

Language-Driven Anchors for Zero-Shot Adversarial Robustness

Xiao Li¹ Wei Zhang¹ Yining Liu² Zhanhao Hu¹ Bo Zhang¹ Xiaolin Hu^{1,3*}

¹Department of Computer Science and Technology, BNRist, Institute for Artificial Intelligence, IDG/McGovern Institute for Brain Research, Tsinghua Laboratory of Brain and Intelligence, Tsinghua University, Beijing, China

²Harbin Institute of Technology, Weihai, China

³Chinese Institute for Brain Research (CIBR), Beijing, China

{lixiao20, zhang-w19, huzhanha17}@mails.tsinghua.edu.cn

22s030184@stu.hit.edu.cn

{dcszb, xlhu}@mail.tsinghua.edu.cn

Abstract

Deep Neural Networks (DNNs) are known to be susceptible to adversarial attacks. Previous researches mainly focus on improving adversarial robustness in the fully supervised setting, leaving the challenging domain of zero-shot adversarial robustness an open question. In this work, we investigate this domain by leveraging the recent advances in large vision-language models, such as CLIP, to introduce zero-shot adversarial robustness to DNNs. We propose LAAT, a Language-driven, Anchor-based Adversarial Training strategy. LAAT utilizes the features of a text encoder for each category as fixed anchors (normalized feature embeddings) for each category, which are then employed for adversarial training. By leveraging the semantic consistency of the text encoders, LAAT aims to enhance the adversarial robustness of the image model on novel categories. However, naively using text encoders leads to poor results. Through analysis, we identified the issue to be the high cosine similarity between text encoders. We then design an expansion algorithm and an alignment cross-entropy loss to alleviate the problem. Our experimental results demonstrated that LAAT significantly improves zero-shot adversarial robustness over state-of-the-art methods. LAAT has the potential to enhance adversarial robustness by large-scale multimodal models, especially when labeled data is unavailable during training. Code is available at <https://github.com/LixiaoTHU/LAAT>.

1. Introduction

Adversarial attacks [42], by adding deliberately designed perturbations to inputs, have posed a significant threat to

*Corresponding author.

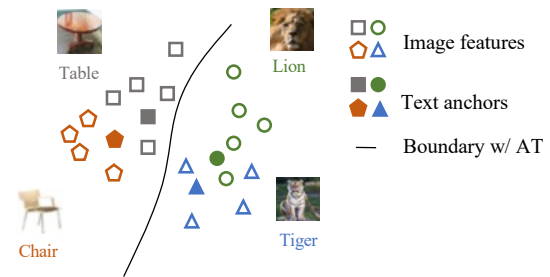


Figure 1. The illustration of zero-shot ability with LAAT. Different colors of the marks indicate different categories. When a model is adversarially trained with the text anchors of *table* and *lion* (seen categories), it can recognize adversarial examples of the two categories (grey and green). Then due to the text anchors of *chair* and *tiger* (novel categories) being close to those of *table* and *lion*, respectively, the model can also recognize the two novel categories.

the application of Deep Neural Networks (DNNs) [16, 57]. In response to the threats, various techniques have been proposed to enhance adversarial robustness of DNNs. Among these, Adversarial Training (AT) and its variants [32, 36, 37, 52] have shown to be highly effective [16, 42, 57]. However, applying AT from scratch to novel datasets or categories can be time-consuming. In real-world scenarios, it is common for models to encounter examples from novel categories, known as zero-shot learning (ZSL) setting [27]. If adversarially trained models possess zero-shot capability, it can significantly enhance the scalability and practical utility of adversarially robust models.

Adversarial robustness in the ZSL setting is challenging. Earlier attempts [51, 55] try to combine attribute-based ZSL methods with AT. However, these methods yielded modest improvements and lacked scalability due to the practical difficulties in obtaining attributes. Recently, large-scale

vision-language models such as CLIP [39] have made significant strides and are considered foundational models for various downstream applications. These models, trained on extensive text-image data, demonstrate strong zero-shot capability by using textual descriptions to recognize new visual categories, eliminating the need for explicit reliance on image attributes. In this work, different from previous attribute-based ZSL methods, we explore the utilization of foundation models to introduce zero-shot adversarial robustness to DNNs.

We propose a Language-driven, Anchor-based Adversarial Training strategy, denoted as LAAT, to enhance the zero-shot adversarial robustness. Specifically, we use the popular text encoder from vision-language models, specifically CLIP [39], to obtain l_2 normalized feature embedding (named *anchor* in this paper) for the text (label name) corresponding to image categories. The text encoder exhibits a property called *semantic consistency*, where semantically similar categories are mapped to neighboring anchors in the feature space. Then we utilize the text anchors to supervise the image classification model with AT. After this process, the image model not only obtains adversarial robustness on seen categories but also aligns the image features with the text anchors. During zero-shot inference, by leveraging the semantic consistency between seen and novel categories, the image model can recognize novel categories while maintaining adversarial robustness. Fig. 1 illustrates the idea. We note that one recent work, TeCoA [34], also utilizes the ability of large-scale vision-language models to enhance the zero-shot adversarial robustness. However, TeCoA was only shown effective for extremely tiny adversarial noise (*i.e.*, $\epsilon = 1/255$).

As CLIP models are not designed for adversarial robustness, directly applying the obtained anchors from CLIP text encoders to anchor-based AT [36, 37] (refer to Sec. 2.1) results in poor performance. We attribute it to the high Cosine Similarity (CoS) problem, *i.e.*, the average cosine similarity between CLIP anchors are too high to be directly used in AT. See Sec. 3.3 for details. This problem has never been discovered in previous work utilizing CLIP models [22, 28, 34, 45]. To address it, we propose several techniques. We first introduce an *expansion* algorithm that maps the original CLIP anchors to new anchors with increased distances while maintaining semantic consistency. We then incorporate an Alignment Cross-Entropy (A-CE) loss as the optimization objective. In addition, we use a smoothness loss to enhance the robustness on unseen categories.

