

Towards Transferable Unrestricted Adversarial Examples with Minimum Changes

Fangcheng Liu
Peking University
equation@stu.pku.edu.cn

Chao Zhang
Peking University
c.zhang@pku.edu.cn

Hongyang Zhang
University of Waterloo
hongyang.zhang@uwaterloo.ca

Abstract—Transfer-based adversarial example is one of the most important classes of black-box attacks. However, there is a trade-off between transferability and imperceptibility of the adversarial perturbation. Prior work in this direction often requires a fixed but large ℓ_p -norm perturbation budget to reach a good transfer success rate, leading to perceptible adversarial perturbations. On the other hand, most of the current unrestricted adversarial attacks that aim to generate semantic-preserving perturbations suffer from weaker transferability to the target model. In this work, we propose a *geometry-aware framework* to generate transferable adversarial examples with minimum changes. Analogous to model selection in statistical machine learning, we leverage a validation model to select the best perturbation budget for each image under both the ℓ_∞ -norm and unrestricted threat models. We propose a principled method for the partition of training and validation models by encouraging intra-group diversity while penalizing extra-group similarity. Extensive experiments verify the effectiveness of our framework on balancing imperceptibility and transferability of the crafted adversarial examples. The methodology is the foundation of our entry to the *CVPR’21 Security AI Challenger: Unrestricted Adversarial Attacks on ImageNet*, in which we ranked 1st place out of 1,559 teams and surpassed the runner-up submissions by 4.59% and 23.91% in terms of final score and average image quality level, respectively. Code is available at <https://github.com/Equationliu/GA-Attack>.

I. INTRODUCTION

Though deep neural networks have exhibited impressive performance in various fields [1], [2], they are vulnerable to adversarial examples [3]–[7], where test inputs that have been modified slightly strategically cause misclassification. Adversarial examples have posed serious threats to various security-critical applications, such as autonomous driving [8] and face recognition [9]. Most positive results on adversarial attacks have focused on white-box settings [10], [11]. However, the problem becomes more challenging when it comes to the black-box setting, where the attacker has no information about the model architecture, hyper-parameters, and even the outputs of the black-box model. In this setting, adversarial examples are typically generated via *transfer-based* methods [3], [12], [13], *e.g.*, attacking an ensemble of accessible source models and hoping that the same adversarial examples are able to fool the unknown target/test model [14], [15].

Despite a large amount of work on transfer-based attacks, many fundamental questions remain unresolved. For example, existing transfer-based attacks [16]–[18] that search for adversarial examples in a fixed-radius ℓ_∞ -norm ball often require

a high perturbation budget to reach a satisfactory transfer success rate. However, such perturbations might be perceptible to humans (see Figs. 1 and 6). On the other hand, unrestricted attacks that aim to generate minimum human-imperceptible perturbations [19]–[21] suffer from weaker transferability to the target model. This is in part due to the difference between the decision boundaries of the source and target models. Given the trade-off between transferability and imperceptibility, one of the long-standing questions is generating transferable adversarial examples by minimum changes of natural examples.

A. Our Methodology and Results

In this work, we propose a novel geometry-aware framework to generate transferable unrestricted adversarial examples with minimum changes. Our intuition is that the smallest perturbation budgets *w.r.t.* distinct images should be different (see Fig. 1) and should depend on their geometrical relationship with the decision boundary of the target model (see Fig. 2). Unfortunately, finding transferable minimum-budget adversarial perturbations is an intractable optimization problem (see Eq. (6)) as the target model is unknown. We approximately solve this problem by discretizing the continuous space of perturbation radius into a finite set and choosing the minimum perturbation budget that is able to fool the test model. The main challenge here is to evaluate whether a given perturbation can transfer well to the unknown target model [22], [23].

To overcome this challenge, we split all accessible white-box source models into training and validation sets, where adversarial perturbations are crafted only on the training set. We use the validation set to select the smallest perturbation radius for each input that suffices to fool the validation model with a certain confidence level through an *early-stopping* mechanism. When the training (or validation) set consists of multiple models, we use their average ensemble [14]. Experimentally, our method yields a significant performance boost on the trade-off (leading to higher S_{total} in Table. II) between transferability and imperceptibility. As shown in Fig. 4, the transfer success rate of our method GA-DTMI-FGSM surpasses the baseline DTMI-FGSM (see Eq. (1)) by up to 16% in absolute value under the same average perturbation reward (see Eq. (5)). Besides, our method GA-DMI-FSA is able to generate semantic-preserving yet transferable unrestricted adversarial examples under the unrestricted threat model (see Eq. (3), Figs. 1, 6, 10 and 11).

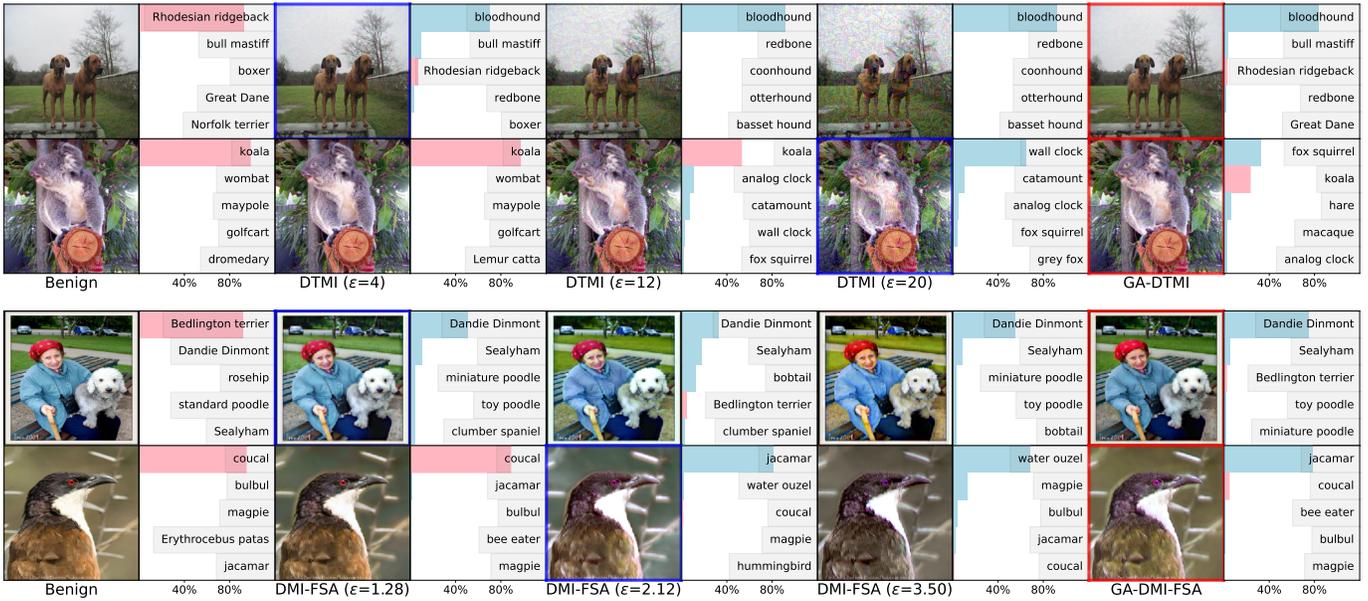


Fig. 1: Comparison between our method and two baselines under both the ℓ_∞ -norm (top) and unrestricted (bottom) threat models using various perturbation radii. In the even columns, we present the top-5 confidence bars of the target model for the images in the left. The ground-truth label is marked by pink and other labels are marked by blue. In each row, the misclassified adversarial example with minimum perturbation radius is highlighted by a blue bounding box, indicating that *the perturbation budgets required for distinct images are different*. Note that the “human-imperceptible” constraint is violated when the ℓ_∞ -norm perturbation radius is too large. However, our GA framework generates transferable unrestricted adversarial examples (highlighted by red bounding boxes) with lower budgets and smaller changes when compared to the benign images.

