

On Generating Transferable Targeted Perturbations

Muzammal Naseer^{*,†}, Salman Khan[†], Munawar Hayat[§], Fahad Shahbaz Khan^{†,◦}, Fatih Porikli[‡]

^{*}Australian National University, Australia, [§]Monash University, Australia, [‡]Qualcomm, USA

[†]Mohamed bin Zayed University of Artificial Intelligence, UAE [◦]Linköping University, Sweden

muzammal.naseer@anu.edu.au, {salman.khan,fahad.khan}@mbzuai.ac.ae, munawar.hayat@monash.edu.au

fatih.porikli@gmail.com

Abstract

While the untargeted black-box transferability of adversarial perturbations has been extensively studied before, changing an unseen model’s decisions to a specific ‘targeted’ class remains a challenging feat. In this paper, we propose a new generative approach for highly transferable targeted perturbations (TTP). We note that the existing methods are less suitable for this task due to their reliance on class-boundary information that changes from one model to another, thus reducing transferability. In contrast, our approach matches the perturbed image ‘distribution’ with that of the target class, leading to high targeted transferability rates. To this end, we propose a new objective function that not only aligns the global distributions of source and target images, but also matches the local neighbourhood structure between the two domains. Based on the proposed objective, we train a generator function that can adaptively synthesize perturbations specific to a given input. Our generative approach is independent of the source or target domain labels, while consistently performs well against state-of-the-art methods on a wide range of attack settings. As an example, we achieve 32.63% target transferability from (an adversarially weak) VGG19_{BN} to (a strong) WideResNet on ImageNet val. set, which is 4× higher than the previous best generative attack and 16× better than instance-specific iterative attack. Code is available at: <https://github.com/Muzammal-Naseer/TTP>.

1. Introduction

We study the challenging problem of *targeted* transferability of adversarial perturbations. In this case, given an input sample from any source category, the goal of the adversary is to change the decision of an *unknown* model to a *specific* target class (e.g., misclassify any painting image to Fire truck, see Fig. 1). This task is significantly more difficult than merely changing the decision to a ran-

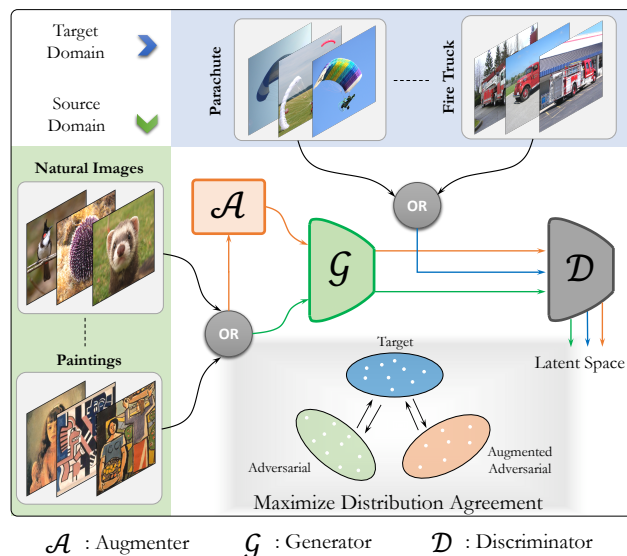


Figure 1: Attack Overview (TTP): Instead of finding perturbations specific to a class-boundary information learned by a model, TTP seeks to match global distribution statistics between the source and the target domains. Specifically, our generator function is trained to maximize agreement between the *perturbed* source distribution, its augmented version and the target distribution in the feature space. Importantly, our attack can function in an unsupervised fashion and does not require source domain to be the same as target (e.g., perturbations can be learned from paintings to transfer on natural images).

dom target class or any similar class (e.g., changing ‘cat’ to ‘aeroplane’ is more difficult than altering the decision to ‘dog’). Target transferability can therefore lead to *goal-driven* adversarial perturbations that provide desired control over the attacked model. However, target transferability remains challenging for the current adversarial attacks [25, 23, 4, 43, 14, 13, 12, 20, 41] that transfer adversarial noise in a *black-box* setting, where architecture and training mechanism of the attacked model remain unknown, and the attack is restricted within a certain perturbation budget.

We observe that modest performance of existing meth-

ods on targeted transferability is due to their reliance on class-boundary information learned by the model which lacks generalizability. For example, iterative *instance-specific* attacks rely on the classification score information to perturb a given sample, thereby ignoring the global *class-specific* information [25, 4, 43, 13]. Such adversarial directions also vary across different models [23], leading to poor target transferability [23, 4]. On the other hand, although universal and generative perturbations are designed to encode global noise patterns [26, 35, 31], they still exploit the class impressions learned by a neural network which alone are not fully representative of the target distribution, thereby achieving only modest black-box fooling rates [36]. Furthermore, they are dependent on the classification information, necessitating a supervised pretrained model for generator’s guidance and therefore cannot directly work with unsupervised features [1, 9]. Another group of techniques exploit intermediate features, but they either find untargeted perturbations by design [28, 37] or are limited in their capacity to transfer targeted perturbations [20, 14, 12, 13].

We introduce a novel generative training framework which maps a given source distribution to a specific target distribution by maximizing the mutual agreement between the two in the latent space of a pretrained discriminator. Our main contributions are:

- **Generative Targeted Transferability:** We propose a novel generative approach to learn transferable targeted adversarial perturbations. Our unique training mechanism allows the generator to explore augmented adversarial space during training which enhances the transferability of adversarial examples during inference (Sec. 3.1).
- **Mutual Distribution Matching:** Our training approach is based on maximizing the mutual agreement between the given source and the target distribution. Therefore, our method can provide targeted guidance to train the generator without the need of classification boundary information. This allows an attacker to learn targeted generative perturbations from the unsupervised features [1, 9] and eliminate the cost of labelled data (Sec. 3.2).
- **Neighbourhood Similarity Matching:** Alongside global distribution matching, we introduce batch-wise neighbourhood similarity matching objective between adversarial and target class samples to maximize the local alignment between the two distributions (Sec. 3.3).