The experimental results showed that the models trained with LAAT achieved remarkable zero-shot performance against strong adversarial perturbations, even surpassing the previous state-of-the-art (SOTA) adversarially robust few-shot methods. We also trained several models with LAAT on a large dataset ImageNet-1K [40]. These mod-

els showed substantial adversarial robustness across various downstream datasets in the zero-shot setting, significantly outperforming the recent TeCoA method [34]. The main contributions can be summarized as follows:

- We identify the high CoS problem in CLIP anchors for adversarial robustness and propose effective techniques to alleviate it.
- The models trained by LAAT show substantial adversarial robustness across downstream datasets, suggesting that AT in the ZSL setting could be a promising way to improve the practical usefulness of AT.

2. Related Work

2.1. Anchor-based AT

AT has become one of the most effective strategies to improve adversarial robustness [3]. Several studies try to improve AT from various aspects [26, 29, 30, 36, 37, 52]. One series of works focuses on refining cross-entropy loss, which is the most commonly used loss in classification, to obtain better robustness with AT [36, 37]. Pang et al. [36] boosts AT by introducing the Max-Mahalanobis center (MMC) loss, a Euclidean loss where there is a fixed feature vector μ_y for each category y , and the optimization object is $\|\mathbf{z} - \mu_y\|_2^2$, where \mathbf{z} is the output feature. Pang et al. [37] uses Hypersphere Embedding to boost AT. It normalizes both weights \mathbf{W} and features \mathbf{z} of the output layer, and proposes two losses in the form of $\cos \theta$ and θ , where θ is the angle between $\frac{\mathbf{W}}{\|\mathbf{W}\|_2}$ and $\frac{\mathbf{z}}{\|\mathbf{z}\|_2}$. We call these AT methods *anchor-based* methods as they all used l_2 normalized feature embeddings (anchors) for training. Besides, several metric learning works [33, 56] also used a similar paradigm to anchor-based AT for enhancing robustness.

2.2. Language-driven Recognition.

Earlier ZSL models exploit auxiliary information, commonly in the form of image attributes, to transfer knowledge between seen and novel categories [47, 50]. Recently, language-driven recognition has caught attention in the ZSL field, especially with the recent advance of large-scale pre-trained vision-language models [24, 39]. CLIP [39] demonstrated that classic recognition tasks can strongly benefit from millions of raw texts pertaining to images. CLIP uses an independent image encoder and text encoder to perform contrastive learning on 400M image-text pairs. After pre-training, natural language is used to reference visual concepts by computing the CoS between text features and the visual feature, enabling the zero-shot transfer of the model to downstream classification datasets. The basic paradigm of CLIP can be extended to several tasks such as semantic segmentation and object detection [22, 28, 45]. However, recent works show that CLIP models lack adversarial robustness and can be easily attacked [4, 34, 53].

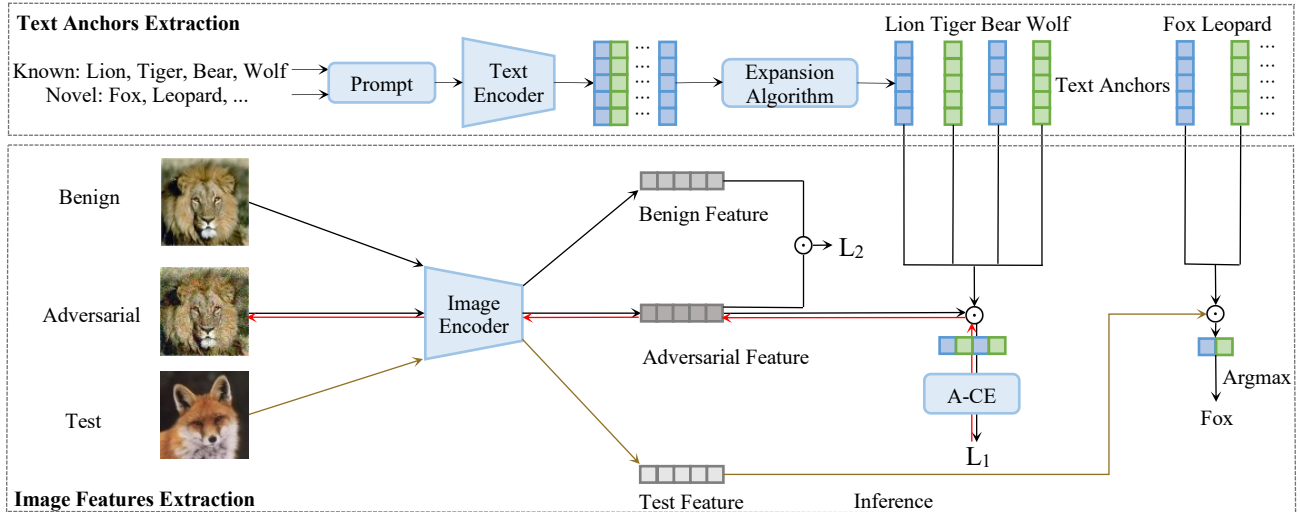


Figure 2. The pipeline of LAAT. Only the image encoder is trainable in the figure. \odot indicates computing the CoS. Red arrows indicate the adversarial example generation process and brown arrows indicate the inference.

2.3. Adversarially Robust Zero/few-shot Learning

An early preprint work [55] combines AT with attribute-based ZSL methods but only gets a mild improvement over non-AT models. In addition, the method is hard to generalize without the attribute annotations. Yucel et al. [51] evaluates the adversarial robustness of classic attribute-based zero-shot models and concludes that adversarial robustness for ZSL field is suffering an immature state. Mao et al. [34] proposes several techniques such as visual prompting and fine-tuning to adapt large-scale models, *e.g.*, the image models of CLIP, for zero-shot adversarial robustness, while their method called TeCoA was only shown to be effective under extremely tiny adversarial noise. Different from TeCoA, LAAT only uses the anchors from the CLIP text encoders.

A closely related area is adversarially robust few-shot learning, which has caught attention recently, too. We also compare LAAT with recent few-shot methods directly to comprehensively evaluate it. AQ [19], R-MAML [46], and ITS-MAML [15] study meta-learning jointly with AT. GR [13] tries to learn a generalizable robust representation by the combination of several AT techniques.