B. Summary of Our Contributions

- We propose a Geometry-Aware (GA) framework, where fixed-budget attacking methods can be integrated, to generate transferable unrestricted adversarial examples with approximately minimum changes. To the best of our knowledge, we are the first to explore transfer-based black-box attacks with adaptive perturbation budgets.
- Under ℓ_∞ -norm setting, our GA framework improves the imperceptibility of the crafted adversarial examples by a large margin without the decrease of transfer success rate (see Fig. 4). By applying our method GA-DTMI-FGSM to the CVPR’21 Security AI Challenger [24], we ranked 1st place out of 1,559 teams and surpassed the runner-up submissions by 4.59% and 23.91% in terms of final score and average image quality level, respectively.
- Under unrestricted setting, we propose a transfer-based unrestricted attack (see Eq. (3)) by combining the *white-box* feature space attack [25] with transfer-based ℓ_∞ -norm attacks to generate semantic-preserving yet transferable adversarial examples (see Figs. 1 and 10). Moreover, the crafted adversarial examples transfer well to adversarially robust models (see Table. III, Figs. 6 and 11).

II. RELATED WORK

A. ℓ_p -norm Adversarial Examples

Existing gradient-based white-box attacks either search for adversarial examples in a fixed ℓ_p -norm ball [26], [27], or optimize the perturbation for each image independently to get

a *minimum-norm* solution such as DeepFool [28], CW [29], and fast adaptive boundary attack [30]. However, white-box assumption usually does not hold in real-world scenarios. In query-based black-box setting, attackers utilize output logits [31], [32] or predicted label [33], [34] of the target model to generate adversarial examples. But these attacks typically suffer from high query complexity, making it easy to be detected [35]. Transfer-based black-box attacks [36]–[39] can pose serious threats in practice as they need no information about the defense models. Dong *et al.* [16] boosted transferability by integrating momentum into gradient-based methods. Liu *et al.* [14] found that attacking a group of substituted source models simultaneously can improve transferability. Besides, transferability benefits from input transformations such as input diversity [17] and translation-invariant method [18].

B. Unrestricted Adversarial Examples

The ℓ_p -norm distance is not an ideal perceptual similarity metric [40], [41], which oversimplifies the diversity of real-world perturbations. Unrestricted adversarial examples have received significant attention in recent years [42]. Most of the current unrestricted attacks aim to generate imperceptible adversarial examples under *white-box* setting, such as geometric transformations [19], [43], [44] and distance metrics beyond ℓ_p norm [20], [21]. Color-based attacks [45]–[50] were also proposed to generate large but imperceptible perturbations, however, the modified color can sometimes be unnatural. Instead of optimizing in the input space, generative approaches [51]–[54] search for adversarial embeddings in

the latent space. Style transfer [50], [55] is inherently an unrestricted attack as it preserves the semantic of the content image. However, constructing transferable unrestricted adversarial examples is still less explored. In this work, we will fill this gap by combining the white-box feature space attack [25] with transfer-based ℓ_∞ -norm attacks to generate semantic-preserving yet transferable unrestricted adversarial examples.

C. Adversarial Defenses

There have been long-standing arms races between defenders and attackers. *Adversarial training* [4] is one of the most promising defense methods. Many variants of adversarial training framework were proposed, e.g., ensemble adversarial training [14] for transfer-based attacks, PGD-based adversarial training [26], and TRADES [56] with a new robust loss based on the trade-off between robustness and accuracy. Geometry-aware instance-reweighted adversarial training [57], which is proven falling into gradient masking [58], shares similar insights with us that the importance of distinct inputs in adversarial training should be different. Laidlaw *et al.* [21] integrate adversarial training with Learned Perceptual Image Patch Similarity (LPIPS) [59], aiming to improve robustness against perturbations that were unseen during training. Unlike empirical defenses, Certified defenses [60]–[63] could provide robustness guarantee under a certain ℓ_p -norm budget.

III. PRELIMINARIES

Notation. A deep neural network classifier can be described as a function $f(\mathbf{x}; \boldsymbol{\theta}) : \mathcal{X} \rightarrow \mathbb{R}^C$, parameterized by weights $\boldsymbol{\theta}$, which maps a vector $\mathbf{x} \in \mathcal{X}$ to its output logits. Given an input \mathbf{x} of class $y \in \{1, 2, \dots, C\}$, the predicted label of $f(\mathbf{x}; \boldsymbol{\theta})$ is $\hat{f}(\mathbf{x}) := \arg \max_j f_j(\mathbf{x}; \boldsymbol{\theta})$, where $f_j(\mathbf{x}; \boldsymbol{\theta})$ represents the j -th entry of $f(\mathbf{x}; \boldsymbol{\theta})$. We use $L(f(\mathbf{x}; \boldsymbol{\theta}), y)$ to represent the cross-entropy loss and denote the ε -neighborhood of \mathbf{x} by $\mathbb{B}(\mathbf{x}, \varepsilon) := \{\mathbf{x}' \in \mathcal{X} : \mathcal{D}(\mathbf{x}, \mathbf{x}') \leq \varepsilon\}$, where \mathcal{D} is a distance metric that describes the change between the adversarial example \mathbf{x}' and the nature example \mathbf{x} . We denote the black-box test model by g , and split the set of accessible source models $\Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ into the set of training models f and the set of validation models h .

A. Transfer-based ℓ_∞ -norm Attacks

Existing transfer-based attacks typically search for adversarial examples in a fixed-radius ℓ_p -norm ball, i.e., $\mathcal{D}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x}' - \mathbf{x}\|_p \leq \varepsilon$. Various methods were proposed to boost transferability of the generated adversarial examples, such as input Diversity Iterative Fast Gradient Sign Method (DI-FGSM) [17], Momentum-based Iterative (MI-FGSM) method [16] and Translation-invariant Iterative (TI-FGSM) method [18]. We formulate a strong ℓ_∞ -norm baseline DTMI-FGSM by combining all these techniques, i.e.,

$$\mathbf{m}_{t+1} = \gamma \cdot \mathbf{m}_t + \frac{\mathbf{W} * \nabla_{\mathbf{x}_t} L(f(T(\mathbf{x}_t, p); \boldsymbol{\theta}), y)}{\|\mathbf{W} * \nabla_{\mathbf{x}_t} L(f(T(\mathbf{x}_t, p); \boldsymbol{\theta}), y)\|_1}, \quad (1)$$

$$\mathbf{x}_{t+1} = \Pi_{\mathbb{B}(\mathbf{x}, \varepsilon)}(\mathbf{x}_t + \alpha \cdot \text{sign}(\mathbf{m}_{t+1})),$$

where $\mathbf{m}_0 = \mathbf{0}$, \mathbf{W} is a pre-defined kernel with a convolution operation $*$, α is the step size, Π is the projection operator,

and γ is the decay factor for the momentum term. $T(\mathbf{x}_t, p)$ represents the input transformation on \mathbf{x}_t with probability p . When $\gamma = 0$, DTMI-FGSM attack degenerates to the DTI-FGSM attack. When $p = 0$, DTMI-FGSM attack degenerates to the DMI-FGSM attack.

