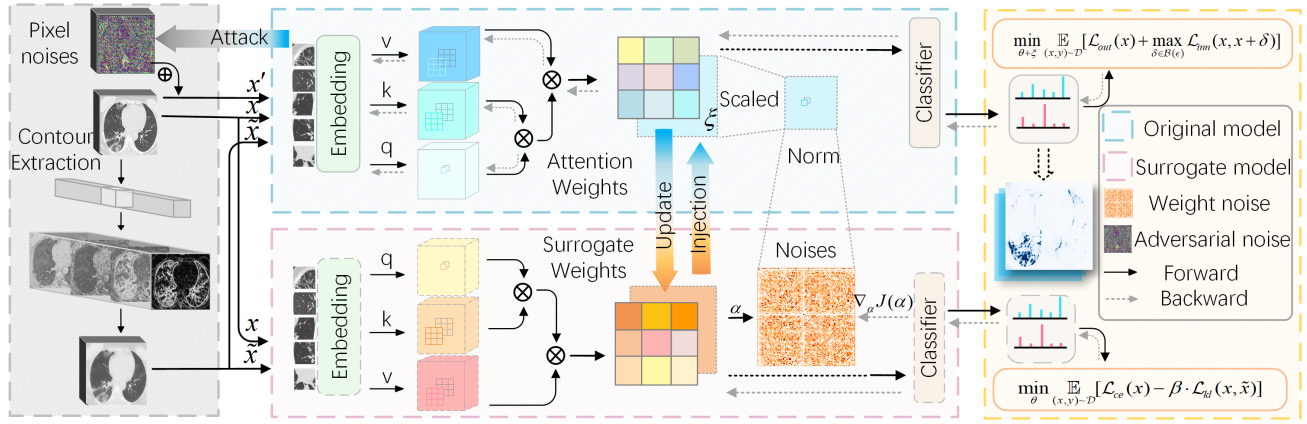


Figure 2: The overall framework of the proposed CAP. In the gray area of the left figure, we first employ a contour extraction module to obtain the prior information of lung cavity. Then, a surrogate model (the middle red area) generates attention noise to attack the weight of target model (the middle blue area). Finally, we optimize the adversarial loss with hybrid distance metric (the right yellow area).

toward distribution shift is not taken into account in previous studies, which causes damage to universality and safety of AT methods in practical deployment.

Methodology

Definition and Overview

Formally, for a given data point (x_i, y_i) sampled from current C -class image dataset $D \subseteq (\mathcal{X}, \mathcal{Y})$, its augmentation is denoted as $(\tilde{x}_i, \tilde{y}_i) \sim \tilde{D}$ where $i \in \{1, \dots, N\}$ is the number of training examples. We represent f to be the discriminative model parameterized by θ , which can be perturbed by attention weight noise ξ . The input adversarial noise is denoted as δ which is bounded by $\|\delta\|_p \leq \epsilon$.

In this section, we first employ a self-guided filter based transformation for edge extraction. The prior information of lung cavity is injected via a well designed attention parameter regularization method. Then we integrate hybrid distance metrics to formulate the empirical risk of min-max optimization. The diagram of our method is shown in Figure 2. We now discuss the technical details of our method.

Contour-preserving Data Augmentation

We first consider to leveraging strong prior knowledge acquired from practical diagnostic experience. Since the tissue lesion caused by COVID-19 often occurs in the interior of lung, we intend to extract this position correlation from lung cavity contour feature. Inspired by the previous work, Guided Filter (He, Sun, and Tang 2010), we adopt a local linear augmentation to smooth small fluctuations without losing edge details. We assume the guidance image in pixel j is $I^{(j)}$ and the filtered output image is \tilde{x} . The transformation can be formulated as:

$$\tilde{x}^{(j)} = \bar{a}_k x^{(j)} + \bar{b}_k, \quad \forall j \in \omega_k \quad (1)$$

where the k is the position index of a local rectangle filter window ω with size $s \times s$. \bar{a}_k and \bar{b}_k are the average of

linear coefficients a_k and b_k respectively, which can be optimized by minimizing the difference between $\tilde{x}^{(j)}$ and the input pixel x . In order to avoid introducing deviations outside the current class, we assume the natural input image x as the self-guidance term I . The solution of the above linear ridge regression is given as:

$$a_k = \frac{\frac{1}{|\omega|} \sum_{i \in \omega_k} x_i^2 - \mu_k \bar{x}_k}{\sigma_k^2 + t}, \quad b_k = \bar{x}_k - a_k \mu_k \quad (2)$$

Here μ_k and σ_k are mean and variance of I in the window k_{th} . $\bar{x}^{(k)}$ is the average of x in ω_k and we use a scaling temperature t to control the degree of smoothness. Since the filter is a edge-preserving function, the machine noise and ground glass in original scans will be dropped and there only left a few trachea tissues and pulmonary nodules. Compared with the frequency domain based and interpolation based methods (Chen et al. 2021; Zhang et al. 2017), we highlight the semantic structures and global visual cues.

Attention Parameter Regularization

The contour-preserving data augmentation allows us to get prior information about the approximate location of lesions. However, simply adopting consistency regularization may neglect the condition that intrinsic information of the instance could be moved during augmentation. As a result, inspired by previous works (Wu, Xia, and Wang 2020; He, Rakin, and Fan 2019; Wen et al. 2018), we propose an attention parameter regularization mechanism with local equivariance to exert a softer constraint. The regularization generates negative weight noise from enhanced images and implicitly introduces the prior knowledge of organ location to attention map.