3. LAAT

3.1. Overall Pipeline

LAAT first obtains fixed text anchors from the text encoder, then uses them to adversarially supervise the training of image classification model. After AT, the image model will obtain adversarial robustness on seen categories and also align the image features with the text anchors to obtain zero-shot ability (see Fig. 1). The overall pipeline of LAAT is

illustrated in Fig. 2. Given the original text anchors, LAAT uses an expansion algorithm to obtain fixed Ground-Truth (GT) anchors suitable for AT, then uses an image encoder to extract both benign and adversarial features. The CoS between adversarial features and GT anchors are maximized by minimizing the A-CE loss (L_1). Besides, a smoothness loss (L_2) is applied to the adversarial and benign features. We then describe each module of LAAT.

3.2. Text Anchor Extraction

Given the names of N categories in the training set, they are first filled into N sentences with a fixed prompt text, such as “A photo of { }”. After that, these sentences are encoded into N anchors $\{\mathbf{a}_i\}_{i=1}^N$ by a fixed pre-trained text encoder, where $\mathbf{a}_i \in \mathbb{R}^n$ and $\|\mathbf{a}_i\|_2 = 1$. Many architectures are feasible if they have semantic consistency as mentioned before and in this work, we use the large-scale pre-trained CLIP text encoder. However, CLIP anchors cannot be used directly due to the high CoS problem, as analyzed below.

3.3. High Cosine Similarity Problem

We first make an intriguing observation that directly using CLIP anchors for anchor-based AT poses significant challenges in achieving convergence. This phenomenon is universal among different CLIP text encoders and different models. To illustrate this, we adversarially trained a ResNet18 [23] on Cifar100 [25] by maximizing the $\cos \theta$ objective [37], *i.e.*, the CoS between the model output \mathbf{z} and the GT anchor \mathbf{a}_y [37], which is computed as $\frac{\mathbf{z}^T \mathbf{a}_y}{\|\mathbf{z}\|_2}$ (note that $\|\mathbf{a}_y\|_2 = 1$). The experimental details can be found in Appendix A. Fig. 3 shows the learning curves with CLIP anchors (the orange curves). For comparison, here we also




Figure 3. Learning curves of AT supervised by $\cos \theta$ with fixed anchors generated from MMC method, CLIP text encoder, and the expansion algorithm (see Sec. 3.4).

show the learning curves with MMC [36] anchors (the blue curves) in the same training setting. MMC anchors are pre-computed by maximizing the distance between anchors [36] and thus the average CoS between MMC anchors is close to 0 (e.g., -0.01 on Cifar100). The accuracy with CLIP anchors increased quite slowly. Even after 200 epoch training, the accuracy was far lower than that of MMC anchors (note that zero-shot recognition cannot be achieved with MMC anchors for the lack of semantic consistency).

The stark discrepancy between the two learning curves motivates us to analyze the essential difference between the two sets of anchors. Inspired by the anchor generation process of MMC, we calculated the CoS between CLIP anchors and found them to be considerably larger than that of MMC anchors (-0.01). The first row of Tab. 1 presents the average CoS between anchors from Cifar100 categories generated by different CLIP models. Note that the high CoS between CLIP anchors is irrelevant to the prompt text. We found that even using several random sentences from Wikipedia to encode anchors by CLIP text encoders, the average CoS remained significantly large.

Although high CoS seems to have no effect on downstream tasks on clean images [22, 28, 45], we discovered that AT was remarkably affected because higher CoS implies that the anchors are closer to the decision boundary, which could leave more room for adversarial examples [36, 52]. This problem may also explain why TeCoA [34], which also used the text anchors during adaptation, failed to convert CLIP image models to adversarially robust models against strong perturbations. Importantly, random or MMC anchors cannot mitigate this issue as they lack the semantic consistency provided by CLIP anchors and thus cannot be used in the ZSL setting. Therefore, the micro designs in our LAAT approach primarily aim to address the challenges associated with CLIP anchors.

3.4. Expansion Algorithm

We first propose an expansion algorithm to increase the distances between anchors (reducing the CoS). The expansion algorithm can remap the anchors to increase the dis-

Type	RN50x4	RN50x16	ViT-B/32	ViT-B/16	ViT-L/14
Original	0.700	0.710	0.779	0.761	0.746
Expanded	0.238	0.253	0.199	0.195	0.222

Table 1. Average CoS between anchors before and after expansion algorithm (see Sec. 3.4) from Cifar100 categories generated by different CLIP text encoders.

tances between anchors while preserving the semantic consistency between them. High CoS means the normalized anchors $\{\mathbf{a}_i\}_{i=1}^N$ are dispersed over a cluster of the unit hypersphere (see Fig. 4). As we expect the remapped anchors to be also on the unit hyper-sphere, the expansion algorithm should be designed under spherical coordinates. Let us recap the transformation between n -dimensional spherical coordinate system and Cartesian coordinates. Suppose $(x^{(1)}, x^{(2)}, \dots, x^{(n)})$ is the Cartesian coordinate, we can compute them by:

$$\begin{aligned}
 x^{(1)} &= r \cos \phi^{(1)}; & x^{(2)} &= r \sin \phi^{(1)} \cos \phi^{(2)}; & \dots; \\
 x^{(n-1)} &= r \sin \phi^{(1)} \dots \sin \phi^{(n-2)} \cos \phi^{(n-1)}; \\
 x^{(n)} &= r \sin \phi^{(1)} \dots \sin \phi^{(n-2)} \sin \phi^{(n-1)},
 \end{aligned} \tag{1}$$

where r is the radial coordinate and $\phi^{(1)}, \dots, \phi^{(n-1)}$ are angular coordinates [6]. Here $\phi^{(1)}$ is the polar angle (taking the three-dimensional spherical coordinate system like longitude and latitude system as an example, the polar angle is the angle with respect to the z axis). If the anchors are clustered around the polar, i.e., $\phi_0 := \max_i \phi_i^{(1)} < \frac{\pi}{2}$, where $\phi_i^{(1)}$ denotes the polar angle of anchor \mathbf{a}_i , a natural expansion method is to enlarge the polar angle $\phi_i^{(1)}$ to $\frac{\pi}{2} \cdot \phi_i^{(1)} / \phi_0$ and keep the other angular coordinates unchanged. Then the cluster can be expanded to the whole hemisphere. Note that we cannot expand them larger than $\frac{\pi}{2}$ as anchors may be close to each other from another hemisphere, which may hurt the semantic consistency of anchors.