B. Transfer-based Unrestricted Attack

Inspired by prior work [64] in style transfer, Xu *et al.* [25] tries to find stylized adversarial examples by assuming that the image pairs from the same class share consistent content and differ mainly in their styles. Here we propose to generate semantic-preserving yet transferable unrestricted adversarial examples by combining the Feature Space Attack (FSA) [25] with transfer-based ℓ_∞ -norm attacks [16], [17]. Given an encoder ϕ , we extract the style features of input \mathbf{x} as channel-wise mean $\boldsymbol{\mu}(\phi(\mathbf{x})) \in \mathbb{R}^C$ and channel-wise standard deviation $\boldsymbol{\sigma}(\phi(\mathbf{x})) \in \mathbb{R}^C$. Specifically,

$$\mu_c = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W \phi_c(\mathbf{x})_{hw},$$

$$\sigma_c = \sqrt{\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W (\phi_c(\mathbf{x})_{hw} - \mu_c)^2}, \quad (2)$$

where $\phi(\mathbf{x}) \in \mathbb{R}^{C \times H \times W}$ represents the latent embedding. Xu *et al.* [25] adds adversarial perturbations on $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ before projecting $\phi(\mathbf{x})$ back to the input space \mathcal{X} with a pre-trained¹ decoder ϕ^{-1} , namely,

$$\tilde{\phi}(\mathbf{x}) = e^{\boldsymbol{\tau}^\sigma} \cdot (\phi(\mathbf{x}) - \boldsymbol{\mu}) + e^{\boldsymbol{\tau}^\mu} \cdot \boldsymbol{\mu},$$

$$\mathbf{x}' = \phi^{-1}(\tilde{\phi}(\mathbf{x})), \quad \|\boldsymbol{\tau}^\mu\|_\infty \leq \ln \varepsilon, \|\boldsymbol{\tau}^\sigma\|_\infty \leq \ln \varepsilon, \quad (3)$$

where $\tilde{\phi}(\mathbf{x})$ enlarges or shrinks the mean $\boldsymbol{\mu}$ and the standard deviation $\boldsymbol{\sigma}$ of the embedding $\phi(\mathbf{x})$ by a factor of $e^{\boldsymbol{\tau}^\mu}$ and $e^{\boldsymbol{\tau}^\sigma}$, respectively. In this way, the distance metric $\mathcal{D}(\mathbf{x}, \mathbf{x}') = \max(e^{\|\boldsymbol{\tau}^\mu\|_\infty}, e^{\|\boldsymbol{\tau}^\sigma\|_\infty}) \leq \varepsilon$. In order to preserve the semantic of the unrestricted adversarial example \mathbf{x}' , a content loss was added during the attacking process, i.e.,

$$\min_{\boldsymbol{\tau}^\mu, \boldsymbol{\tau}^\sigma} \mathcal{L}(\mathbf{x}', y) = \lambda \cdot \mathcal{L}_{\text{top-5}}(f(\mathbf{x}'; \boldsymbol{\theta}), y) + \|\phi(\mathbf{x}') - \tilde{\phi}(\mathbf{x})\|_2,$$

where λ balance the trade-off between adversarial and the content loss. Following Xu *et al.* [25], we set $\lambda = 128$ and use the top-5 margin loss for adversarial attack. With all above, the unrestricted attack (see Eq. (3)) can be solved by conventional ℓ_∞ -norm attack on parameters $\boldsymbol{\tau}^\mu$ and $\boldsymbol{\tau}^\sigma$. Moreover, the same techniques in Sec. III-A such as input diversity $T(\mathbf{x}', p)$ (only for the margin loss) and momentum-based method can be integrated to improve transferability, i.e.,

$$\mathbf{m}_{t+1} = \gamma \cdot \mathbf{m}_t + \frac{\nabla_{\boldsymbol{\tau}_t} \mathcal{L}(T(\mathbf{x}', p), y)}{\|\nabla_{\boldsymbol{\tau}_t} \mathcal{L}(T(\mathbf{x}', p), y)\|_1}, \quad (4)$$

$$\boldsymbol{\tau}_{t+1} = \Pi_{\mathbb{B}(\boldsymbol{\tau}, \varepsilon)}(\boldsymbol{\tau}_t - \alpha \cdot \text{sign}(\mathbf{m}_{t+1})).$$

When $\gamma = 0$, the DMI-FSA attack degenerates to the DI-FSA attack.

¹We use the official pre-trained shallowest decoder: <https://github.com/qiulingxu/FeatureSpaceAttack>.

TABLE I: **An overview of all considered networks for generating adversarial examples.** Top-1_{ImageNet} represents the accuracy on the ILSVRC 2012 validation set while Top-1₁₀₀₀ represents the accuracy on the randomly selected 1000 images.

Training	Index	Model Name	Top-1 _{ImageNet}	Top-1 ₁₀₀₀	Index	Model Name	Top-1 _{ImageNet}	Top-1 ₁₀₀₀
Normal	0	ViT-S/16	76.01%	99.8%	1	ViT-B/16	81.08%	99.1%
	2	Swin-B/patch4-window7	84.23%	99.4%	3	ResNeXt101-32x8d-sws1	83.62%	99.9%
	4	ResNeXt50-32x4d-ssl	78.90%	99.7%	5	ResNet50-sws1	79.97%	99.5%
	6	Inception-v3	76.94%	100.0%	7	Inception-ResNet-v2	79.85%	99.9%
Ensemble	8	Ens3-adv-Inception-v3	76.49%	100.0%	9	Ens-adv-Inception-ResNet-v2	78.98%	99.9%

as DMI-FSA (see Eq. (4)) under the perturbation budget

$$\varepsilon_k = \frac{k}{K} \times \varepsilon, \quad k = 1, 2, \dots, K.$$

Each sub-procedure starts from the solution of last sub-procedure to accelerate the convergence. To obtain a minimum-radius solution, we perform an early-stopping mechanism at the end of each sub-procedure if the probability of true class on the validation model h is smaller than a threshold η , i.e.,

$$\mathbb{P}(\hat{h}(\mathbf{x}) = y) = \frac{\exp(h_y(\mathbf{x}; \boldsymbol{\theta}))}{\sum_j \exp(h_j(\mathbf{x}; \boldsymbol{\theta}))} < \eta. \quad (7)$$

Our GA framework is summarized in Algorithm 1 and illustrated in Fig. 2. Note that the output of GA framework is related to the choice of the training model f and the validation model h . Thus it is important to figure out *which partition of the source models performs better*.