Our extensive experiments on various ImageNet splits and CNN architectures show state-of-the-art targeted transferability against naturally and adversarially trained models, stylized models and input-processing based defenses. The results demonstrate our benefit compared to recent targeted instance-specific as well as other generative methods. Further, our attack demonstrates rapid convergence.

2. Related Work

Iterative Instance-Specific Perturbations: After Szegedy *et al.* [40] highlighted the vulnerability of neural networks, many adversarial attacks have been introduced to study if the adversarial examples are transferable from one model to another, when a target model is unknown. Among these, iterative instance-specific attacks [4, 43, 5] perturb a given sample by iteratively using gradient information. Target transferability of such attacks is very poor [4, 23] (as shown in Sec. 4). Other attacks also use feature space either by maximizing the feature difference [44, 11, 21] or applying attention [42] or avoiding non-linearity while back-propagating gradients [8] or exploiting skip-connections [41]. However, these attacks are mainly designed to enhance non-targeted transferability which is an easier problem. Recently, different instance-specific (transferable) targeted attacks have been proposed including [20] which introduces a triplet loss to push adversarial examples towards the target label while increasing their distance from the original label. Inkawich *et al.* [12, 13] proposed to exploit feature space [14] along with the classifier information [13] to generate target adversaries that are shown to transfer relatively better than other instance-specific attacks. These attacks [12, 13] have the following limitations. **a)** They need access to a *labeled* dataset *e.g.*, ImageNet [38] in order to train one-vs-all binary classifiers for attacked target classes. **b)** They need to identify best performing single layer [12] or a combination of layers [13] which adds further complexity to attack optimization. **c)** Finally, the attack performance degrades significantly with quality of features, *e.g.*, it struggles to transfer target perturbations from VGG models [12].

Universal Perturbation: In contrast to instance-specific perturbations, [26] learns a single universal noise pattern which is representative of the entire data distribution and can fool a model on majority of samples. Li *et al.* [22] introduce gradient transformation module to find smooth universal patterns while [28] shows that such patterns can be found without any training data. Although universal perturbations [26, 27, 28, 22] based attacks are efficient (the attacker just needs to add the noise to any given sample at inference), they are limited in their capacity to yield transferable adversaries which can generalize across different data distributions and models [35, 31].

Generative Perturbations: Generative adversarial perturbations perform better than directly optimizing universal noise [37, 35, 31]. Poursaeed *et al.* [35] proposed the first generative approach to adapt perturbations to an input sample. Naseer *et al.* [31] improved this framework with relativistic training objective which also allows cross-domain transferability. Our method belongs to the generative category and can adapt to an input sample with a single forward pass. Unlike [37, 35, 31], we seek to fool the model by matching *distributions* of source and targets with distribu-

tion matching and neighbourhood similarity criteria. Our proposed framework does not require labeled source or target data and can extract target perturbations from a discriminator model trained in an unsupervised manner while previous generative methods are dependent on class-boundary information learned by the model. Further, our method converges faster (Sec. 4) and provides improved targeted transferability owing to its novel loss and training mechanism.

3. Generating Targeted Adversaries

Our goal is to craft adversarial perturbations δ that can fool a model to misclassify any given input to a specific target class t . We assume access to source and target domain data represented by P and Q , from which the source and target class samples are obtained i.e., $\mathbf{x}_s \sim P$, $\mathbf{x}_t \sim Q$. The source and target domains are likely to be non-aligned i.e., $P \neq Q$, making it challenging to achieve targeted transferability of adversarial perturbations. We also consider a perturbed source data P' that comprises of adversarially manipulated samples $\mathbf{x}'_s \sim P'$ where $\mathbf{x}'_s = \mathbf{x}_s + \delta$. \mathbf{x}_s , \mathbf{x}'_s and \mathbf{x}_t represent source, adversarial and target domain samples while $\mathcal{D}_\psi(\mathbf{x}_s)$, $\mathcal{D}_\psi(\mathbf{x}'_s)$ and $\mathcal{D}_\psi(\mathbf{x}_t)$ are their corresponding latent distributions.

3.1. Generative Model

We propose a generative approach to perturb the source domain samples \mathbf{x}_s to a specified target class. The framework (see Fig. 2) consists of a generator \mathcal{G}_θ and a discriminator \mathcal{D}_ψ parameterized by θ and ψ , respectively. The generator function \mathcal{G}_θ learns a mapping from the source images to the target category such that the input images are minimally changed i.e., adversarial noise δ is strictly constrained under a norm distance $l_\infty \leq \epsilon$. This is ensured by projecting the unbounded adversaries from \mathcal{G}_θ within fixed norm distance of \mathbf{x}_s using a differentiable clipping operation,

$$\mathbf{x}'_s = \text{clip}(\min(\mathbf{x}_s + \epsilon, \max(\mathcal{W} * \mathcal{G}_\theta(\mathbf{x}_s), \mathbf{x}_s - \epsilon))), \quad (1)$$

where, \mathcal{W} is a smoothing operator with fixed weights that reduces high frequencies without violating the l_∞ distance constraint. The smooth projection in Eq. 1 (denoted by \mathcal{P} in Fig. 2) not only tightly bounds generator’s output within l_∞ norm but also encourages avoiding redundant high frequencies [34] during the optimization process. This allows the generator to converge to a more meaningful solution.

The existing generative designs for adversarial attacks [35, 31] leverage the decision space of the discriminator to craft perturbations. In such cases, the *class-boundary* information learned by the discriminator is used to fool DNN models (e.g. for ImageNet, discriminator is pretrained on 1k classes). This dependence is problematic since an attacker must have access to a discriminator trained on large-scale labeled dataset [3]. Attacker then tries to learn target

Algorithm 1 Generating TTP

Require: Source data \mathcal{X}_s , Target data \mathcal{X}_t , pretrained discriminator \mathcal{D}_ψ , perturbation budget ϵ , loss criteria \mathcal{L}_G .