Thus, given the *original anchor* \mathbf{a}_i , the first step of the expansion is to calculate a rotation matrix \mathbf{R} to transform the center \mathbf{v} of the clustered \mathbf{a}_i 's ($\mathbf{v} := \sum_{i=1}^N \mathbf{a}_i / \|\sum_{i=1}^N \mathbf{a}_i\|_2$) to the polar $\mathbf{p} = [1, 0, \dots, 0]^T$, and the *rotated anchor* $\tilde{\mathbf{a}}_i$ is computed by: $\tilde{\mathbf{a}}_i = \mathbf{R} \mathbf{a}_i$. Note that computing \mathbf{R} between two unit vectors has several ways and we use the Aguilera and Pérez-Aguila [1] algorithm. As $\tilde{\mathbf{a}}_i$'s are clustered around the polar \mathbf{p} , we then enlarge the polar angle of $\tilde{\mathbf{a}}_i$ as described above. Denoting the j -th element of \mathbf{a}_i as $a_i^{(j)}$, the largest polar angle between $\tilde{\mathbf{a}}_i$'s and \mathbf{p} can be computed as $\phi_0 = \max_{1 \leq i \leq N} \{\arccos \tilde{a}_i^{(1)}\}$. According to Eq. (1), the Cartesian coordinates of *expanded anchor* $\bar{\mathbf{a}}_i$ is given by:

$$\begin{aligned}
 \bar{a}_i^{(1)} &= \cos \bar{\phi}_i^{(1)}; \\
 \bar{a}_i^{(j)} &= \tilde{a}_i^{(j)} \cdot \sin \bar{\phi}_i^{(1)} / \sin \tilde{\phi}_i^{(1)}, \quad j = 2, \dots, n,
 \end{aligned} \tag{2}$$




Figure 4. An illustration of the expansion operation in 3D space.

where $\tilde{\phi}_i^{(1)} = \arccos \tilde{a}_i^{(1)}$ denotes the polar angle of each \tilde{a}_i and $\bar{\phi}_i^{(1)} = \frac{\pi}{2} \cdot \tilde{\phi}_i^{(1)} / \phi_0$ denotes the expanded polar angle. Fig. 4 illustrates this operation and the detailed derivation is shown in Appendix B. In the end, we use the inverse of \mathbf{R} to map the expanded anchor \tilde{a}_i to the original locations and obtain the *final anchor*: $\hat{a}_i = \mathbf{R}^T \tilde{a}_i$. The pseudo-code of the whole expansion algorithm is provided in Appendix C. As shown, the expansion algorithm can increase the distances between different anchors while maintaining the semantic consistency of the original CLIP anchors (close anchors are still close after expansion).

The second row of Tab. 1 shows the average CoS between anchors after the expansion algorithm. Compared with the original CLIP anchors, the average CoS between the expanded CLIP anchors is greatly reduced. Fig. 3 shows the learning curves with the expanded CLIP anchors (the green curves). The final robust accuracy (Fig. 3, right) is quite close to that of MMC anchors.

3.5. Supervision Objective

We propose to improve the anchor-based AT methods further by introducing an A-CE loss into the supervision objective. The A-CE loss is a composition of CE loss and the anchor alignment process (the $\cos \theta$ objective), which can also be seen as a special case of InfoNCE [35] with using l_2 normalized features. This could be confusing considering that previous works have shown that CE is inferior to other anchor-based objectives such as the θ and $\cos \theta$ [37] in the standard AT setting (see Sec. 2.1). However, we find that with the high CoS problem, A-CE can boost the performance of anchor-based AT methods in that the softmax function of A-CE could make a relaxation on the anchor alignment objective of the original anchor-based AT methods. Assuming that the image encoder $f_\Theta(\cdot)$ outputs normalized features $\mathbf{z} = f_\Theta(\mathbf{x}) \in \mathbb{R}^n$ and $\|\mathbf{z}\|_2 = 1$ and taking the $\cos \theta$ objective as an example, the $\cos \theta$ objective is to maximize $\mathbf{z}^T \hat{\mathbf{a}}_y$, where $\hat{\mathbf{a}}_y$ denotes the GT anchor. It encourages the visual feature’s CoS to be as close as possible to the GT anchor, while the A-CE loss, by putting $\cos \theta$ into the softmax function, only encourages the visual feature’s CoS with GT to be larger than those with other anchors. With A-CE supervision, the output feature could be opti-

mized to be far away from all anchors, but it can still be correctly classified as long as it is closest to the GT anchor. This will be discussed further in Sec. 4.4. Overall, the supervision objective is:

$$L_1 = \mathbb{E}_{(\mathbf{x}, y)} \left[-\log \frac{\exp(f_\Theta(\mathbf{x} + \delta)^T \hat{\mathbf{a}}_y)}{\sum_{i=1}^N \exp(f_\Theta(\mathbf{x} + \delta)^T \hat{\mathbf{a}}_i)} \right], \quad (3)$$

where δ is adversarial perturbation generated by $\frac{\partial L_1}{\partial \mathbf{x}}$. Unlike previous works [39] on standard training, we do not use a temperature parameter τ to scale the CoS. This is discussed in Appendix D.

3.6. Smoothness Loss

Since we prefer the adversarial robustness of the model in the zero-shot transfer setting, we additionally introduce a smoothing loss, encouraging the CoS between adversarial features to be similar to the benign features of an image:

$$L_2 = \mathbb{E}_{(\mathbf{x}, y)} [-f_\Theta(\mathbf{x})^T f_\Theta(\mathbf{x} + \delta)]. \quad (4)$$

The smoothing loss is independent of the category labels of training examples. It can be expected to improve the adversarial robustness on novel categories. Finally, the whole loss in the training stage is given by $L = L_1 + \alpha L_2$, where α is a hyper-parameter.