Considering a simple vision transformer, we add the noise only to the multi-headed self-attention (MSA) module which focuses on the positional relationship between patches. Considering the edge preserving ability of guided filtering, the distributions between the original and augmented images

have an extensive similarity. We expect the decision boundary toward both two types of samples to be close to maintain an implicit consistency. At the same time, the noise confuses the prediction of natural samples by enlarging the expectation of Wasserstein distance between conditional distributions for each clean class. We utilize a mirror image of current model as a surrogate and the surrogate loss is constructed for maximizing this diversity:

$$\mathcal{L}_{sur} = \mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{L}_{ce}(x, y; \theta) - \beta \cdot \mathcal{L}_{kl}((x, \tilde{x}), y; \theta)] \quad (3)$$

where \mathcal{L}_{kl} is Kullback–Leibler divergence and \mathcal{L}_{ce} is cross-entropy loss function. We re-write the weight of attention layer as α . Then with a batch size m , the gradient of surrogate attention weight can be calculated by:

$$\nabla_{\alpha} J(\alpha) = \frac{1}{m} \sum_{i=1}^m \nabla_{\alpha} \mathcal{L}_{sur}((x_i, \tilde{x}_i), y_i; \alpha) \quad (4)$$

To overcome the numerical instability problem, we adopt gradient normalization and control the perturbation amplitude with scaled ratio γ . The final parameter noise can be formulated by:

$$\xi = \eta \frac{\nabla_{\alpha} J(\alpha)}{\|\nabla_{\alpha} J(\alpha)\|_2} \cdot \gamma \|\alpha\|_2 \quad (5)$$

Eq. 3, 4, 5 are all computed in the surrogate model. Similar with (Wu, Xia, and Wang 2020), we move the attention weight back to the center with the same step size after back propagation. Then the final gradient update after simplification can be expressed as:

$$\theta \leftarrow \theta - \frac{\eta}{N} \sum_{i=1}^N \nabla_{\theta+\xi} \mathcal{L}_{out}((x_i, x_i + \delta), \hat{y}_i; \theta + \xi) \quad (6)$$

where $\delta = \arg \max_{\delta \in \mathcal{B}(\epsilon)} \mathcal{L}_{inn}(x_i + \delta, y_i; \theta + \xi)$

Note that in Eq. 6, the outer label \hat{y}_i in \mathcal{L}_{out} is softened by $\hat{y} = \epsilon \cdot f(\tilde{x}_i) + (1 - \epsilon) \cdot y_i$. Since ξ is directly proportional to \mathcal{L}_{sur} , weight perturbation is added to original θ in the process of adversarial learning for maximizing the surrogate error. The weight perturbation implicitly introduces a consistency constraint and narrows the distributions between original and augmented images. Besides, since the weight noise is negative and it forces the model to make wrong decisions for both two types of inputs. The intuition behind is that, the gradient-descent optimizer attempts to find an equilibrium point that perceives model-wise worst-case and keeps contour consistency.

Hybrid Distance Metric based Optimization

In CAP, we improve the adversarial loss function with a hybrid distance metric to substitute \mathcal{L}_{inn} and \mathcal{L}_{out} in Eq. 6. Previous studies have shown that a metric satisfying the distance axioms can help regulate the input inductive bias when maximizing the between-class distances (Pang et al. 2022; Wang et al. 2020a). Considering a well-trained model that achieves high accuracy on both natural and adversarial data distributions, we believe that the input robustness should not

completely depend on how the model measures a specific distribution shift in the adversarial training process. Starting from this point, we attempt to make some changes to the original adversarial framework with different distance metrics in the CAP framework, with the purpose of relieving the inductive bias in optimization.

We denote the class posterior probability as $p(y_i|x_i)$. Following (Golik, Doetsch, and Ney 2013), the square error can be written by:

$$\mathcal{L}_{se} = \sum_i^N \sum_c^C [p(y_k|x_i) - \psi(y_k, y_i)]^2 \quad (7)$$

where $\psi(\cdot, \cdot)$ is the Kronecker function. Compared with KL divergence, square error does not introduce the negative logarithm of posterior probability. Thus there is a smooth transition in the training process. However, our empirical observations demonstrate that when the output is close to the optimal or the worst, the activation value tends to be saturated which manifests a flat loss plane, even though learning is not over. In particular, since the step size and magnitude of perturbation are usually small while generating strong rival classes in the update of adversarial attacks, it brings more difficulties to converge to the farthest position when facing gradient vanishing. Consequently, we consider applying the square error function only in the external minimization problem as \mathcal{L}_{out} . Our final optimization objectives can be summarized as follows:

$$\min_{\theta+\xi} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{L}_{se}(x, y; \theta + \xi) + \beta \cdot \mathcal{L}_{se}(x, x + \delta; \theta + \xi)] \quad (8)$$

where $\delta = \arg \max_{\delta \in \mathcal{B}(\epsilon)} \mathcal{L}_{kl}((x, x + \delta), y; \theta + \xi)$

We also provide some theoretical evidences to bound the formulation. We remove $\mathcal{L}_{se}(x_i, y_i)$ in sample x_i with label y_i by:

$$\sum_i^N [p(y|x_i) - \psi(y, y_i)]^2 = [1 - p(y_i|x_i)]^2 + \sum_{y_i \neq y} q^2(y|x_i) \quad (9)$$