Ensure: Randomly initialize the generator, \mathcal{G}_θ

- 1: **repeat**
 - 2: Randomly sample mini-batches $\mathbf{x}_s \sim \mathcal{X}_s$ and $\mathbf{x}_t \sim \mathcal{X}_t$
 - 3: Create augmented copy of the source mini-batch $\tilde{\mathbf{x}}_s$.
 - 4: Forward-pass \mathbf{x}_s and $\tilde{\mathbf{x}}_s$ through the generator and generate unbounded adversaries; $\mathbf{x}'_s, \tilde{\mathbf{x}}'_s$.
 - 5: Bound the adversaries using Eq. 1 such that:

$$\|\mathbf{x}'_s - \mathbf{x}_s\|_\infty \leq \epsilon \quad \text{and} \quad \|\tilde{\mathbf{x}}'_s - \tilde{\mathbf{x}}_s\|_\infty \leq \epsilon$$
 - 6: Forward pass $\mathbf{x}'_s, \tilde{\mathbf{x}}'_s$ and \mathbf{x}_t through \mathcal{D}_ψ .
 - 7: Compute the matching losses; \mathcal{L} , \mathcal{L}^{aug} and \mathcal{L}^{sim} using Eq. 3, 4 and 8, respectively.
 - 8: Compute the generator loss given in Eq. 9.
 - 9: Backward pass and update \mathcal{G}_θ
 - 10: **until** \mathcal{G}_θ converges.
-

class impressions using either cross-entropy (CE) [35] or relativistic CE [31]. Thus, the generated perturbations are directly dependent on the quality of the discriminator’s classification space. Furthermore, the generated adversaries are dependent on the input instance-specific features and do not model the global properties of the target distribution, resulting in only limited transferability.

To address above limitations, our generative design models the target distribution Q and pushes the perturbed source distribution P' closer to Q using the latent space of \mathcal{D}_ψ ,

$$\|\delta\|_\infty \leq \epsilon, \quad s.t., \quad \mathcal{D}_\psi(\mathbf{x}'_s) \approx \mathcal{D}_\psi(\mathbf{x}_t). \quad (2)$$

This global objective provides two crucial benefits. *First*, reducing mismatch between perturbed and target distributions provides an improved guidance to the generator. The resulting perturbations well align the input samples with the target distribution, leading to transferable adversaries. *Second*, the distributions alignment task makes us independent of the \mathcal{D}_ψ ’s classification information. In turn, our approach can function equally well with a discriminator trained in a self-supervised manner on unlabelled data [1, 9]. In our case, we simply align the feature distributions from \mathcal{D}_ψ to match P' and Q . Thus, for a given sample \mathbf{x} , n -dimensional features are obtained i.e., $\mathcal{D}_\psi(\mathbf{x}) \in \mathbb{R}^n$. If \mathcal{D}_ψ is trained in a supervised manner on ImageNet then $n = 1000$, and if \mathcal{D}_ψ is trained in an unsupervised fashion then n is equal to the output feature dimension.

3.2. Distribution Matching

We measure the mutual agreement between P' and Q using Kullback Leibler (KL) divergence defined on discrimi-

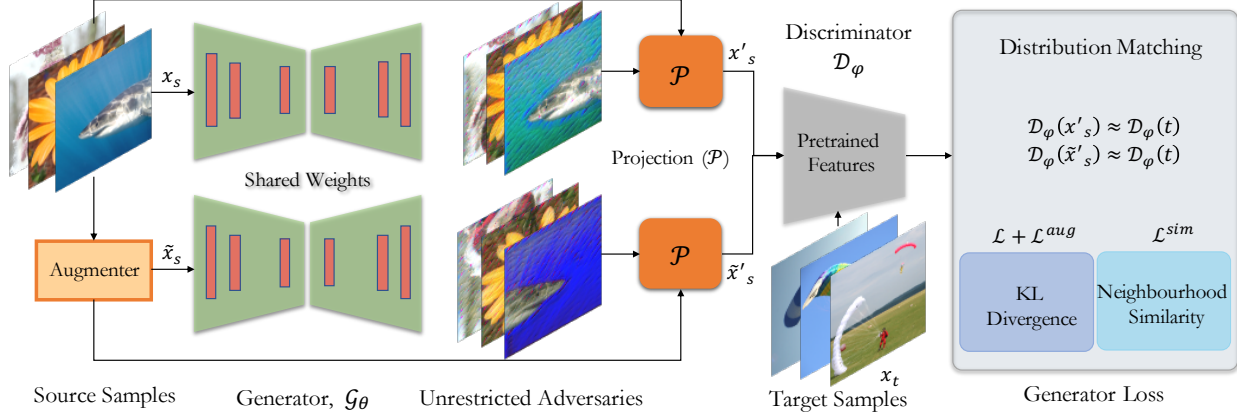


Figure 2: *Targeted Transferable Perturbations*: During training, TTP matches adversarial and augmented adversarial samples to a target domain within discriminator’s latent space for improved transferability. The adversarial samples corresponding to original and augmented images are bounded (via projection) around their source samples to explore adversarial space around natural as well as augmented samples.

nator features $\mathcal{D}_\psi(\mathbf{x}'_s)$ and $\mathcal{D}_\psi(\mathbf{x}_t)$,

$$D_{KL}(P' \| Q) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^n \sigma(\mathcal{D}_\psi(\mathbf{x}'_s{}^i))_j \log \frac{\sigma(\mathcal{D}_\psi(\mathbf{x}'_s{}^i))_j}{\sigma(\mathcal{D}_\psi(\mathbf{x}_t^i))_j},$$

where N represents the number of samples, n is the discriminator’s output dimension, and σ denotes the softmax operation. In simple terms, KL divergence measures the difference between two distributions in terms of the average surprise in experiencing \mathbf{x}_t when we expected to see \mathbf{x}'_s . Since KL divergence is asymmetric *i.e.* $D_{KL}(P' \| Q) \neq D_{KL}(Q \| P')$, and not a valid distance measure, we define our loss function for distribution matching [19] as follows:

$$\mathcal{L} = D_{KL}(P' \| Q) + D_{KL}(Q \| P'). \quad (3)$$

As a regularization measure, we add augmented versions of the source domain samples during distribution matching. This enables the generator to focus specifically on adding target class-specific patterns that are robust to input transformations. To this end, we randomly apply rotation, crop resize, horizontal flip, color jittering or gray scale transformation to create augmented samples $\tilde{\mathbf{x}}_s$ from the original \mathbf{x}_s . The $\tilde{\mathbf{x}}_s \sim \tilde{P}$ are passed through the \mathcal{G}_θ and the perturbed augmented samples $\tilde{\mathbf{x}}'_s \sim \tilde{P}'$ are projected using Eq. 1 to stay close to the augmented samples *i.e.*, $\|\tilde{\mathbf{x}}'_s - \tilde{\mathbf{x}}_s\|_\infty \leq \epsilon$. No augmentation is applied to the target domain samples. We then pass $\tilde{\mathbf{x}}'_s$ through the discriminator and compute the mutual agreement between $\mathcal{D}_\psi(\tilde{\mathbf{x}}'_s)$ and $\mathcal{D}_\psi(\mathbf{x}_t)$ as follows:

$$\mathcal{L}^{aug} = D_{KL}(\tilde{P}' \| Q) + D_{KL}(Q \| \tilde{P}'). \quad (4)$$

The impact of data augmentations and their effectiveness for our proposed targeted attack is studied in Sec. 4.

3.3. Neighbourhood Similarity Matching

The above objective promotes alignment between the distributions but does not consider the local structure *e.g.*, the relationship between a sample and its augmented versions. For a faithful alignment between perturbed source samples and the target class samples, we propose to also match the *neighbourhood similarity distributions* between the two domains. Specifically, consider a batch of target domain samples $\{\mathbf{x}_t^i\}_{i=1}^N$ and a batch of perturbed source domain samples $\{\mathbf{x}'_s{}^i\}_{i=1}^N$. For the case of \mathbf{x}'_s , in a given training batch, we compute a similarity matrix \mathcal{S}^s whose elements encode the cosine similarity between the original sample and its augmented version $\tilde{\mathbf{x}}'_s$, *i.e.*,

$$\mathcal{S}_{i,j}^s = \frac{\mathcal{D}_\psi(\mathbf{x}'_s{}^i) \cdot \mathcal{D}_\psi(\tilde{\mathbf{x}}'_s{}^j)}{\|\mathcal{D}_\psi(\mathbf{x}'_s{}^i)\| \|\mathcal{D}_\psi(\tilde{\mathbf{x}}'_s{}^j)\|}. \quad (5)$$

In contrast, for the case of \mathbf{x}_t , we compute similarity between only the original target samples (no augmentations) as we need to model the local neighbourhood connectivity in the target domain. This choice is impractical for the source domain case where many categories co-exist, while for the target distribution, we assume a single category. Thus the target similarity matrix \mathcal{S}^t is computed as,

$$\mathcal{S}_{i,j}^t = \frac{\mathcal{D}_\psi(\mathbf{x}_t^i) \cdot \mathcal{D}_\psi(\mathbf{x}_t^j)}{\|\mathcal{D}_\psi(\mathbf{x}_t^i)\| \|\mathcal{D}_\psi(\mathbf{x}_t^j)\|}. \quad (6)$$

The resulting similarity matrices are normalized along the row dimension with softmax to obtain probability estimates,

$$\bar{\mathcal{S}}_{i,j} = \frac{\exp(\mathcal{S}_{i,j})}{\sum_k \exp(\mathcal{S}_{i,k})}, \text{ where, } \mathcal{S} \in \{\mathcal{S}^s, \mathcal{S}^t\}. \quad (7)$$

Here, each term shows the probability with which the two sample pairs are related to each other. Given $\bar{\mathcal{S}}^s$ and $\bar{\mathcal{S}}^t$, we

compute the KL divergence to enforce a loss term that seeks to match the local neighbourhood patterns between source and target domains,

$$\mathcal{L}^{sim} = \sum_{i,j} \bar{S}_{i,j}^t \log \frac{\bar{S}_{i,j}^t}{\bar{S}_{i,j}^s} + \sum_{i,j} \bar{S}_{i,j}^s \log \frac{\bar{S}_{i,j}^s}{\bar{S}_{i,j}^t}. \quad (8)$$

3.4. Overall loss function

Finally, the generator parameters are updated by minimizing the following loss (Algorithm 1):

$$\mathcal{L}_G = \mathcal{L} + \mathcal{L}^{aug} + \mathcal{L}^{sim}. \quad (9)$$

This loss encourages the generator to perturb source samples that not only match the global characteristics of the target distribution ($\mathcal{L} + \mathcal{L}^{aug}$), but also the local information based on neighbourhood connectivity (\mathcal{L}^{sim}).

4. Experiments

Our generator \mathcal{G}_θ is based on ResNet architecture [16], and outputs an adversarial sample with the same size as of input (Fig. 3). This generator architecture is the same as in the baseline generative attacks [35, 31]. Our discriminator \mathcal{D}_ψ is pre-trained in a supervised or self-supervised manner. For training \mathcal{G}_θ , we freeze \mathcal{D}_ψ . We use Adam optimizer [18] with a learning rate of 10^{-4} ($\beta_1 = .5, \beta_2 = .999$) for 20 epochs. For source domain data, we use 50k random images from ImageNet train set. Our method is not sensitive to the choice of source samples since it can learn transferable perturbations even from other domains *e.g.* Paintings. Similar to other generative methods [35, 31], we fix source data. For target domain data, we use 1300 images for each target collected from ImageNet training set (without their original labels). We used default settings or implementations as provided by the authors of baseline attacks. Similarly, we used open-sourced (pretrained) stylized [7], adversarial [39] and purifier (NRP) [29] models to evaluate robustness.

4.1. Evaluation Settings

We perform inference on ImageNet validation set (50k samples). *No augmentations are applied at inference time.* The perturbation budget is tightly bounded and clearly mentioned in each experiment following the standard practices $l_\infty \leq 16$ [4, 14, 13] and $l_\infty \leq 32$ [31, 22]. We perturb all the ImageNet val. samples (except the target samples) to the pre-defined target class. We repeat this process for all the given targets and report Top-1 (%) accuracy averaged across all targets. We compare our method under two main settings (10-Targets and 100-Targets), as described below.