3.7. Performing Zero-Shot Recognition

When performing zero-shot recognition on novel categories, the image features of a test image are generated by the image encoder, and the anchors for novel categories are generated by the text encoder and the expansion algorithm. Note that the internal parameters of the expansion algorithm, such as ϕ_0 , are pre-computed using anchors from the training categories. Finally, the zero-shot prediction is the category whose text anchor exhibits the highest CoS with the image features.

4. Experiments

4.1. Experimental Setup

Baseline. We first consider the standard ZSL setting. Due to the lack of adversarially robust methods in this setting, we compared LAAT with several robust few-shot methods directly: AQ [19], R-MAML [46], and GR [13]. ITS-MAML [15] was not compared as it used a different attack setting and no code was released. Note that zero-shot is harder than few-shot as the latter can have access to a few examples from the novel categories.

We then considered a more pragmatic version of ZSL, which requires recognizing examples from both seen and novel categories. This is called the generalized zero-shot learning (GZSL) setting [8]. We trained several models with

Dataset	Method	Setting	Clean	FGSM	PGD	CW	AA
CIFAR-FS	AQ [19]	1-shot	54.31 ± 0.50	32.47 ± 0.47	30.41 ± 0.46	30.14 ± 0.46	27.09 ± 2.07
	R-MAML [46]	1-shot	38.53 ± 0.47	26.76 ± 0.44	26.21 ± 0.43	29.76 ± 0.46	16.67 ± 5.48
	GR [13]	1-shot	45.27 ± 0.49	39.60 ± 0.46	38.03 ± 0.46	37.00 ± 0.46	-
	LAAT	zero-shot	55.60 ± 0.46	41.17 ± 0.44	40.12 ± 0.44	37.35 ± 0.44	36.45 ± 1.90
	LAAT	1-shot	57.88 ± 0.45	43.73 ± 0.44	42.74 ± 0.44	40.10 ± 0.45	39.23 ± 2.12
MiniImageNet	AQ [19]	1-shot	39.17 ± 0.40	22.53 ± 0.35	20.95 ± 0.34	19.33 ± 0.38	17.85 ± 1.41
	R-MAML [46]	1-shot	36.18 ± 0.36	23.85 ± 0.46	21.74 ± 0.34	19.68 ± 0.40	13.69 ± 3.95
	GR [13]	1-shot	36.14 ± 0.45	29.23 ± 0.33	27.57 ± 0.38	26.61 ± 0.33	-
	LAAT	zero-shot	46.97 ± 0.35	30.69 ± 0.30	29.27 ± 0.31	27.91 ± 0.29	25.24 ± 1.42
	LAAT	1-shot	47.37 ± 0.35	31.55 ± 0.31	29.93 ± 0.31	28.19 ± 0.29	25.75 ± 1.45

Table 2. The comparison of LAAT in 5-way zero-shot setting and previous works in 5-way 1-shot setting under $\epsilon = 8/255$. The 95% confidence interval was reported. Here all methods used Conv4-512 as the image model. The accuracy of GR under AA is not shown as it was not reported and no code was released. We also list the performance of LAAT in the 1-shot setting with image-text blended anchors (discussed in Sec. 4.3). The results higher than previous methods in each column are in **bold**.

Method	Model	Training	Cifar100	aPY	AwA2	COCO Obj.	STL10	OxfordPet	DTD	Caltech101	Caltech256	SUN
			100	32	50	80	10	37	47	101	257	397
TeCoA [34]	ViT-B/32*	$\epsilon = 1/255$	0.00	4.39	0.88	1.07	0.98	0.00	1.43	1.66	1.33	0.00
	ViT-B/32	$\epsilon = 8/255$	0.00	5.27	1.45	1.33	1.46	0.28	1.89	2.97	2.36	0.14
	ViT-B/16	$\epsilon = 8/255$	0.20	6.45	1.95	1.56	1.66	0.49	2.25	7.03	3.81	0.29
LAAT	ViT-B/32	$\epsilon = 8/255$	6.64	24.32	23.34	18.56	41.12	20.31	4.59	34.43	21.39	4.30
	ViT-B/16	$\epsilon = 8/255$	7.32	34.97	26.21	20.61	49.41	28.81	4.79	37.89	24.02	5.18

Table 3. Adversarial accuracy (%) of models trained on ImageNet-1K on ten downstream datasets in GZSL setting. The number of categories is shown below the dataset name. * indicates the results of the originally released checkpoint of TeCoA with the training perturbation $\epsilon = 1/255$. The robustness was all evaluated by PGD-20 under $\epsilon = 8/255$.

LAAT on ImageNet-1K and evaluated them directly on ten downstream datasets. In this setting, we mainly compare LAAT with TeCoA [34]. We did not compare with Yucel et al. [51] and Zhang et al. [55] as they relied on image attributes and were unsuitable for this setting.

Datasets. When comparing with few-shot methods, we used two popular few-shot benchmarks CIFAR-FS [5] and MiniImageNet [43]. CIFAR-FS, a variant of Cifar100 [25], contains 64, 16, and 20 categories for training, validation, and testing, respectively. MiniImageNet, a subset of ImageNet-1k [40], contains 600 images of 84×84 size per category. We used the same dataset splitting setting as previous works [13, 19, 46], for training and testing.

In the GZSL setting, we tested models trained on ImageNet-1K on ten downstream datasets such as AwA2 [48] and aPY [17]. Note that all of these datasets include several novel categories that have never appeared in ImageNet-1K, e.g., *statue* and *ass* in aPY [17], *blue whale* and *dolphin* in AwA2 [48]. More introduction to the datasets is shown in the Appendix E.

Implementation details. Unless otherwise stated, we used the CLIP ViT-B/16 text encoder and the prompt text was set as “This is a photo of a { }”. α was set to 3.

In the standard ZSL setting, we followed previous few-

shot works [13, 46] and used the widely applied model Conv4-512 [54] as the image model. The training recipe strictly followed GR [13]. We performed AT by generating adversarial examples via PGD method with maximum adversarial perturbation $\epsilon = 8/255$ in l_∞ -norm bound, with $T = 7$ iterative steps (step size $s = 2/255$).