As such, we substitute it into Eq. 9 to bound the external optimization:

$$0 \leq \mathcal{L}_{out}(x_i, y_i) \leq 2 + 2\beta \quad (10)$$

Here \mathcal{L}_{out} is the external loss in Eq. 8. Detailed proof is shown in Appendix A. Note that the notation x_i' is used as an abbreviation for adversarial sample $x_i + \delta_i$. The square loss reaches the upper bound when there is a strong rival class and the decision boundary confuses both natural and adversarial samples. We provide some empirical evidence for the above theory in Table 1. The compound metrics significantly improve the performance in both TRADES and CAP frameworks. Our final pipeline is summarized in Figure 2.




Figure 3: Four datasets we used for evaluation. The test sets include images with both weak (COVID19-C) and strong (MosMed-M) distribution shift.

Method	Outer	Inner	Clean	PGD-10
TRADES	KL	KL	88.73	77.67
	SE	SE	84.51	41.85
	KL	SE	84.50	50.70
	SE	KL	91.95	84.31
CAP (Ours)	KL	KL	86.72	58.35
	SE	SE	93.96	63.68
	KL	SE	93.96	41.85
	SE	KL	92.96	85.92

Table 1: Effect of different distance metrics on optimization problems. Outer/Inner represent the distance metric used in minimum/maximum optimization problem, respectively.

Experiments

Datasets We use two public COVID-19 CT databases for evaluation, i.e., **SARS-COV-2** (Morozov et al. 2020) and **MosMed** (Soares et al. 2020). SARS-COV-2 contains 2482 2D lung window images, including positive and negative patients. MosMed provides a total of 1110 3D scans, which is used for infection severity quantification. It contains 5 types of labels (healthy, mild, moderate, severe and critical), which are subdivided by the infection percentage of CT volume. We adopt Hounsfield transformation to convert it into two modalities of CT images, i.e., lung window and mediastinal window. They are named as **MosMed-L** and **MosMed-M** respectively.

To evaluate the generalization capacity under distribution shift, we establish a comprehensive database named **COVID19-C**. COVID19-C is collected from multiple public datasets (Wang, Lin, and Wong 2020; Hemdan, Shouman, and Karar 2020; Morozov et al. 2020; Soares et al. 2020) and involves 1747 images from patients in 10 different countries. Each case is chosen by experienced radiologists to cover as much varieties from different background and texture shift as possible. We divide COVID19-C into 4 subsets according to regions, which are abbreviated as **COV-C** (China), **COV-R** (Russia), **COV-I** (Iran) and **COV-M**. Note that COV-M contains patients from multiple countries and samples in

four subsets are independent. Examples are shown in Figure. 3 and Table. 2.

Settings The proposed method is compared with a series of AT and OoD baselines, i.e., FAST (Wong, Rice, and Kolter 2020), AT (Madry et al. 2017), TRADES (Zhang et al. 2019), AVmixup (Lee, Lee, and Yoon 2020), HAT (Rade and Moosavi-Dezfooli 2021), FAT (Zhang et al. 2020), AWP (Wu, Xia, and Wang 2020), MMA (Ding et al. 2018), SCORE (Pang et al. 2022); Mixup (Zhang et al. 2017), AutoAug (Cubuk et al. 2018), CutOut (DeVries and Taylor 2017), CutMix (Yun et al. 2019), FACT (Xu et al. 2021) and APRS (Chen et al. 2021). For SARS-COV-2, we use Visformer-Tiny and Deit-Tiny as our backbones and train them for 200 epochs. We choose AdamW optimizer with cosine annealing learning rate 0.0005 for Visformer while 0.0001 for Deit. For MosMed-L, we adopt Visformer-Tiny with AdwmW, learning rate of 0.0002 with 20 steps decay. Models are trained for 200 epochs with batch size 64. Attacks are conducted for 10 steps in SARS-COV-2 (5 steps in MosMed-L) with step size of $2/255$ and l_∞ perturbation budget $\epsilon = 8/255$. In CAP, we set scaled ratio $\gamma = 0.0001$, $t = 0.003$ with kernel size 5. Training and test sets are divided by 8:2 and images are resized to $224 \cdot 224$. We implement experiments with PyTorch framework on Nvidia RTX 3090 GPUs and more details are shown in Appendix C.

Dataset	Positive	Negative	Total	Country	Slice
SARS-COV-2	1252	1229	2481	Brazil	Expert
COV-C	1429	1648	3077	China	Expert
COV-R	236	702	938	Russia	Expert
COV-I	105	105	210	Iran	Expert
COV-M	182	170	352	-	Auto

Table 2: Datasets for COVID-19 binary detection task.