10-Targets: We further consider two settings. **(a)** *10-Targets (subset-source)* which is consistent with [12] and has a subset of source classes at inference. **(b)** *10-Targets (all-source)* which is a more challenging large-scale setting

Src.	Attack	Naturally Trained (IN) Models				
		VGG19 _{BN}	Dense121	ResNet50	ResNet152	WRN-50-2
VGG19 _{BN}	PGD [25]	95.67*	0.31	0.30	0.20	0.25
	MIM [4]	99.91*	0.92	0.68	0.36	0.47
	DIM [43]	99.38*	3.10	2.08	1.02	1.29
	DIM-TI [5]	89.71*	1.08	0.66	0.42	0.45
	Po-TRIP [20]	99.40*	4.61	3.21	1.78	2.01
	GAP [35]	98.23*	16.19	15.83	5.89	7.78
	CDA [31]	98.30*	16.26	16.22	5.73	8.35
	Ours-P	97.38*	45.53	42.90	26.72	31.00
	Ours	98.54*	45.77	45.87	27.18	32.63
Ven _s	Ours	97.34*	71.41	71.68	50.78	48.03
Dense121	PGD [25]	1.28	97.40*	1.78	1.01	1.37
	MIM [4]	1.85	99.90*	2.71	1.68	1.88
	DIM [43]	7.31	98.81*	9.06	5.78	6.29
	DIM-TI [5]	0.91	88.59*	1.18	0.77	0.86
	Po-TRIP [20]	8.10	99.00*	11.21	7.83	8.50
	GAP [35]	39.01	97.30*	47.85	39.25	34.79
	CDA [31]	42.77	97.22*	54.28	44.11	46.01
	Ours-P	57.91	97.41*	71.35	55.57	53.45
	Ours	58.90	97.61*	68.72	57.11	56.80
Den _s	Ours	76.96	96.25*	88.81	83.48	81.85
ResNet50	PGD [25]	0.92	1.38	93.74*	1.86	1.89
	MIM [4]	1.58	3.37	98.76*	3.39	3.17
	DIM [43]	9.14	15.47	99.01*	12.45	12.61
	DIM-TI [5]	0.79	2.12	88.91*	1.47	1.45
	Po-TRIP [20]	12.01	19.43	99.22*	14.41	15.10
	GAP [35]	58.47	71.72	96.81*	64.89	61.82
	CDA [31]	64.58	73.57	96.30*	70.30	69.27
	Ours-P	73.09	84.76	96.63*	76.27	75.92
	Ours	78.15	81.64	97.02*	80.56	78.25
Res _s	Ours	90.43	94.39	96.67*	95.48*	92.63

Table 1: **Target Transferability:** {10-Targets (all-source)} Top-1 target accuracy (%) averaged across 10 targets with 49.95K ImageNet val. samples. Perturbation budget: $l_\infty \leq 16$. Our method outperforms previous instance-specific as well as generative approaches by a large margin. '*' indicates white-box attack. **Ours-P** represents TTP trained on Paintings.

as source images can come from all the ImageNet classes except the target class. For consistency and direct comparison, the ten target classes are same as in [12].

–*10-Targets (subset-source):* Following [12], for each target class, 450 source samples belonging to remaining 9 classes (except target class) become inputs to \mathcal{G}_θ to be transferred to the selected target.

–*10-Targets (all-source):* For each target class, samples of all 999 source classes (except the target class) in ImageNet val. set are considered i.e., for each target class, 49,950 samples of 999 classes become inputs to \mathcal{G}_θ .

100-Targets (all-source): We divide ImageNet 1k classes into 100 mutually exclusive sets. Each set contains 10 classes. We randomly sample 1 target from each set to create 100 targets (see Appendix E for more details). Generators are trained against these targets and evaluated on ImageNet val. set in 100-Targets (all-source) setting with the same protocol as described for 10-Targets (all-source).

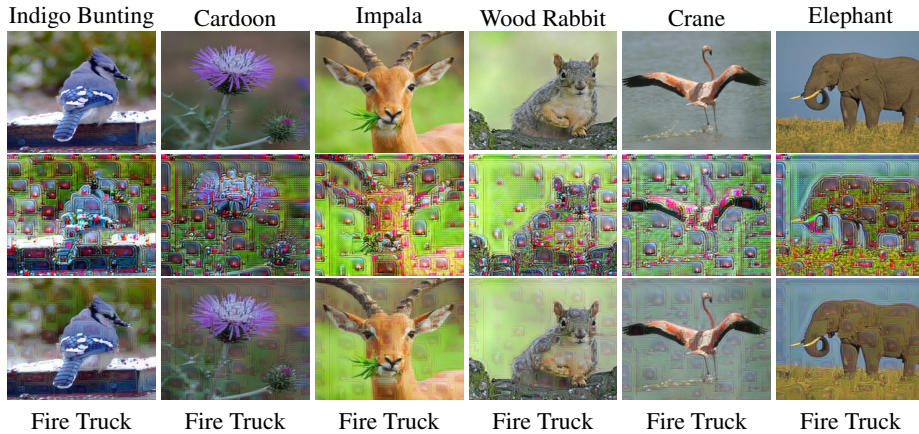


Figure 3: Targeted adversaries produced by a TTP generator learned to maximize the agreement with 'Fire Truck' distribution against Dense121 ImageNet model. 1st and 2nd rows show clean images and unrestricted outputs of the adversarial generator, respectively. 3rd row shows adversaries after valid projection. See Appendix F for more qualitative examples including comparisons between targeted patterns learned by TTP from different source models of a certain family of networks.