In the GZSL setting, we trained two image models ViT-B/32 and ViT-B/16 [14] with LAAT on the large ImageNet-1K. The training basically followed the recipe proposed by Debenedetti et al. [12]. Here we performed AT by generating adversarial examples via PGD method with $\epsilon = 8/255$ in l_∞ -norm bound and $T = 2$ iterative steps ($s = 8/255$).

4.2. Comparison Results

The Standard ZSL Setting. We performed experiments in the standard and usual 5-way one/zero-shot setting [13, 19, 46]. Following previous works [19, 46], we report the robustness of our method against common adversarial attacks, including FGSM [20], PGD [32] with 20 iterative steps and step size $s = 2/255$ (denoted as PGD-20 in the following text), CW [7] (optimized by PGD for 30 steps with step size $s = 0.8/255$), and AutoAttack (AA) [11], an effective ensemble attack (several variants of PGD and one query-based black-box Square attack [2]) used to assess adversarial robustness. The maximal perturbation was set to commonly used $\epsilon = 8/255$ in l_∞ setting. The experi-

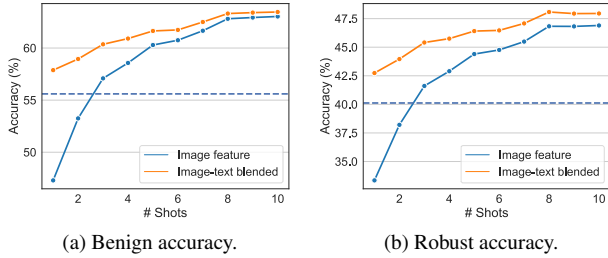


Figure 5. Classification accuracy on both benign and adversarial examples in 5-way few-shot setting on CIFAR-FS, with image feature anchors or with image-text blended anchors. The dashed line denotes 5-way zero-shot accuracy.

mental results are reported from 2000 randomly sampled 5-way tasks. The results presented in Tab. 2 demonstrate that LAAT in the zero-shot setting achieved competitive performance with several methods in the few-shot setting. Notably, the adversarial robustness of LAAT models even surpassed that of the previous SOTA adversarially robust few-shot method GR [13]. Additionally, our method exhibited significantly better clean accuracy than all three few-shot methods. We note that these results should be interpreted cautiously because LAAT and these few-shot works used different auxiliary information: the former used text information of the test category names and the latter relied on a small number of example images from the test categories.

GZSL Setting across Datasets. In the GZSL setting, the models were evaluated in N -way classification setting, where N is the number of categories in each dataset. All images on downstream datasets are resized to 224×224 in inference. Here we compared LAAT with TeCoA. TeCoA originally performed adversarial fine-tuning on ImageNet-1K with maximum perturbation $\epsilon = 1/255$. For a fair comparison, here adversarial fine-tuning was also performed with the released code of TeCoA [34]. Tab. 3 shows the zero-shot results on ten downstream datasets under PGD-20. Here the accuracy is reported by sampling 100 examples per category. We can see that TeCoA almost failed to train adversarially robust models in the common ϵ setting, while the models trained with LAAT can obtain substantial adversarial robustness on several downstream datasets.

Note that the models trained with LAAT did not perform quite well on Cifar100, DTD [9], and SUN [49]. This could be caused by the original small resolution of Cifar100 (32×32) [12] and the large inter-class differences between DTD, SUN, and ImageNet-1K.

4.3. Extending LAAT to Few-Shot Setting

Our zero-shot models supervised by LAAT can be naturally extended to the few-shot setting and the robustness can be further improved. Given the images of novel categories as

Objective	Clean	Robust
$Eucl$	39.17	28.04
θ	37.92	27.84
$\cos \theta$	41.91	27.91

Table 4. Classification accuracy (%) on CIFAR-FS supervised by the original CLIP anchors with different optimization objectives.

EXP	A-CE	SM	Clean	Robust
			41.91	27.91
✓			46.98	32.71
✓	✓		54.02	37.69
✓	✓	✓	55.60	40.12

Table 5. Classification accuracy (%) of Conv4-512 on CIFAR-FS in different ablation experiments. The first two lines were the results obtained with $\cos \theta$ as the objective.

the support set, we can follow the previous few-shot SOTA method [13] and use the prototype-based metric learning [41] to build image feature-based anchors. That is, in the K -shot case, the image anchor \mathbf{v}_y for a novel category y can be computed as $\mathbf{v}_y = \text{Norm}_2\{\sum_{i=1}^K f_{\Theta}(\mathbf{x}_i)\}$, where \mathbf{x}_i denotes the images of y in the support set and $\text{Norm}_2\{\cdot\}$ denotes the l_2 normalization. We can also weight the original text anchor \mathbf{a}_y from the text encoder and the image feature from the support set to build an image-text blended anchor for y : $\hat{\mathbf{a}}_y^* = \text{Norm}_2\{\beta \cdot \hat{\mathbf{a}}_y + \sum_{i=1}^K f_{\Theta}(\mathbf{x}_i)\}$. Here the text anchor can be regarded as the *prior knowledge* of y . We set $\beta = 2$ in all experiments.

Tab. 2 shows the results of different models with image-text blended anchors in the 1-shot setting. The 1-shot performance of LAAT was better than the zero-shot performance. Fig. 5 further shows the results with two different anchors $\hat{\mathbf{a}}^*$ and \mathbf{v} in K -shot settings. We can see that the text anchor used to perform zero-shot recognition boosted the performance in various K -shot settings, especially when the support set is small.

4.4. Ablation Study

Limited by computing resources, the ablation study is not studied on ImageNet-1K. Unless specified, it was performed with Conv4-512 in the 5-way zero-shot setting. The robustness was evaluated by the PGD-20 attack method.