Visformer-Tiny												
Method	Clean	FGSM	R-FGSM	PGD	MIM	AutoAttack	Avg	COV-C	COV-R	COV-I	COV-M	Avg
FAST	81.69	81.49	79.88	76.26	72.84	72.43	76.58	32.98	27.84	46.67	43.51	37.75
AT	84.51	79.68	82.29	79.28	79.48	79.28	80.00	42.23	44.03	46.19	48.97	45.36
TRADES	88.73	80.08	84.31	77.67	78.74	74.25	79.01	44.24	45.74	49.05	47.38	46.60
AVmixup	93.16	77.67	85.71	74.85	75.25	74.25	77.55	30.97	38.07	31.43	42.37	35.71
HAT	84.71	67.81	76.06	67.27	67.61	66.6	69.07	40.35	39.77	23.33	38.50	35.49
FAT	92.76	82.09	87.73	81.09	81.29	80.89	82.62	39.01	48.58	38.10	47.15	43.21
AWP	90.54	80.08	84.51	77.87	78.47	76.26	79.44	41.55	44.32	50.00	49.89	46.44
MMA	75.25	63.38	67.81	61.17	60.97	57.34	62.13	30.56	40.91	46.67	46.47	41.15
SCORE	84.51	42.05	42.05	41.85	41.65	15.49	36.62	37.67	36.65	30.48	44.65	37.36
TRADES+Mixup	82.49	58.55	70.62	56.34	56.54	56.34	59.68	30.97	38.07	15.24	38.72	30.75
TRADES+AutoAug	71.63	65.39	67.81	64.79	64.99	64.39	65.47	32.17	39.49	42.38	43.05	39.27
TRADES+CutOut	85.51	75.86	78.04	73.44	68.01	62.78	71.63	31.37	37.78	46.46	46.47	40.52
TRADES+CutMix	76.66	56.34	64.99	53.92	54.33	53.52	56.62	20.51	24.15	31.9	21.64	24.55
TRADES+FACT	84.10	78.07	80.28	77.46	77.46	77.46	78.15	39.54	44.32	35.24	39.41	39.63
TRADES+APRS	87.32	76.66	82.49	75.05	75.65	74.45	76.86	38.20	32.10	43.33	45.79	39.86
CAP(Ours)	92.96	86.52	90.54	85.92	85.92	85.92	86.96	45.58	50.57	48.10	53.99	49.56

Deit-Tiny												
Method	Clean	FGSM	R-FGSM	PGD	MIM	AutoAttack	Avg	COV-C	COV-R	COV-I	COV-M	Avg
FAST	93.76	60.36	75.25	39.44	47.48	23.54	49.21	8.85	5.68	0.01	2.73	4.32
AT	91.55	83.5	87.12	81.49	82.09	81.09	83.06	35.79	44.03	46.19	39.18	41.30
TRADES	94.37	84.91	89.34	82.49	83.10	82.09	84.39	36.73	40.91	44.76	46.01	42.10
AVmixup	94.37	85.51	90.54	82.70	82.90	82.29	84.79	36.86	40.91	44.29	45.10	41.79
HAT	95.17	82.70	89.94	79.68	80.28	79.07	82.33	26.14	21.59	38.57	20.73	26.76
AWP	94.16	85.11	89.94	82.70	83.10	81.09	84.39	39.01	43.18	44.76	45.10	43.01
MMA	83.50	68.01	73.84	64.19	64.99	64.19	67.04	30.29	23.86	44.76	35.99	33.73
SCORE	82.09	73.74	77.06	60.36	55.73	36.42	60.66	35.25	38.07	12.86	32.35	29.63
CAP (Ours)	95.57	87.12	90.74	84.31	85.51	83.70	86.28	44.88	42.90	47.62	47.38	45.70

Table 3: Clean accuracy, adversarial accuracy and OoD adversarial accuracy (%) on SARS-COV-2 and COVID-19-C. We perturb OoD data with PGD-10 and all the attacks are generated with 10 iterations ($\epsilon = 8/255$ and step size $2/255$).

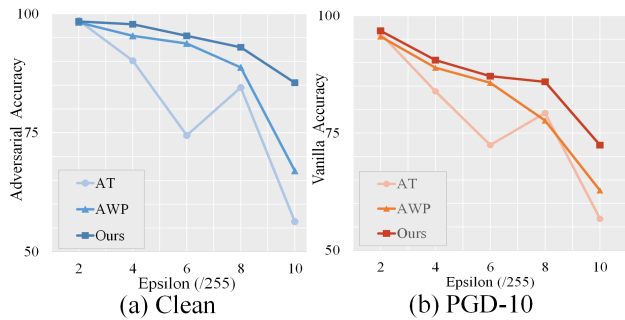


Figure 4: Effect of a wide range of perturbation budget (ϵ) in adversarial training procedure with different method on SARS-COV-2.

Results

COVID-19 Detection on SARS-COV-2

White-box Attack To evaluate the performance of our proposed method, we first employ a wide range of white-box adversarial attacks with l_∞ threat model: FGSM (Goodfellow, Shlens, and Szegedy 2014), R-FGSM (Tramèr et al. 2017), PGD (Madry et al. 2017), MIM (Dong et al. 2018) and Auto Attack (Croce and Hein 2020). As reported in Table

3, we achieve the highest adversarial accuracy in Visformer under all five attacks and improve the SOTA result (FAT) by 4.34% on average while maintaining natural accuracy. A narrower gap between clean and adversarial accuracy is observed in CAP (6% vs 10.14% compared with FAT). In Deit, we get 1.2% and 1.49% improvement than AVmixup in clean and adversarial samples, respectively. Moreover, Figure 4 shows CAP consistently improves the robustness with different perturbation budgets.