Src.	Attack	VGG19 _{BN}	Dense121	ResNet50
VGG19 _{BN}	AA [14]	–	0.8	0.6
	FDA-fd [12]	–	3.0	2.1
	FDA ^N [13]	–	6.0	5.4
	CDA [31]	–	17.82	17.09
	Ours-P	–	48.56	44.47
	Ours	–	48.29	47.07
Dense121	AA [14]	0.0	–	0.0
	FDA-fd [12]	34.0	–	34.0
	FDA ^N [13]	42.0	–	48.3
	CDA [31]	44.84	–	53.73
	Ours-P	59.81	–	71.32
	Ours	61.75	–	69.60
ResNet50	AA [14]	1.1	2.0	–
	FDA-fd [12]	16.0	21.0	–
	FDA ^N [13]	32.1	48.3	–
	CDA [31]	68.55	75.68	–
	Ours-P	75.18	85.71	–
	Ours	79.04	84.42	–

Table 2: **Target Transferability:**{10-Targets (sub-source)} Top-1 accuracy (%) across 10 targets. Our method shows significant improvements in trasfering target perturbations compared to generative as well as feature based instance-specific method [12, 13]. Perturbation budget: $l_\infty \leq 16$. Only black-box attack results are shown. **Ours-P** represents TTP trained on Paintings.

4.2. Attack Protocols and Results

We evaluate black-box target transferability in the following scenarios. **(a) Unknown Target Model:** Attacker has access to a pretrained discriminator trained on labeled data but has no knowledge about the architecture of the target model. **(b) Unknown Decision Space:** Attacker has access to the pre-trained discriminator trained on unlabeled data in an unsupervised manner but does not know about the architecture and the class-boundary information learned by the target model. **(c) Unknown Defense:** Attacker is unaware of the type of defense deployed at the target model, or if any defense is applied at all, e.g., the defense can be an input processing approach or a robust training mechanism such as adversarial training.

4.2.1 Unknown Target Model

Natural Training: We evaluate naturally trained ImageNet models and show strong empirical results in Tables 1, 2 & 3 demonstrating that generative methods are far superior than sample-specific targeted attacks based on boundary information [20, 4, 43] or feature exploitation [14, 12, 13]. Our approach has significantly higher target transferability rates than previous generative methods [31, 35]. To highlight an example from Table 2, our method achieves 47.07% transferability from VGG19_{BN} to ResNet50 which is 175% and 771% better than the previous best generative [31] and sample-specific [13] target attacks, respectively.

Ensemble Effect: We also train generators with our algorithm on the ensembles of same-family discriminators. Specifically, we define the following ensembles: $V_{ens} : \text{VGG}\{11, 13, 16, 19\}_{BN}$, $R_{ens} : \text{ResNet}\{18, 50, 101, 152\}$, and $D_{ens} : \text{DenseNet}\{121, 161, 169, 201\}$.

The purpose of such ensembles is to understand if the combination of weak individual models from the same family can provide strong learning for the target distributions. From Table 1, we observe that modeling target distribution from an ensemble provides significantly better transferability than any individual discriminator (see Appendix A for more analysis). This signifies that an attacker can use multiple variants of the same network to boost the attack.

Target Transferability and Model Disparity: We note that within a specific family, transferring targeted perturbations from a smaller model to a larger one (e.g. ResNet18 \rightarrow ResNet152 or VGG11_{BN} \rightarrow VGG19) is difficult as we increase the size discrepancy. Interestingly, this trend remains the same even from larger to smaller models i.e., the attack strength will increase with the disparity between models rather than only depending upon the strength of target model. For example, target transferability ResNet152 \rightarrow ResNet50 is higher than ResNet152 \rightarrow ResNet18 even though ResNet18 is weaker than ResNet50 (Fig. 4). Similar behaviour can be observed within cross-family models i.e.,

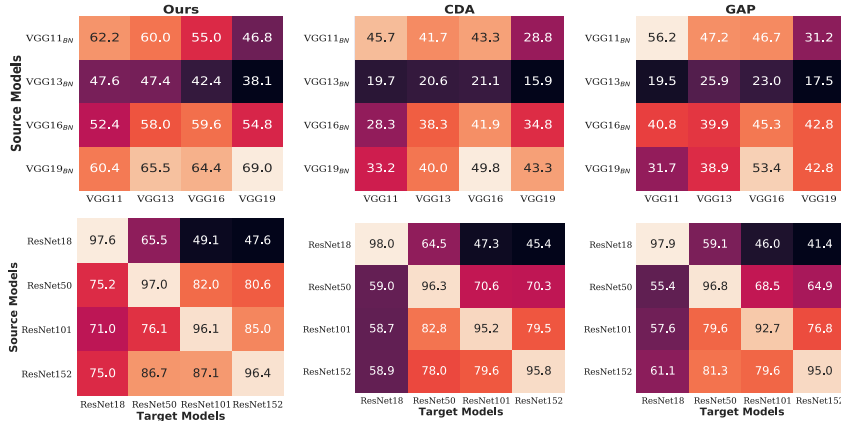


Figure 4: **Within Family Target Transferability:** {10-Targets (all source) settings} These results indicate that our approach boosts target transferability within different models of the same family with or without batch-norm and favorably beats the previous generative approaches (GAP [35], CDA [31]) by a large margin. Each value is averaged across 10 targets (Sec. 4) with 49.95k ImageNet val. samples for each target. Perturbation budget is set to $l_\infty = 16$.




Figure 5: **Target Transferability of Unsupervised Features:** {10-Targets (all-source) settings}. Our approach when applied to unsupervised features, MoCo [9], surpasses GAP [35] and CDA [31] that are dependent on classification layer by design. Perturbation budget is $l_\infty = 16$.

target transferability from ResNet50 to Dense121 and vice versa is higher than VGG19_{BN} as both models share skip connections (Table 1). See Appendix B for vulnerability of models with and without batch-norm [15].

4.2.2 Unknown Decision Space

Here we investigate the question, “Can unsupervised features provide targeted adversarial perturbations?” A unique property of our proposed approach is that it can be applied to feature space without any class boundary information to achieve target adversarial direction. This allows an attacker to benefit from recently proposed unsupervised feature learning methods [9, 1]. Rather than using a discriminator trained on large scale labelled data, attack can be learned and launched from features of a discriminator trained purely in an unsupervised fashion. Therefore, our attack can eliminate the cost of label annotations. Results in Fig. 5 demonstrate that our method learned from unsupervised features, MoCo [9], not only provides target transferability but surpasses the previous generative methods which are dependent on the discriminator trained on labelled data.