We split n pre-trained models $\Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ into k training models and $n - k$ validation models. Instead of traversing all possible partitions to select the optimal split by querying the test model g by C_n^k times, we propose a query-free approach that only utilizes the information of transferability between the pre-trained models (see Fig. 3). Let w_{ij} ($w_{ij} \neq w_{ji}$) be the transfer success rate ($\frac{N_{ij}}{N}$ in Eq. (5)) from the source model ϕ_i to the target model ϕ_j under a fixed-radius attack (e.g., DTMI-FGSM). Denote the binary partition function as \mathcal{G} . The training set and the validation set can be formulated as $\mathcal{T} = \{i \mid \mathcal{G}(\phi_i) = 0\}$ and $\mathcal{V} = \{j \mid \mathcal{G}(\phi_j) = 1\}$, respectively. We define the *partition loss* $\ell_{\mathcal{G}}$ as:

$$\begin{aligned} \ell_{\mathcal{G}} &= \frac{1}{k} \sum_{i \in \mathcal{T}} \ell_{\mathcal{G}}(\phi_i) + \frac{1}{n-k} \sum_{j \in \mathcal{V}} \ell_{\mathcal{G}}(\phi_j), \\ \ell_{\mathcal{G}}(\phi_i) &= \frac{1}{k-1} \sum_{t \neq i, t \in \mathcal{T}} w_{it} + \frac{1}{n-k} \sum_{t \in \mathcal{V}} w_{it}, \\ \ell_{\mathcal{G}}(\phi_j) &= \frac{1}{n-k-1} \sum_{t \neq j, t \in \mathcal{V}} w_{jt} + \frac{1}{k} \sum_{t \in \mathcal{T}} w_{tj}. \end{aligned} \quad (8)$$

For both the training set loss $\ell_{\mathcal{G}}(\phi_i)$ and the validation set loss $\ell_{\mathcal{G}}(\phi_j)$: 1) Minimizing the first formula on the right of Eq. (8) encourages intra-group diversity. To make the decision boundary of the ensemble model (f or h) more general and effective, we minimize the transfer success rate between any two pre-trained models inside the group. 2) Minimizing the second formula on the right of Eq. (8) is penalizing extra-group similarity. If $f = h$, the early-stopping mechanism in Eq. (7) will be triggered too early, leading to small adversarial

Training Model	ViT-S/16	ViT-B/16	Swin-B/patch4-window7	ResNeXt101-32x8d-sws1	ResNeXt50-32x4d-ssl	ResNet50-sws1	Inception-v3	Inception-ResNet-v2	Ens3-adv-Inception-v3	Ens-adv-Inception-ResNet-v2
ViT-S/16	1	0.97	0.68	0.61	0.81	0.75	0.74	0.68	0.58	0.38
ViT-B/16	0.87	1	0.56	0.52	0.69	0.61	0.54	0.49	0.44	0.3
Swin-B/patch4-window7	0.66	0.69	1	0.64	0.73	0.69	0.69	0.63	0.34	0.17
ResNeXt101-32x8d-sws1	0.35	0.39	0.32	1	0.99	0.89	0.89	0.66	0.64	0.31
ResNeXt50-32x4d-ssl	0.44	0.47	0.37	0.86	1	0.92	0.75	0.71	0.34	0.18
ResNet50-sws1	0.41	0.43	0.36	0.89	0.93	1	0.76	0.72	0.37	0.21
Inception-v3	0.3	0.26	0.17	0.41	0.63	0.58	1	0.76	0.33	0.12
Inception-ResNet-v2	0.35	0.32	0.24	0.57	0.72	0.69	0.86	1	0.41	0.21
Ens3-adv-Inception-v3	0.41	0.37	0.26	0.35	0.53	0.48	0.62	0.54	1	0.21
Ens-adv-Inception-ResNet-v2	0.35	0.35	0.27	0.4	0.5	0.47	0.52	0.56	0.44	1

Fig. 3: **Transferability between networks under DTMI-FGSM attack.** The rows stand for source models and the columns stand for target models. Adversarial examples transfer well between models with similar architectures.

perturbations for all inputs. Reducing the transferability from the training model f to the validation model h might improve generalization of the adversarial examples to unknown models. Empirically, we find that the proposed partition loss $\ell_{\mathcal{G}}$ negatively correlates with the average total score S_{total} .

V. EXPERIMENTS

A. Experimental Setup

Datasets & Networks. Similar to Xie *et al.* [17], we randomly select 1,000 images from ILSVRC 2012 validation set [67], which are almost correctly classified by all the attacking models. All these images are resized to $229 \times 229 \times 3$ beforehand. We consider eight normally trained models and two ensemble adversarially trained models [15]. The weights of all these models are publicly available [68]. More details about the networks are summarised in Table. I. The transferability between these models under ℓ_{∞} -norm setting is summarised in Fig. 3. It is much easier for the generated adversarial examples to transfer from vision transformers to convolutional neural networks (CNNs), which is consistent with the empirical observation in Shao *et al.* [5]. Surprisingly, the robustness of naturally trained vision transformers under transfer attack is even on par with two ensemble adversarially trained CNNs.

Implementation Details. Given the maximum perturbation size ε and number of sub-procedures K (5 as default) in our geometry-aware framework, we set the step size $\alpha = \frac{1.25 \times \varepsilon_k}{T}$ in the k -th sub-procedure, where the number of iterations T is set to 10 in the ℓ_∞ setting and 50 in the unrestricted setting. ε is set to 20 in the ℓ_∞ -norm setting and 3.5 in the unrestricted setting. When running a fixed-radius baseline at perturbation budget ε_k , we set the number of iteration as $\frac{T}{2} (1 + \frac{K\varepsilon_k}{\varepsilon})$ with step size α to keep the same total perturbation budget (the sum of step size across all iterations) as our geometry-aware framework for fair comparison. The reward function $\mathcal{F}_{\text{reward}}(\varepsilon_0)$ is set to $1/\varepsilon_0$ as smaller perturbation radius exhibits significantly higher image quality. For the momentum term, we set the decay factor $\mu = 1$ as in Dong *et al.* [16]. For DI-FGSM [17], we set the transformation probability to $p = 0.7$. The input is first randomly resized to be an $r \times r \times 3$ image with $r \in [(1 - \gamma)s, (1 + \gamma)s]$, and then padded to size $(1 + \gamma)s \times (1 + \gamma)s \times 3$. The transformed input is then resized to $s \times s \times 3$ for different input size s of various models, i.e., 224, 299 and 384. We set $\gamma = 0.1$ as default. For TI-FGSM [18], we use Gaussian kernel with kernel size 5×5 .