Distribution Shift In addition to surpassing numerous previous methods in source data distribution, from Table 3, it is evident that CAP also outperforms them under distribution shift. We introduce COVID19-C for generalization ability evaluation and images are perturbed by PGD-10 with $\epsilon = 8/255$ and step size $2/255$. Results demonstrate that we gain an average improvement of 3.12% and 2.69% than SOTA baselines with Visformer and Deit respectively. More intriguingly, existing augmentations, e.g., Cutout (9.04% lower than ours) and AutoAug (10.29% lower than ours), cannot transfer adversarial robustness to unseen distribution as they do in natural images. Furthermore, the means \pm S.D of clean accuracy, average adversarial accuracy and OoD adversarial accuracy for Visformer are listed as follows (%): $\{82.13\pm 9.89, 70.10\pm 11.91, 38.91\pm 5.68\}$ and $\{91.12\pm 4.92, 74.48\pm 12.86, 32.83\pm 12.26\}$ for Deit. The results show that

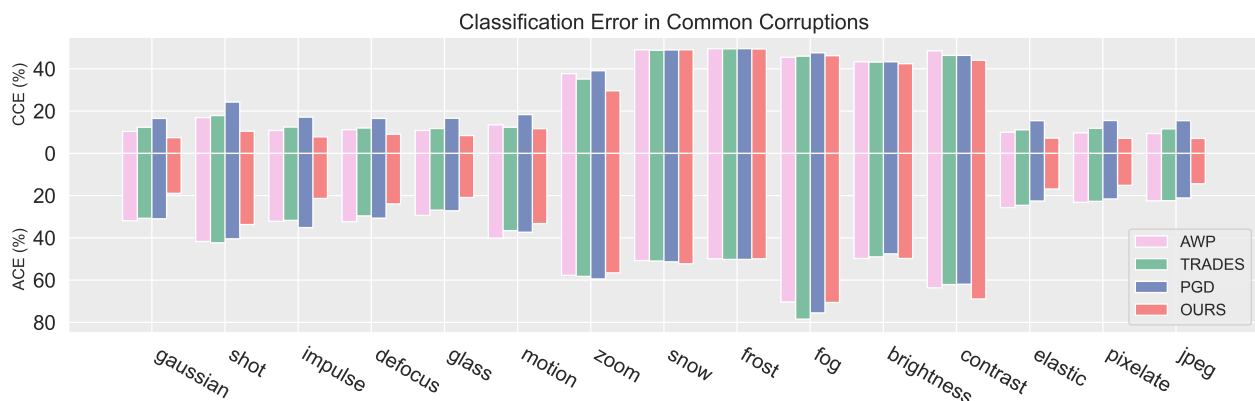


Figure 5: Clean Corruption Error (CCE) and Adversarial Corruption Error (ACE) for different methods with Visformer-Tiny in corrupt SARS-COV-2.

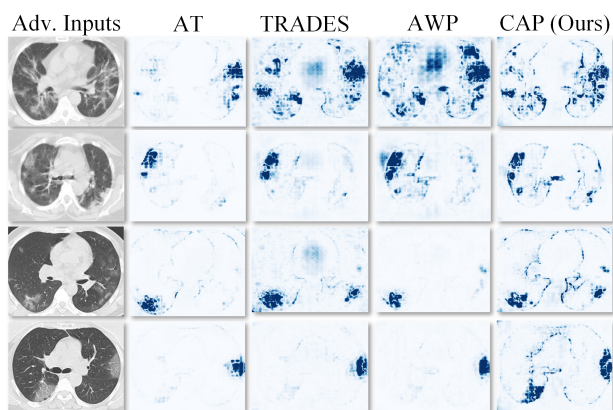


Figure 6: Saliency maps on models with different AT methods. The input images are perturbed by PGD-10 with $\epsilon = 8/255$.

the SOTAs in Deit achieve better adversarial robustness while bring larger generalization error. Overall, the experimental results prove that our enhancement avoids excessive destruction of structural information and accurately represents the ontology semantics.

Common Corruption Due to the discrepancy in hardware and shooting parameters, there exists surface variations between input scans. Hence we evaluate the generalization behavior with 15 common corruptions (Hendrycks and Dietterich 2019), e.g., gaussian noise and elastic. All the models are tested with CCE and ACE, which can be calculated by $CCE_c = \frac{1}{5} \sum_{s=1}^5 E_{s,c}(x)$ and $ACE_c = \frac{1}{5} \sum_{s=1}^5 E_{s,c}(x')$. Here x' is corrupt image with PGD attack like Table 3. Figure 5 shows our CAP decreases both CCE and ACE from SOTA method (AWP) in most corruptions, e.g. shot noise CCE 10.42% vs 16.86% and ACE 33.76% vs 41.77%. The mean CCE and mean ACE in CAP is 22.42% and 36.48%, which are both better than AWP (25.06% and 41.44%). Notice that despite the performance of CAP in some cases (such

as weather) is slightly decreased, these corruptions will not appear in practical CT scans.

Visualization Figure 6 indicates that our CAP provides a sharper and more visually coherent gradient-based interpretability for perturbed CT scans. It is obvious that vanilla trained model cannot learn robust features from adversarial samples (second column). We alleviate the local omission of previous AT methods for large-scale infection (last row of the first two columns). On further investigation, the model trained by CAP highlights the contour of lung cavity and eliminates the attention outside the lung (last row of the last two columns). It proves our method may indeed learn the prior structure information and achieve a visual robustness. More figures are provided in Appendix D.4.