4.2.3 Unknown Defense Mechanisms

Input Processing as a Defense: We evaluate robustness of different input processing based adversarial defense methods in Fig. 6. We consider the following four representative defenses: a) JPEG with compression quality set to 50% [2], b) DNN-Oriented JPEG compression [24], c) Median Blur with window size set to 5×5 [30], and d) Neural representa-

Attack	VGG19 _{BN}	Dense121	ResNet-152	WRN-50-2	SIN [7]
GAP [35]	47.87	58.10	54.72	49.65	7.1
CDA [31]	53.41	60.34	57.67	51.23	7.6
Ours	69.55	77.48	75.74	74.61	31.0

Table 3: **Target Transferability:** {100-Targets (all-source)} Top-1 target accuracy (%) averaged across 100 targets with 49.95K ImageNet val. samples per target. Generators are trained against ResNet50. Perturbation budget is $l_\infty \leq 16$.

ϵ	Attack	Augmix [10]	Stylized [7]		Adversarial [39]			
			SIN-IN	SIN	l_∞		l_2	
					$\epsilon=.5$	$\epsilon=1$	$\epsilon=.1$	$\epsilon=.5$
16	GAP [35]	51.57	76.92	12.96	1.88	0.34	23.41	0.92
	CDA [31]	59.79	75.93	9.21	2.10	0.39	23.89	1.18
	Ours	73.09	87.40	30.17	4.63	0.56	45.40	1.99
	Ours _{ens}	88.79	92.96	57.75	14.23	1.24	74.95	7.62
32	GAP [35]	54.86	81.15	28.07	26.32	6.36	59.04	16.53
	CDA [31]	63.18	76.81	19.65	27.60	6.74	57.54	16.07
	Ours	78.66	91.27	41.52	46.82	16.35	75.97	30.94
	Ours _{ens}	89.96	94.15	70.70	70.22	34.21	90.42	58.25

Table 4: **Target Transferability:** {10-Targets (all source) settings} Top-1 (%) target accuracy. Generators are trained against naturally trained ResNet50 or ResNet ensemble. Perturbation are then transferred to ResNet50 trained using different methods including Augmix [10], Stylized [7] or adversarial [39].

tion purifier (NRP) [29] which is a state-of-the-art defense. Generators are trained against naturally trained ResNet50 and target perturbations are then transferred to VGG19_{BN} and Dense121 which are protected by the input processing defenses. We observe (Fig. 6) that JPEG is the least effective

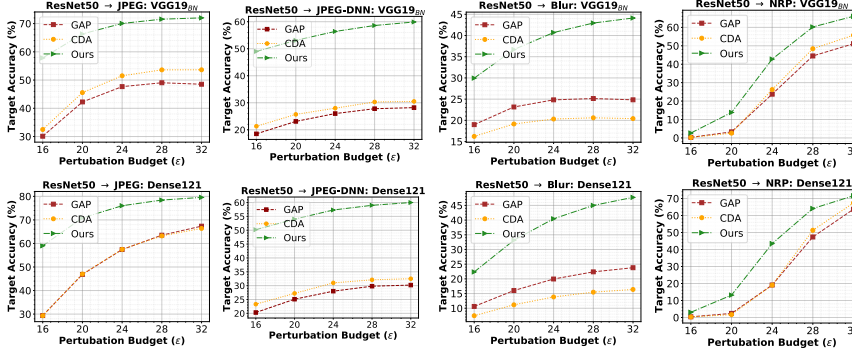


Figure 6: **Target Transferability against Input Processing Defenses:** {10-Targets (all-source) settings} Input processing including NRP [29] are broken under targeted black-box attacks. Our method outperforms GAP [35] and CDA [31] on all the considered defenses including JPEG, JPEG-DNN [24], Median Blur and NRP [29]. Each point is an averaged across 10 targets (Sec. 4) with 49.95k ImageNet val. samples for each target. Generators are trained against ResNet50.

tive method against target attacks while JPEG-DNN [24] performs relatively better than JPEG. Compared to JPEG, JPEG-DNN and Median blur, NRP shows better resistance to target attacks at $l_\infty \leq 16$ but quickly breaks as perturbation is increased. Median blur shows more resistance than JPEG, JPEG-DNN and NRP at higher perturbation rates ($l_\infty \leq 32$)¹. Success rate of our method is much better than previous generative attacks [35, 31] even when the target model and the input processing remain unknown (Fig. 6).

Robust Training Mechanism: Here we study the transferability of our approach against various robust training methods (augmented vs. stylized vs. adversarial) based defense strategies. Augmentation based training can make the model robust to natural corruptions [10] while training on stylized ImageNet [7] improves shape bias and training on adversarial examples can improve robustness against adversarial attacks at the cost of computation, clean accuracy, and generalization to global changes [6]. We evaluate the vulnerability of these training methods in Table 4. Generators are trained against naturally trained ResNet50 or ResNet ensemble and adversarial perturbations are then transferred to ResNet50 trained using Augmix [10], Stylized ImageNet (SIN) [7], mixture of Stylized and natural ImageNet (SIN-IN) and adversarial examples [39]. Target transferability can easily be achieved against models trained on mixture (SIN-IN), however, the model trained on stylized images (SIN) shows higher resistance but remains vulnerable as our target attack (ensemble) achieves $\approx 71\%$ success at perturbation of $l_\infty = 32$ (Table 4). Adversarially trained models using Madry’s method [25] are more robust to target attacks.

4.3. Ablative Analysis

In order to understand the effect of each component of our approach, we present an ablative study in Fig. 7. Target perturbations are transferred from ResNet50 to VGG16 (SIN) trained on stylized ImageNet which is a much harder task than transferring to naturally trained VGG16. We observe that training TTP on only distribution matching loss (Eq. 3) increases the transferability by more than 100% in




Figure 7: **Ablation:** We dissect effect of each component of our method including novel losses, augmentation, smooth projection and epochs. Results are presented with 10-Target (all source) settings. Perturbation budget is set to $l_\infty = 16$.

comparison to GAP [35] (with cross-entropy) or CDA [31] (with relativistic cross-entropy). Adding smoothing operator \mathcal{W} enhances the efficiency of TTP. \mathcal{W} is a differentiable Gaussian kernel with size 3×3 . We then noticed a significant jump in transferability when augmentations are introduced and TTP is trained using both distribution matching losses (Eq. 3 & 4) which is further complemented by neighbor similarity loss (Eq. 8). Our generator trained for only one epoch outperforms GAP and CDA trained for 20 epochs (Fig. 7) which highlights our rapid convergence rate.