B. Balancing Transfer Success Rate and Perturbation Reward

Implementation Details. Benefiting from the adaptive choice of perturbation budgets, our geometry-aware framework can generate transferable unrestricted adversarial examples with smaller changes. In this experiment, the training model f and validation model h are an ensemble of models $\{2, 3, 5\}$ and $\{1, 4, 6\}$ in Table. I, respectively. The test model is Inception-ResNet-v2. The optimal threshold η (see Eq (7)) is searched from a finite set ranging from 0.001 to 0.9 by querying² the test model g to achieve the best average total score S_{total} . For each η , we execute our method and compute the average perturbation reward S_{APR} (the x -axis of each red point in Fig. 4). Then the corresponding fixed-radius baseline is run at the same x -axis. We conduct experiments on two threat models. For the ℓ_∞ -norm setting, we combine our Geometry-Aware (GA) framework with DI-FGSM, DTI-FGSM, DTMI-FGSM, and Admix-DTI-FGSM [38] (limited by the memory of a single NVIDIA RTX 3090, we set the number of admixed images $m_1 = 3$ and the number of randomly sampled images from other categories $m_2 = 2$), named GA-DI-FGSM, GA-DTI-FGSM, GA-DTMI-FGSM, and GA-Admix-DTI-FGSM, respectively; For the unrestricted setting, we combine our GA framework with DMI-FSA and DI-FSA, named GA-DMI-FSA and GA-DI-FSA, respectively.

Experimental Results. We present the contour plot of average total score in Fig. 4, where the improvement of our method upon baselines depends on the choice of hyper-parameter η (leading to different S_{APR}). Fixing S_{APR} as 0.115, our

²In contrast to conventional query-based attacks that need the logits or predicted label on the target model, we query whether an adversarial example transfers to the target model successfully. Besides, we have prior information on η which depends on the similarity between f, h and g , making the query complexity rather limited.

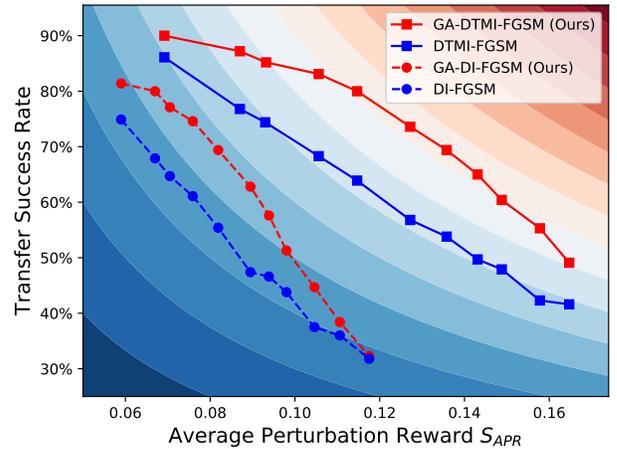


Fig. 4: **Contour of average total score S_{total} (higher is better)**. Fixing transfer success rate as 80%, our approach GA-DTMI-FGSM surpasses the baseline DTMI-FGSM (see Eq. (1)) by up to 43.35% in terms of average perturbation reward.

approach GA-DTMI-FGSM surpasses the baseline DTMI-FGSM by up to 16.1% in terms of transfer success rate. As shown in Table. II, our approach yields a significant performance boost on the average total score S_{total} across various threat models, especially in the ℓ_∞ -norm setting where both the transfer success rate and S_{APR} are improved.

TABLE II: **Comparison of our method with baselines.** We report the results when both our approach and baselines achieve highest average total score S_{total} . TSR: Transfer Success Rate.

Method	TSR (\uparrow)	S_{APR} (\uparrow)	S_{total} (\uparrow)
DI-FGSM [17]	61.1%	0.0759	4.64%
GA-DI-FGSM	69.4%	0.0819	5.68%
DTI-FGSM [18]	57.3%	0.1101	5.68%
GA-DTI-FGSM	67.9%	0.1176	7.98%
DTMI-FGSM	63.9%	0.1147	7.33%
GA-DTMI-FGSM	69.4%	0.1358	9.42%
Admix-DTI-FGSM [38]	68.1%	0.1248	8.50%
GA-Admix-DTI-FGSM	82.5%	0.1299	10.72%
DI-FSA	48.3%	0.5328	25.73%
GA-DI-FSA	50.4%	0.5541	27.93%
DMI-FSA	51.3%	0.5616	28.81%
GA-DMI-FSA	58.3%	0.5355	31.32%

C. Case Study: CVPR'21 Security AI Challenger

In the CVPR'21 Security AI Challenger: Unrestricted Adversarial Attacks on ImageNet [24], contestants were asked to submit adversarial examples without any access to the defense models. The dataset is a subset of ILSVRC 2012 validation set [67], which consists of 5,000 images with 5 images per class. The final score of each submission was manually scored from two aspects: 1) image semantic and 2) quality. If the semantic of the submitted image changes (judged by human referees), then $S_s = 0$, otherwise $S_s = 1$. The image quality S_q (equivalent to our reward function

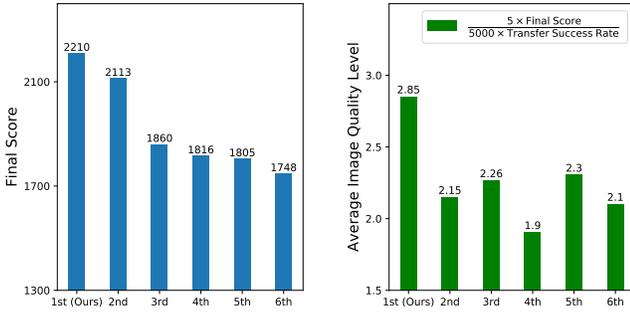


Fig. 5: **Top-6 results** in the CVPR’21 Security AI Challenge: Unrestricted Adversarial Attacks on ImageNet. The final scores were manually scored by multiple human referees.

$\mathcal{F}_{\text{reward}}$) was quantified with five levels $S_q \in \{1, 2, 3, 4, 5\}$ by multiple human referees. The final score is given by $\sum_i \mathbb{1}\{\hat{g}(\mathbf{x}'_i) \neq y_i\} \times S_s(\mathbf{x}'_i) \times \frac{S_q(\|\mathbf{x}'_i - \mathbf{x}_i\|)}{5}$.

We apply our method GA-DTMI-FGSM ($\eta = 0.01$) to the competition, where our entry ranked 1st place out of 1,559 teams. In the adversarial competition, our training and validation models are both an ensemble of eight high-performance models. We report the final score and average image quality level (equivalent to our average perturbation reward) in Fig. 5. It shows that our method outperforms other approaches by a large margin. In particular, we surpass the runner-up submissions by 4.59% and 23.91% in terms of final score and average image quality level, respectively.

D. Transferable Unrestricted Adversarial Examples

Most of current defenses can be easily broken by unseen attacks in a white-box manner. Adversarial training against multiple ℓ_p -norm attacks [76] solved this issue partially, however, at the cost of robustness against single ℓ_p -norm attack. Laidlaw *et al.* [21] integrated adversarial training with Learned Perceptual Image Patch Similarity (LPIPS) [59], aiming to improve robustness against perturbations that were unseen during training. However, the proposed attack [21], similar to other unrestricted attacks [44], [46], suffers from weaker transferability to the target model. In practice, attackers

typically have no information about the defense models and the defenders do not have the ground truth to make pixel-level comparison (perturbation can be large as long as the generated adversarial examples are semantic-preserving). Therefore, we propose to benchmark classification models on ImageNet under *transfer-based unrestricted attacks*.