Resource Requirement To measure the resource requirement of different AT frameworks, we compute the seconds per training epoch for several high-accuracy baselines. In Visformer, the training time of PGD, TRADES, AWP, FAT and CAP are {102.53, 105.49, 117.26, 135.09, 113.49}. In Deit, the values of them are {80.86, 119.75, 83.36, 111.39, 82.73}. The results show that our algorithm only needs medium-sized computational resource when achieving the best performance.

Severity Classification on MosMed

We then consider a more challenging severity classification task and the results are shown in Table 4. In source data distribution (Lung window scans), the accuracy, specificity and F1 score of CAP for adversarial samples are improved consistently. In particular, it outperforms the SOTA method, standard AT, by a large margin of 4.7% of adversarial accuracy in total five classes. In class 0, 1 and 2, our clean accuracy slightly decreases since the trade-off (Zhang et al. 2019) still exists in medical data.

Furthermore, when deploying previous techniques to data with larger distribution shift (i.e., Mediastinal window), their adversarial robustness deteriorates a lot. For example, the clean and adversarial accuracy in TRADES are surprisingly 68.1%/65.2% worse than they perform in MosMed-L, while

MosMed-L										
		AT			TRADES			CAP (Ours)		
Class	Infection	Accuracy	Specificity	F1 Score	Accuracy	Specificity	F1 Score	Accuracy	Specificity	F1 Score
0	Healthy	93.8/90.1	97.0/94.6	85.9/77.6	97.2 /89.0	98.4 /93.1	93.9 /75.8	96.2/ 90.4	98.3/ 95.0	91.5/ 78.0
1	$0 < x < 25$	89.2/82.7	79.1/70.2	91.6/86.6	94.9 /81.7	91.8 /73.5	95.9 /85.4	94.2/ 85.7	88.7/ 75.8	95.4/ 88.8
2	$25 < x < 50$	95.2/91.7	98.7/97	76.2/57.8	97.8 /92.2	99.1/96.4	90.2 /63.0	97.4/ 94.0	99.1 /97.6	87.9/ 71.0
3	$50 < x < 75$	98.0/97.3	99.6/99.2	71.1/61.5	98.7/96.6	99.5/98.4	83.7/55.8	98.8 /97.5	99.8 /99.1	84.0 /65.9
4	$75 < x < 100$	100/100	100/ 100	100/100	100/100	100/100	100/ 100	100/100	100/100	100 /80
Total	-	88.1/80.9	-	-	94.3 /79.7	-	-	93.3/ 85.6	-	-

MosMed-L → MosMed-M										
0	Healthy	50.0/54.4	52.8/66.8	27.1/11.2	26.7/15.5	6.0/2.2	37.6/24.6	76.5 /74.0	98.7 /95.9	3.0/NAN
1	$0 < x < 25$	40.5/36.5	94.6/90.6	12.5/5.1	41.3/30.0	96.7/75.1	12.4/2.0	61.4 /58.8	1.4/0.5	75.9 /74.0
2	$25 < x < 50$	88.7/88.7	100/100	NAN	88.7/88.7	100/100	NAN	88.7 /88.7	100 /100	NAN
3	$50 < x < 75$	96.0/96.0	100/100	NAN	96.0/96.0	100/99.4	NAN	96.0 /96.0	100 /100	NAN
4	$75 < x < 100$	52.2/34.0	52.1/33.8	0.7 /0.5	99.8/99.8	100/100	NAN	99.8 /99.8	100 /100	NAN
Total	-	13.7/4.8	-	-	26.2/14.5	-	-	61.2 /58.7	-	-

Table 4: Clean/Adversarial performance of models trained with MosMed-L. We use both MosMed-L and MosMed-M for evaluation.

CAP narrows the gap by only 32.1%/26.9% respectively. Note that there is a numerical explosion in some evaluations, such as F1 score from class 2 to 4. The underlying reason is that the lower contrast in mediastinal window images affects the recognition of large-area diffuse patchy infections. Fortunately, we can still distinguish positive and negative patients with a better F1 score (75.9/74.0 vs the highest 12.5/5.1), even if it could underestimate the infection volumes.

Ablation Study

We perform an ablation study for each component of our CAP framework, including self-guided filter (SGF), attention parameter regularization (APR) and hybrid distance metric (HDM). As shown in Table 5, each component of CAP is indeed effective. Firstly, applying our edge-preserving augmentation in vanilla TRADES improves both adversarial accuracy and generalization ability by 2.01% and 2.73%. After attention parameter regularization, all the three metrics are improved especially adversarial accuracy on SARS-COV-2, which indicates the importance of the way to embed the prior information. We further validate the necessity of HDM and results suggest the effectiveness of incorporating all the three components. More ablations about the parameter-sensitivity and the comparison with consistency regularization (Tack et al. 2022) are carried out in Appendix D.1 and D.2, respectively.

SGF	AWR	HDM	Clean	PGD-10	COV-M
			88.73	77.67	47.38
✓			86.52	79.68	50.11
✓	✓		90.14	83.50	51.71
✓	✓	✓	92.96	85.92	53.08

Table 5: Ablation study on different components of our method on SARS-COV-2 and COV-M with Visformer-T.

Conclusion

In this paper, we propose a contour attention preserving (CAP) framework to improve the model-wise adversarial robustness in COVID-19 CT classification. An implicit parameter regularization is first introduced to inject the lung cavity prior feature into vision transformer. We further optimize the min-max problem via a hybrid distance metric. Then we propose a new COVID-19 CT dataset to evaluate the robustness under distribution shift. Extensive experiments demonstrate our method significantly improves the adversarial robustness and interpretability in medical imaging diagnosis.