5. Conclusion

We proposed a new generative approach that can learn to model transferable target-specific perturbations. Given an image from any source class, our approach can synthesize perturbations that lead to its misclassification on a variety of black-box target models. The core of our approach is an instance-adaptive generator function that is learned using a novel loss formulation. Our loss focuses on matching the distribution-level statistics of perturbed source and target samples. By its design, our approach can work with both supervised and unsupervised representations. We demonstrate impressive transferability rates across a range of attack settings compared to state-of-the-art. Our results advocate for the use of global loss functions defined over distributions to craft highly transferable adversarial patterns. In future work, we plan to extend proposed method to other model families such as vision transformers [33, 32, 17].

¹Blur defense causes large drop in clean accuracy (see Appendix D).

References

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. **2, 3, 7**
- [2] Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Siwei Li, Li Chen, Michael E Kounavis, and Duen Horng Chau. Shield: Fast, practical defense and vaccination for deep learning using jpeg compression. 2018. **7**
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. **3**
- [4] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. **1, 2, 5, 6, 13**
- [5] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019. **2, 5**
- [6] Nic Ford, Justin Gilmer, Nicolas Carlini, and Dogus Cubuk. Adversarial examples are a natural consequence of test error in noise. *arXiv preprint arXiv:1901.10513*, 2019. **8**
- [7] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. **5, 7, 8, 11, 12, 13**
- [8] Yiwen Guo, Qizhang Li, and Hao Chen. Backpropagating linearly improves transferability of adversarial examples. *arXiv preprint arXiv:2012.03528*, 2020. **2, 11, 13**
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. **2, 3, 7**
- [10] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. **7, 8, 13**
- [11] Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4733–4742, 2019. **2**
- [12] Nathan Inkawhich, Kevin Liang, Lawrence Carin, and Yiran Chen. Transferable perturbations of deep feature distributions. In *International Conference on Learning Representations*, 2020. **1, 2, 5, 6, 13**
- [13] Nathan Inkawhich, Kevin J Liang, Binghui Wang, Matthew Inkawhich, Lawrence Carin, and Yiran Chen. Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. *arXiv preprint arXiv:2004.14861*, 2020. **1, 2, 5, 6, 13**
- [14] Nathan Inkawhich, Wei Wen, Hai Helen Li, and Yiran Chen. Feature space perturbations yield more transferable adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7066–7074, 2019. **1, 2, 5, 6**
- [15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. **7, 11, 12**
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. **5**
- [17] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021. **8**
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. **5**
- [19] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951. **4**
- [20] Maosen Li, Cheng Deng, Tengjiao Li, Junchi Yan, Xinbo Gao, and Heng Huang. Towards transferable targeted attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 641–649, 2020. **1, 2, 5, 6, 13**
- [21] Qizhang Li, Yiwen Guo, and Hao Chen. Yet another intermediate-level attack. In *European Conference on Computer Vision*, pages 241–257. Springer, 2020. **2**
- [22] Yingwei Li, Song Bai, Cihang Xie, Zhenyu Liao, Xiaohui Shen, and Alan L Yuille. Regional homogeneity: Towards learning transferable universal adversarial perturbations against defenses. *arXiv preprint arXiv:1904.00979*, 2019. **2, 5**
- [23] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016. **1, 2**
- [24] Zihao Liu, Qi Liu, Tao Liu, Yanzhi Wang, and Wujie Wen. Feature distillation: Dnn-oriented jpeg compression against adversarial examples. *arXiv preprint arXiv:1803.05787*, 2018. **7, 8**
- [25] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. **1, 2, 5, 8, 13**
- [26] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1765–1773, 2017. **2**
- [27] Konda Reddy Mopuri, Aditya Ganeshan, and R Venkatesh Babu. Generalizable data-free objective for crafting universal adversarial perturbations. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2452–2465, 2018. **2**

- [28] Konda Reddy Mopuri, Utsav Garg, and R Venkatesh Babu. Fast feature fool: A data independent approach to universal adversarial perturbations. *arXiv preprint arXiv:1707.05572*, 2017. 2
- [29] Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. A self-supervised approach for adversarial robustness. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 5, 7, 8, 13
- [30] Muzammal Naseer, Salman Khan, and Fatih Porikli. Local gradients smoothing: Defense against localized adversarial attacks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1300–1307. IEEE, 2019. 7, 13
- [31] Muzammal Naseer, Salman H Khan, Harris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-domain transferability of adversarial perturbations. *Advances in Neural Information Processing Systems*, 2019. 2, 3, 5, 6, 7, 8
- [32] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *arXiv preprint arXiv:2105.10497*, 2021. 8
- [33] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169*, 2021. 8
- [34] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2017. <https://distill.pub/2017/feature-visualization>. 3
- [35] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4422–4431, 2018. 2, 3, 5, 6, 7, 8
- [36] Konda Reddy Mopuri, Phani Krishna Uppala, and R Venkatesh Babu. Ask, acquire, and attack: Data-free uap generation using class impressions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018. 2
- [37] Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, and R Venkatesh Babu. Nag: Network for adversary generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 742–751, 2018. 2
- [38] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 2
- [39] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *ArXiv preprint arXiv:2007.08489*, 2020. 5, 7, 8, 13
- [40] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 2
- [41] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections matter: On the transferability of adversarial examples generated with resnets. In *ICLR*, 2020. 1, 2, 11, 13
- [42] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R Lyu, and Yu-Wing Tai. Boosting the transferability of adversarial samples via attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1161–1170, 2020. 2
- [43] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan Yuille. Improving transferability of adversarial examples with input diversity. In *Computer Vision and Pattern Recognition*. IEEE, 2019. 1, 2, 5, 6, 13
- [44] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xiangqi Huang, Xiang Gan, and Yong Yang. Transferable adversarial perturbations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–467, 2018. 2