Effectiveness of each design. In Fig. 3, we have shown that if the model is supervised by CoS between the original CLIP anchors and the output feature, say, $\cos \theta$, it is difficult to converge under AT. In Tab. 4, we show the results of other anchor-based AT optimization objectives, such as the angle [37] and the Euclidean distance [36], denoted as θ and $Eucl$ respectively, in the ZSL setting. The results indicate that none of these three optimization objectives led to good classification accuracy. Note that random guessing in

AT Method	Hyper-Parameter	Clean	Robust
EXP + A-CE + TRADES	$1/\lambda = 1$	55.26	19.22
	$1/\lambda = 6$	55.82	25.02
EXP + A-CE + SM	$\alpha = 1.0$	55.62	39.95
	$\alpha = 3.0$	55.60	40.12
	$\alpha = 6.0$	53.01	40.29

Table 6. Classification accuracy (%) on CIFAR-FS with different AT methods.

Method	Setting	Conv4-64		ResNet12	
		Clean	Robust	Clean	Robust
AQ	1-shot	43.74	27.54	43.58	31.45
R-MAML	1-shot	32.46	24.61	42.54	32.54
GR	1-shot	44.51	37.45	48.13	39.29
LAAT	zero-shot	49.84	38.74	58.37	42.88

Table 7. Classification accuracy (%) on CIFAR-FS with different image models.

this setting can obtain 20% accuracy.

We then conducted an ablation study to show the effectiveness of each design in LAAT: the expansion algorithm, the A-CE loss, and the smoothness loss, denoted as EXP, A-CE, and SM, respectively. The results in Tab. 5 show that all of these designs are helpful in improving the zero-shot adversarial robustness.

As stated in Sec. 3.5, A-CE could make a relaxation on the original anchor-based AT objectives. To validate this, we computed the average CoS between GT anchors and the output visual features of the two models trained: EXP and EXP + A-CE (EXP was supervised by $\cos\theta$; see Tab. 5). With A-CE loss, the visual feature’s CoS with GT anchors dropped from 0.455 to 0.205 on average, while the accuracies on both benign and adversarial examples were improved (see Tab. 5). These results support our conjecture that some examples not quite close to the GT anchor can still be classified correctly with A-CE supervision.

Smoothness v.s. TRADES. It is seen that SM has similar objectives with TRADES [52], but they are different in many aspects (see Appendix F). Here we also compare SM with TRADES empirically. The experiments were performed based on expanded anchors with A-CE supervision. The results are shown in Tab. 6, together with the most important hyper-parameter of each method. We can clearly see that TRADES performed much worse than our smoothness loss. In addition, the results show that the effectiveness of smoothing loss was not quite sensitive to α .

Different image models. We also performed experiments with other popular image models in the few-shot setting other than Conv4-512, e.g., Conv4-64 [44] and ResNet12 [19]. Tab. 7 shows the results of the two image models on CIFAR-FS. Conv4-64 is smaller than Conv4-512 while

Text Encoder	Embed-dim	Clean	Robust
None (MMC, one-shot)	512	47.55	32.79
RN50x4	640	52.28	38.65
RN50x16	768	53.86	38.02
ViT-B/32	512	52.68	38.39
ViT-B/16	512	55.60	40.12
ViT-L/14	768	52.78	39.48

Table 8. Classification accuracy (%) supervised by MMC anchors in one-shot setting and the full LAAT with different CLIP text encoders in zero-shot setting on CIFAR-FS. Embed-dim denotes the dimension of the text anchors.

ResNet12 is slightly larger. The models trained with LAAT showed stronger zero-shot adversarial robustness than those few-shot methods regardless of the model size.

Different text encoders. Without text encoders, it is impossible for LAAT to achieve zero-shot robustness. In principle, arbitrary text encoders are feasible for LAAT if they have semantic consistency. We show the influence of using different CLIP text encoders in Tab. 8. We also show the results without using any text encoders. Instead, the image model was supervised by the MMC anchors (here zero-shot recognition cannot be performed while one-shot results are reported). We can see that using any text encoders can obtain strong zero-shot robustness, significantly surpassing the few-shot robustness without using text encoders. We also studied the relationship between the distances of different CLIP models’ anchors and Cifar100 super-categories. The results in Appendix G further show that semantic consistency of the text encoder is one of the keys to zero-shot adversarial robustness.

5. Conclusion and Discussion

In this work, we propose a novel LAAT strategy for adversarially robust image classification in a zero-shot setting, inspired by recent vision-language models. Extensive experiments demonstrate that our method has strong adversarial robustness across different models and several zero-shot settings. We hope our work could encourage more researchers to investigate the robustness in the zero-shot setting, which could be a promising way to improve the scalability and flexibility of adversarially robust models. Limited by the data and computing resources, we cannot perform AT on large-scale image-text pairs directly like CLIP [39]. We believe it could be another promising way to achieve zero-shot adversarial robustness.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China (Nos. U2341228 and U19B2034).