Implementation Details. For the ReColor attack [46], we set $\varepsilon = 1.0$ and iterations $T = 100$ which achieves 89.7% attack success rate on training model f (the same as Sec. V-B) and 9.1% transfer success rate on the test model Inception-ResNet-v2. For FSA attack [25], we set $\varepsilon = 3.0$ and $T = 500$ which achieves 54.2% attack success rate and 12.4% transfer success rate on the same training and test models (Note that our method GA-DMI-FSA achieves 95.5% attack success rate and 58.3% transfer success rate, indicating that the input diversity and momentum techniques in Eq. (4) boost both the attacking ability and transferability.). Besides six adversarially trained and two high-performance classification models, we select two state-of-the-art models on the ImageNet-R dataset [73].

Experimental Results. From Table. III, we can conclude the following observations: a) $\text{GA}_{\text{Admix-DTI}}$ transfers better than $\text{GA}_{\text{DTMI-FGSM}}$. b) ℓ_∞ -norm transfer attack can hardly break ℓ_∞ -norm adversarially trained models while the unrestricted attack (GA-DMI-FSA) reduces the accuracy of these models by a large margin (see also in Fig. 6). c) DeepAugment [73], which utilizes semantic-preserving augmentations during training, exhibits non-trivial robustness against GA-DMI-FSA attack. d) Efficientnet-l2-ns [74] performs well under all the transfer-based attacks and enjoys 86% accuracy against GA-DMI-FSA attack, showing that the distribution of our generated adversarial examples is close to the natural examples’. Note that Efficientnet-l2-ns is also the best-performing model on the ImageNet-V2 dataset [77], [78]. We visualize part of the transfer attack results on adversarially trained Resnext101-DenoiseAll [71] in Figs. 6, where our method GA-DMI-FSA is able to generate semantic-preserving yet transferable unrestricted adversarial examples. For more visualization results, please see Fig. 11 in Sec. VII-A.

TABLE III: **Benchmarking classification on Imagenet under transfer-based unrestricted attacks.** PGD_{40}^* indicates the PGD attack with 40 steps ($\varepsilon = \frac{4}{255}$). We denote the adversarial examples crafted by GA-DTMI-FGSM, $\text{GA}_{\text{Admix-DTI}}$, and GA-DMI-FSA in Table. II as $\text{GA}_{\text{DTMI-FGSM}}$, $\text{GA}_{\text{Admix-DTI}}$, and $\text{GA}_{\text{DMI-FSA}}$, respectively. **Bold** and underline indicate the lowest and second lowest in each row.

Defenses	Clean	PGD_{40}^*	ReColor	FSA	$\text{GA}_{\text{DTMI-FGSM}}$	$\text{GA}_{\text{Admix-DTI}}$	$\text{GA}_{\text{DMI-FSA}}$
Inception-ResNet-v2 _{Ens-adv} [15]	99.9%	10.0%	93.2%	90.0%	87.1%	78.1%	<u>43.5%</u>
Fast _{AT} [69]	66.9%	<u>35.7%</u>	62.8%	59.3%	64.7%	64.1%	28.5%
Free _{AT} [70]	77.3%	<u>40.6%</u>	71.6%	68.4%	74.0%	73.0%	35.1%
Resnet152-Base [71]	67.6%	<u>39.0%</u>	64.1%	61.2%	65.1%	65.7%	37.4%
Resnext101-DenoiseAll [71]	80.3%	<u>52.2%</u>	77.0%	73.8%	78.7%	78.1%	47.8%
Resnet152-Denoise [71]	72.2%	<u>41.8%</u>	68.2%	65.2%	70.7%	69.9%	40.3%
RVT-Tiny [72]	96.7%	0.0%	78.9%	81.3%	44.5%	<u>26.3%</u>	33.6%
DeepAugment+AugMix [73]	96.1%	0.0%	82.8%	89.3%	58.9%	<u>41.4%</u>	63.2%
Efficientnet-l2-ns [74]	99.5%	0.0%	95.7%	97.0%	81.3%	<u>78.6%</u>	86.0%
Swin-L/patch4-window-12 [75]	99.1%	0.0%	88.4%	90.7%	66.6%	<u>58.7%</u>	61.8%

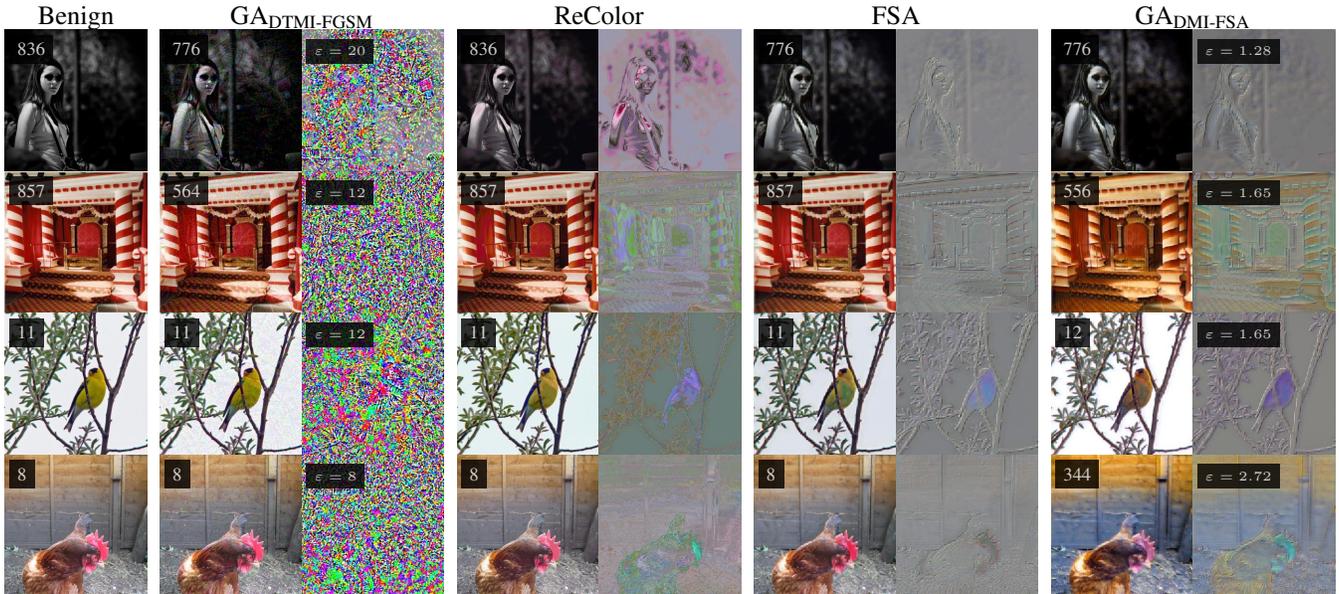


Fig. 6: **Visualization of transfer attack results on Resnext101-DenoiseAll [71].** For each image, we print its predicted label on model Resnext101-DenoiseAll in the upper left corner. For each transfer-based adversarial example, we present the perturbation on its right. For each perturbation crafted via our geometry-aware framework, we print its perturbation budget in the upper left corner. Although the transfer-based ℓ_∞ -norm attack GA-DTMI-FGSM is able to fool the defense test model to a certain extent, the generated perturbations can be “human-perceptible” (the first and third rows of $GA_{DTMI-FGSM}$). Besides, the other two unrestricted attacks suffer from weaker transferability when compared to our method $GA_{DTMI-FSA}$, which adjusts the images’ color and texture that ImageNet-trained CNNs might be biased to [79].