Acknowledgements

This work was supported in part by the National Nature Science and Foundation of China under Grant No. 71801031, the National Key Research and Development Program of China under Grant No. 2020YFB1713800, and the Guangdong Basic and Applied Basic Research Foundation of China project “Research on dermatosis automatic diagnosis system based on multi-type of medical image” under Grant No. 2019A1515011962.

References

- Abuduweili, A.; Li, X.; Shi, H.; Xu, C.-Z.; and Dou, D. 2021. Adaptive consistency regularization for semi-supervised transfer learning. In *Proc. of CVPR*.
- Anwar, S. M.; Majid, M.; Qayyum, A.; Awais, M.; Alnowami, M.; and Khan, M. K. 2018. Medical image analysis using convolutional neural networks: a review. *Journal of medical systems*.
- Athalye, A.; Carlini, N.; and Wagner, D. 2018. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proc. of ICML*.
- Benz, P.; Ham, S.; Zhang, C.; Karjauv, A.; and Kweon, I. S. 2021. Adversarial robustness comparison of vision transformer and mlp-mixer to cnns. *arXiv preprint arXiv:2110.02797*.

- Chan, A.; Tay, Y.; Ong, Y. S.; and Fu, J. 2019. Jacobian adversarially regularized networks for robustness. *arXiv preprint arXiv:1912.10185*.
- Chen, G.; Peng, P.; Ma, L.; Li, J.; Du, L.; and Tian, Y. 2021. Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. In *Proc. of ICCV*.
- Croce, F.; and Hein, M. 2020. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proc. of ICML*.
- Cubuk, E. D.; Zoph, B.; Mane, D.; Vasudevan, V.; and Le, Q. V. 2018. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*.
- Deng, Z.; Zhang, L.; Ghorbani, A.; and Zou, J. 2021. Improving adversarial robustness via unlabeled out-of-domain data. In *Proc. of AISTATS*.
- DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Ding, G. W.; Sharma, Y.; Lui, K. Y. C.; and Huang, R. 2018. Mma training: Direct input space margin maximization through adversarial training. *arXiv preprint arXiv:1812.02637*.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting adversarial attacks with momentum. In *Proc. of CVPR*.
- Finlayson, S. G.; Bowers, J. D.; Ito, J.; Zittrain, J. L.; Beam, A. L.; and Kohane, I. S. 2019. Adversarial attacks on medical machine learning. *Science*.
- Geirhos, R.; Janssen, D. H.; Schütt, H. H.; Rauber, J.; Bethge, M.; and Wichmann, F. A. 2017. Comparing deep neural networks against humans: object recognition when the signal gets weaker. *arXiv preprint arXiv:1706.06969*.
- Golik, P.; Doetsch, P.; and Ney, H. 2013. Cross-entropy vs. squared error training: a theoretical and experimental comparison. In *Proc. of Interspeech*.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- He, K.; Sun, J.; and Tang, X. 2010. Guided image filtering. In *Proc. of ECCV*.
- He, Z.; Rakin, A. S.; and Fan, D. 2019. Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In *Proc. of CVPR*.
- Hemdan, E. E.-D.; Shouman, M. A.; and Karar, M. E. 2020. Covidx-net: A framework of deep learning classifiers to diagnose covid-19 in x-ray images. *arXiv preprint arXiv:2003.11055*.
- Hendrycks, D.; and Dietterich, T. 2019. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*.
- Hendrycks, D.; and Dietterich, T. G. 2018. Benchmarking neural network robustness to common corruptions and surface variations. *arXiv preprint arXiv:1807.01697*.
- Hirano, H.; Minagi, A.; and Takemoto, K. 2021. Universal adversarial attacks on deep neural networks for medical image classification. *BMC medical imaging*.
- Huang, Z.; Wang, H.; Xing, E. P.; and Huang, D. 2020. Self-challenging improves cross-domain generalization. In *Proc. of ECCV*.
- Kaviani, S.; Han, K. J.; and Sohn, I. 2022. Adversarial attacks and defenses on AI in medical imaging informatics: A survey. *Expert Systems with Applications*.
- Ker, J.; Wang, L.; Rao, J.; and Lim, T. 2017. Deep learning applications in medical image analysis. *Ieee Access*.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2018. Adversarial examples in the physical world. In *Artificial intelligence safety and security*.
- Lal, S.; Rehman, S. U.; Shah, J. H.; Meraj, T.; Rauf, H. T.; Damaševičius, R.; Mohammed, M. A.; and Abdulkareem, K. H. 2021. Adversarial attack and defence through adversarial training and feature fusion for diabetic retinopathy recognition. *Sensors*.
- Lee, K.; Lee, K.; Lee, H.; and Shin, J. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Proc. of NeurIPS*.
- Lee, S.; Lee, H.; and Yoon, S. 2020. Adversarial vertex mixup: Toward better adversarially robust generalization. In *Proc. of CVPR*.
- Liao, X.; Qian, Y.; Chen, Y.; Xiong, X.; Wang, Q.; and Heng, P.-A. 2020. MMTLNet: Multi-Modality Transfer Learning Network with adversarial training for 3D whole heart segmentation. *Computerized Medical Imaging and Graphics*.
- Ma, X.; Niu, Y.; Gu, L.; Wang, Y.; Zhao, Y.; Bailey, J.; and Lu, F. 2021. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Mangaokar, N.; Pu, J.; Bhattacharya, P.; Reddy, C. K.; and Viswanath, B. 2020. Jekyll: Attacking medical image diagnostics using deep generative models. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, 139–157. IEEE.
- Morozov, S.; Andreychenko, A.; Pavlov, N.; Vladzmyrskyy, A.; Ledikhova, N.; Gombolevskiy, V.; Blokhin, I. A.; Gelezhe, P.; Gonchar, A.; and Chernina, V. Y. 2020. Mosmeddata: Chest ct scans with covid-19 related findings dataset. *arXiv preprint arXiv:2005.06465*.
- Pang, T.; Lin, M.; Yang, X.; Zhu, J.; and Yan, S. 2022. Robustness and Accuracy Could Be Reconcilable by (Proper) Definition. *arXiv preprint arXiv:2202.10103*.
- Paul, S.; and Chen, P.-Y. 2022. Vision transformers are robust learners. In *Proc. of AAAI*.
- Rade, R.; and Moosavi-Dezfooli, S.-M. 2021. Helper-based adversarial training: Reducing excessive margin to achieve a better accuracy vs. robustness trade-off. In *ICML 2021 Workshop on Adversarial Machine Learning*.