References

- [1] Antonio Aguilera and Ricardo Pérez-Aguila. General n-dimensional rotations. In *WSCG*, pages 1–8, 2004. 4
- [2] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: A query-efficient black-box adversarial attack via random search. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 484–501, 2020. 6
- [3] Anish Athalye et al. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Int. Conf. Mach. Learn. (ICML)*, pages 274–283, 2018. 2
- [4] Yuanhao Ban and Yinpeng Dong. Pre-trained adversarial perturbations. *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2022. 2
- [5] Luca Bertinetto, João F. Henriques, Philip H. S. Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *Int. Conf. Learn. Represent. (ICLR)*, 2019. 6, 3
- [6] LE Blumenson. A derivation of n-dimensional spherical coordinates. *The American Mathematical Monthly*, 67(1):63–66, 1960. 4
- [7] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy, SP*, pages 39–57, 2017. 6
- [8] Wei-Lun Chao, Soravit Changpinyo, Boqing Gong, and Fei Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 52–68, 2016. 5
- [9] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3606–3613, 2014. 7, 2
- [10] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, pages 215–223, 2011. 2
- [11] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Int. Conf. Mach. Learn. (ICML)*, pages 2206–2216, 2020. 6
- [12] Edoardo DeBenedetti et al. A light recipe to train robust vision transformers. *arXiv preprint arXiv:2209.07399*, 2022. 6, 7
- [13] Junhao Dong, Yuan Wang, Jianhuang Lai, and Xiaohua Xie. Improving adversarially robust few-shot image classification with generalizable representations. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 9015–9024. IEEE, 2022. 3, 5, 6, 7
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. 6
- [15] Xiaoyue Duan, Guoliang Kang, Runqi Wang, Shumin Han, Song Xue, Tian Wang, and Baochang Zhang. Rethinking the number of shots in robust model-agnostic meta-learning. *arXiv preprint arXiv:2211.15180*, 2022. 3, 5
- [16] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1625–1634, 2018. 1
- [17] Ali Farhadi, Ian Endres, Derek Hoiem, and David A. Forsyth. Describing objects by their attributes. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1778–1785, 2009. 6, 2
- [18] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPRW*, pages 178–178, 2004. 2
- [19] Micah Goldblum et al. Adversarially robust few-shot learning: A meta-learning approach. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. 3, 5, 6, 8
- [20] Ian J. Goodfellow et al. Explaining and harnessing adversarial examples. In *Int. Conf. Learn. Represent. (ICLR)*, 2015. 6
- [21] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007. 2
- [22] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. 2, 4
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 770–778, 2016. 3
- [24] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Int. Conf. Mach. Learn. (ICML)*, pages 4904–4916, 2021. 2
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *California Institute of Technology*, 2009. 3, 6, 2
- [26] Huafeng Kuang, Hong Liu, Yongjian Wu, and Rongrong Ji. Semantically consistent visual representation for adversarial robustness. *IEEE Transactions on Information Forensics and Security*, 2023. 2
- [27] Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 36(3):453–465, 2014. 1
- [28] Boyi Li, Kilian Q. Weinberger, Serge J. Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. In *Int. Conf. Learn. Represent. (ICLR)*, 2022. 2, 4
- [29] Xiao Li, Hang Chen, and Xiaolin Hu. On the importance of backbone to the adversarial robustness of object detectors. *arXiv preprint arXiv:2305.17438*, 2023. 2
- [30] Xiao Li, Ziqi Wang, Bo Zhang, Fuchun Sun, and Xiaolin Hu. Recognizing object by components with human prior knowledge enhances adversarial robustness of deep neural networks. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 2023. 2
- [31] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in

- context. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 740–755, 2014. [2](#)
- [32] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Int. Conf. Learn. Represent. (ICLR)*, 2018. [1](#), [6](#)
- [33] Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, and Baishakhi Ray. Metric learning for adversarial robustness. *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 32, 2019. [2](#)
- [34] Chengzhi Mao, Scott Geng, Junfeng Yang, Xin Wang, and Carl Vondrick. Understanding zero-shot adversarial robustness for large-scale models. *Int. Conf. Learn. Represent. (ICLR)*, 2023. [2](#), [3](#), [4](#), [6](#), [7](#)
- [35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [5](#)
- [36] Tianyu Pang, Kun Xu, Yinpeng Dong, Chao Du, Ning Chen, and Jun Zhu. Rethinking softmax cross-entropy loss for adversarial robustness. In *Int. Conf. Learn. Represent. (ICLR)*, 2020. [1](#), [2](#), [4](#), [7](#)
- [37] Tianyu Pang, Xiao Yang, Yinpeng Dong, Taufik Xu, Jun Zhu, and Hang Su. Boosting adversarial training with hypersphere embedding. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. [1](#), [2](#), [3](#), [5](#), [7](#)
- [38] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3498–3505, 2012. [2](#)
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, et al. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn. (ICML)*, pages 8748–8763, 2021. [2](#), [5](#), [8](#)
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis. (IJCV)*, 115(3):211–252, 2015. [2](#), [6](#)
- [41] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, pages 4077–4087, 2017. [7](#)
- [42] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Int. Conf. Learn. Represent. (ICLR)*, 2014. [1](#)
- [43] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, pages 3630–3638, 2016. [6](#)
- [44] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, pages 3630–3638, 2016. [8](#)
- [45] Mengmeng Wang et al. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. [2](#), [4](#)
- [46] Ren Wang, Kaidi Xu, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Chuang Gan, and Meng Wang. On fast adversarial robustness adaptation in model-agnostic meta-learning. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. [3](#), [5](#), [6](#)
- [47] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 69–77, 2016. [2](#)
- [48] Yongqin Xian, Christoph H. Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning - A comprehensive evaluation of the good, the bad and the ugly. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 41(9):2251–2265, 2019. [6](#), [2](#)
- [49] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3485–3492, 2010. [7](#), [2](#)
- [50] Yunlong Yu, Zhong Ji, Yanwei Fu, Jichang Guo, Yanwei Pang, and Zhongfei (Mark) Zhang. Stacked semantics-guided attention model for fine-grained zero-shot learning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, pages 5998–6007, 2018. [2](#)
- [51] Mehmet Kerim Yucel et al. How robust are discriminatively trained zero-shot learning models? *Image Vis. Comput.*, 119: 104392, 2022. [1](#), [3](#), [6](#)
- [52] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically principled trade-off between robustness and accuracy. In *Int. Conf. Mach. Learn. (ICML)*, pages 7472–7482, 2019. [1](#), [2](#), [4](#), [8](#)
- [53] Jiaming Zhang, Qi Yi, and Jitao Sang. Towards adversarial attack on vision-language pre-training models. In *ACM Int. Conf. Multimedia*, pages 5005–5013, 2022. [2](#)
- [54] Manli Zhang, Jianhong Zhang, Zhiwu Lu, Tao Xiang, Mingyu Ding, and Songfang Huang. IEPT: instance-level and episode-level pretext tasks for few-shot learning. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. [6](#)
- [55] Xingxing Zhang, Shupeng Gui, Zhenfeng Zhu, Yao Zhao, and Ji Liu. ATZSL: defensive zero-shot recognition in the presence of adversaries. *arXiv preprint arXiv:1910.10994*, 2019. [1](#), [3](#), [6](#)
- [56] Mo Zhou and Vishal M Patel. Enhancing adversarial robustness for deep metric learning. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 15325–15334, 2022. [2](#)
- [57] Xiaopei Zhu, Xiao Li, Jianmin Li, Zheyao Wang, and Xi-aolin Hu. Fooling thermal infrared pedestrian detectors in real world using small bulbs. In *AAAI*, pages 3616–3624, 2021. [1](#)