E. Ablation Studies and Discussions

The optimal η depends on the train-valid splitting. As declared in the implementation details in Sec. V-B, the threshold η is searched from a finite set ranging from 0.001 to 0.9. We now investigate how the hyper-parameter η will affect the average total score S_{total} . From Fig. 7, we observe that the optimal η^* varies across different splittings and can be larger if the transferability from the training model f to the test model g is higher enough. However, as shown in Table IV, the improvement of our GA framework upon fixed-radius baseline (DTMI-FGSM) is *stable* and independent of the splitting.

TABLE IV: **Comparison of the improvement upon fixed-radius baseline under different train-valid splittings.** The test model g is model 7 in Table I and we select three (two of them are highlighted in Fig. 8a) representative splittings according to the partition loss ℓ_G . For each setting, we repeat the experiment three times and report the mean and the standard deviation (in the parenthesis).

f	h	ℓ_G	Method	S_{APR} (\uparrow)	S_{total} (\uparrow)
{4, 5, 6}	{1, 2, 3}	2.22	DTMI	0.1080 (0.001)	8.66% (0.16%)
			GA-DTMI	0.1456 (6e-5)	11.6% (0.21%)
{2, 3, 5}	{1, 4, 6}	2.56	DTMI	0.1150 (0.006)	7.32% (0.07%)
			GA-DTMI	0.1417 (0.001)	9.48% (0.03%)
{1, 2, 3}	{4, 5, 6}	2.70	DTMI	0.1273 (0.007)	6.41% (0.07%)
			GA-DTMI	0.1291 (0.006)	8.23% (0.05%)

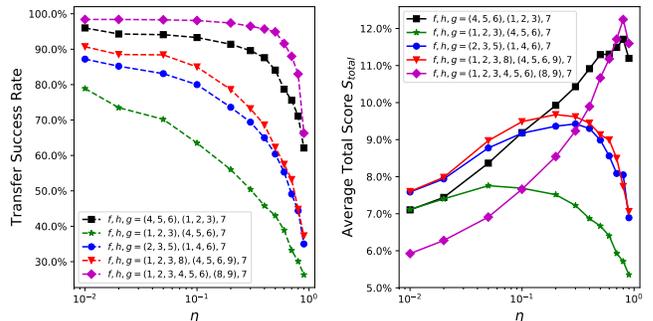


Fig. 7: **Comparison between various train-valid splittings.** The index in the legend corresponds to the model index in Table I. As the threshold η increases, the generated adversarial examples have higher confidence (probability of the true class) on the validation model h , leading to a lower transfer success rate (left). Besides, the optimal η^* that yields the maximum S_{total} is dependent on the partition (right).

The effectiveness of ℓ_G under different numbers of pre-trained models and different kinds of pre-trained models. To investigate the robustness of the proposed partition loss ℓ_G under different settings, we carefully design controlled experiments in Figs. 8 and 9. There are total C_n^k kinds of partitions when selecting k training models from n pre-trained models. Given n and k , we run our method GA-DTMI-FGSM for all the train-valid splittings and obtain a scatter plot with C_n^k points. We observe a strong negative correlation between

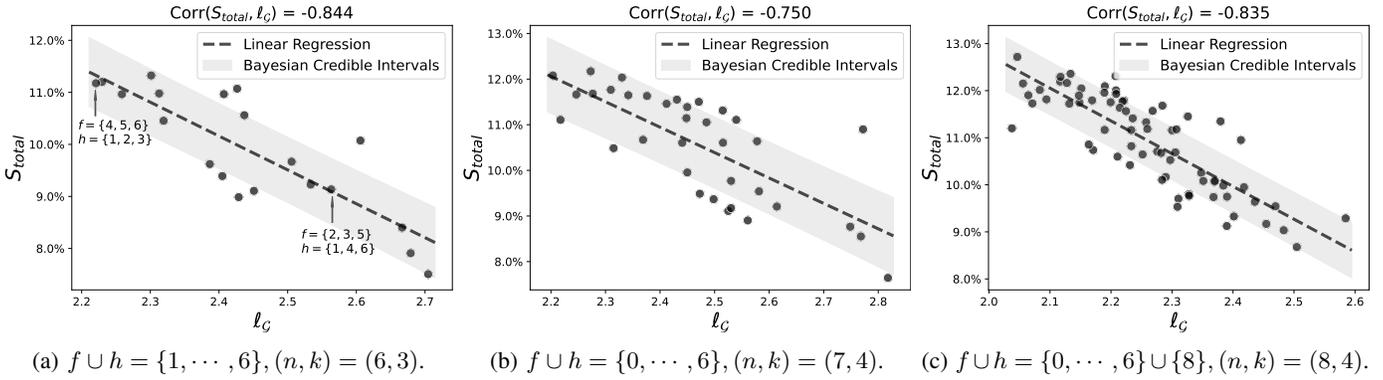


Fig. 8: **The effectiveness of the partition loss ℓ_G under different numbers of pre-trained models.** The test model g is model 7 in Table I. We conduct bayesian ridge regression and plot the mean of the predictive distribution as dashed lines. The Bayesian Credible Intervals range from mean - standard deviation (of the predictive distribution) to mean + standard deviation.

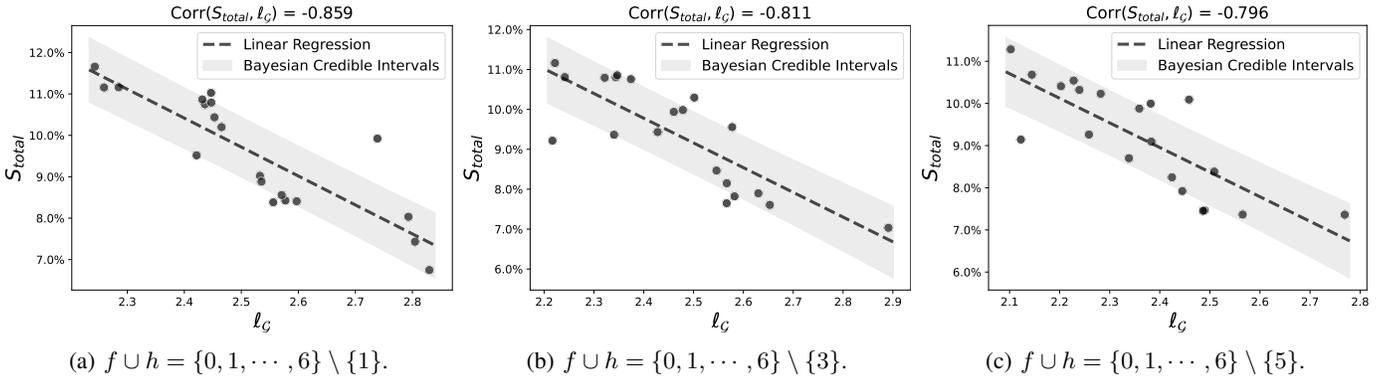


Fig. 9: **The effectiveness of the partition loss ℓ_G under different kinds of pre-trained models.** The test model g is model 7 in Table I. We simulate different kinds of pre-trained models by removing one model from the fixed set such that $(n, k) = (6, 3)$.

the partition loss ℓ_G and the average total score S_{total} . For example, the average Pearson correlation coefficient over the six scatterplots in Figs. 8 and 9 is around -0.82. Moreover, the negative correlation is significant and consistent across different numbers of pre-trained models (see Fig. 8) and different kinds of pre-trained models (see Fig. 9).