- Raghunathan, A.; Xie, S. M.; Yang, F.; Duchi, J. C.; and Liang, P. 2019. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032*.
- Rice, L.; Wong, E.; and Kolter, Z. 2020. Overfitting in adversarially robust deep learning. In *Proc. of ICML*.
- Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Soares, E.; Angelov, P.; Biaso, S.; Froes, M. H.; and Abe, D. K. 2020. SARS-CoV-2 CT-scan dataset: A large dataset of real patients CT scans for SARS-CoV-2 identification. *MedRxiv*.
- Song, C.; He, K.; Wang, L.; and Hopcroft, J. E. 2018. Improving the generalization of adversarial training with domain adaptation. *arXiv preprint arXiv:1810.00740*.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tack, J.; Yu, S.; Jeong, J.; Kim, M.; Hwang, S. J.; and Shin, J. 2022. Consistency regularization for adversarial robustness. In *Proc. of AAAI*.
- Tramèr, F.; Kurakin, A.; Papernot, N.; Goodfellow, I.; Boneh, D.; and McDaniel, P. 2017. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*.
- Wang, H.; Chen, T.; Gui, S.; Hu, T.; Liu, J.; and Wang, Z. 2020a. Once-for-all adversarial training: In-situ tradeoff between robustness and accuracy for free. *Proc. of NeurIPS*.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning robust global representations by penalizing local predictive power. *Proc. of NeurIPS*.
- Wang, H.; Huang, Z.; Wu, X.; and Xing, E. P. 2022. Toward Learning Robust and Invariant Representations with Alignment Regularization and Data Augmentation. *arXiv preprint arXiv:2206.01909*.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020b. High-frequency component helps explain the generalization of convolutional neural networks. In *Proc. of CVPR*.
- Wang, L.; Lin, Z. Q.; and Wong, A. 2020. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports*.
- Wen, Y.; Vicol, P.; Ba, J.; Tran, D.; and Grosse, R. 2018. Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *arXiv preprint arXiv:1803.04386*.
- Wong, E.; Rice, L.; and Kolter, J. Z. 2020. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*.
- Wu, D.; Xia, S.-T.; and Wang, Y. 2020. Adversarial weight perturbation helps robust generalization. *Proc. of NeurIPS*.
- Xiang, K.; Peng, L.; Yang, H.; Li, M.; Cao, Z.; Jiang, S.; and Qu, G. 2021. A novel weight pruning strategy for light weight neural networks with application to the diagnosis of skin disease. *Applied Soft Computing*.
- Xie, Q.; Luong, M.-T.; Hovy, E.; and Le, Q. V. 2020. Self-training with noisy student improves imagenet classification. In *Proc. of CVPR*.
- Xu, M.; Zhang, J.; Ni, B.; Li, T.; Wang, C.; Tian, Q.; and Zhang, W. 2020. Adversarial domain adaptation with domain mixup. In *Proc. of AAAI*.
- Xu, Q.; Zhang, R.; Zhang, Y.; Wang, Y.; and Tian, Q. 2021. A fourier-based framework for domain generalization. In *Proc. of CVPR*.
- Yun, S.; Han, D.; Oh, S. J.; Chun, S.; Choe, J.; and Yoo, Y. 2019. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proc. of ICCV*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Zhang, H.; Yu, Y.; Jiao, J.; Xing, E.; El Ghaoui, L.; and Jordan, M. 2019. Theoretically principled trade-off between robustness and accuracy. In *Proc. of ICML*.
- Zhang, J.; Xu, X.; Han, B.; Niu, G.; Cui, L.; Sugiyama, M.; and Kankanhalli, M. 2020. Attacks which do not kill training make adversarial learning stronger. In *Proc. of ICML*.
- Zhou, D.; Yu, Z.; Xie, E.; Xiao, C.; Anandkumar, A.; Feng, J.; and Alvarez, J. M. 2022. Understanding the robustness in vision transformers. *arXiv preprint arXiv:2204.12451*.