VI. CONCLUSION

In this work, we propose a geometry-aware framework, where fixed-radius methods can be integrated to generate transferable unrestricted adversarial examples with minimum changes. Under ℓ_∞ -norm setting, our framework could improve the imperceptibility of the crafted adversarial examples by a large margin without the decrease of transfer success rate. Besides, we propose a transfer-based unrestricted attack by combining the white-box feature space attack with transfer-based ℓ_∞ -norm attacks to generate semantic-preserving yet transferable unrestricted adversarial examples.

Acknowledgements. Fangcheng Liu and Chao Zhang are supported by the National Nature Science Foundation of China under Grant 62071013 and 61671027, and National Key R&D Program of China under Grant 2018AAA0100300. Hongyang Zhang is supported by NSERC Discovery Grant RGPIN-2022-03215, DGEGR-2022-00357.

REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778. 1
- [2] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021. 1
- [3] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *ICLR*, 2014. 1
- [4] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *ICLR*, 2015. 1, 3
- [5] R. Shao, Z. Shi, J. Yi, P.-Y. Chen, and C.-J. Hsieh, "On the adversarial robustness of visual transformers," *arXiv preprint arXiv:2103.15670*, 2021. 1, 5
- [6] S. Bhojanapalli, A. Chakrabarti, D. Glasner, D. Li, T. Unterthiner, and A. Veit, "Understanding robustness of transformers for image classification," in *ICCV*, October 2021, pp. 10 231–10 241. 1
- [7] Y. Bai, J. Mei, A. Yuille, and C. Xie, "Are transformers more robust than cnns?" in *NeurIPS*, 2021. 1
- [8] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016. 1
- [9] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *BMVC*. BMVA Press, September 2015, pp. 41.1–41.12. 1
- [10] A. Athalye, N. Carlini, and D. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," in *ICML*, vol. 1, 2018, pp. 436–448. 1

- [11] F. Tramer, N. Carlini, W. Brendel, and A. Madry, “On adaptive attacks to adversarial example defenses,” in *NeurIPS*, vol. 33. Curran Associates, Inc., 2020, pp. 1633–1645. 1
- [12] N. Papernot, P. McDaniel, and I. Goodfellow, “Transferability in machine learning: from phenomena to black-box attacks using adversarial samples,” *arXiv preprint arXiv:1605.07277*, 2016. 1
- [13] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM conference on computer and communications security*, 2017, pp. 506–519. 1
- [14] Y. Liu, X. Chen, C. Liu, and D. Song, “Delving into transferable adversarial examples and black-box attacks,” in *ICLR*, 2017. 1, 2, 3
- [15] F. Tramer, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “Ensemble adversarial training: Attacks and defenses,” in *ICLR*, 2018. 1, 5, 7
- [16] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, “Boosting adversarial attacks with momentum,” in *CVPR*, 2018, pp. 9185–9193. 1, 2, 3, 6
- [17] C. Xie, Z. Zhang, Y. Zhou, S. Bai, J. Wang, Z. Ren, and A. Yuille, “Improving transferability of adversarial examples with input diversity,” in *CVPR*, 2019. 1, 2, 3, 5, 6
- [18] Y. Dong, T. Pang, H. Su, and J. Zhu, “Evading defenses to transferable adversarial examples by translation-invariant attacks,” in *CVPR*, 2019, pp. 4312–4321. 1, 2, 3, 6
- [19] C. Xiao, J.-Y. Zhu, B. Li, W. He, M. Liu, and D. Song, “Spatially transformed adversarial examples,” in *ICLR*, 2018. 1, 2
- [20] E. Wong, F. Schmidt, and Z. Kolter, “Wasserstein adversarial examples via projected sinkhorn iterations,” in *ICML*. PMLR, 2019, pp. 6808–6817. 1, 2
- [21] C. Laidlaw, S. Singla, and S. Feizi, “Perceptual adversarial robustness: Defense against unseen threat models,” in *ICLR*, 2021. 1, 2, 3, 7
- [22] S. Cheng, Y. Dong, T. Pang, H. Su, and J. Zhu, “Improving black-box adversarial attacks with a transfer-based prior,” in *NeurIPS*, vol. 32, 2019. 1, 4
- [23] Z. Katzir and Y. Elovici, “Who’s afraid of adversarial transferability?” *arXiv preprint arXiv:2105.00433*, 2021. 1, 4
- [24] Y. Chen, X. Mao, Y. He, H. Xue, C. Li, Y. Dong, Q.-A. Fu, X. Yang, W. Xiang, T. Pang *et al.*, “Unrestricted adversarial attacks on ImageNet competition,” *arXiv preprint arXiv:2110.09903*, 2021. 2, 6
- [25] Q. Xu, G. Tao, S. Cheng, and X. Zhang, “Towards feature space adversarial attack by style perturbation,” *AAAI*, vol. 35, no. 12, pp. 10523–10531, May 2021. 2, 3, 7
- [26] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *ICLR*, 2018. 2, 3
- [27] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *ICLR 2017 - Workshop Track Proceedings*, no. c, pp. 1–14, 2019. 2
- [28] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: a simple and accurate method to fool deep neural networks,” in *CVPR*, 2016, pp. 2574–2582. 2
- [29] N. Carlini and D. Wagner, “Towards evaluating the robustness of neural networks,” in *2017 IEEE symposium on security and privacy (sp)*. IEEE, 2017, pp. 39–57. 2
- [30] F. Croce and M. Hein, “Minimally distorted adversarial examples with a fast adaptive boundary attack,” in *ICML*. PMLR, 2020, pp. 2196–2205. 2
- [31] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, “Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models,” in *Proceedings of the 10th ACM workshop on artificial intelligence and security*, 2017, pp. 15–26. 2
- [32] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, “Square attack: A query-efficient black-box adversarial attack via random search,” in *ECCV*, 2020, pp. 484–501. 2
- [33] W. Brendel, J. Rauber, and M. Bethge, “Decision-based adversarial attacks: Reliable attacks against black-box machine learning models,” in *ICLR*, 2018. 2
- [34] M. Cheng, T. Le, P.-Y. Chen, H. Zhang, J. Yi, and C.-J. Hsieh, “Query-efficient hard-label black-box attack: An optimization-based approach,” in *ICLR*, 2019. 2
- [35] D. Willmott, A. K. Sahu, F. Sheikholeslami, F. Condessa, and Z. Kolter, “You only query once: Effective black box adversarial attacks with minimal repeated queries,” *arXiv preprint arXiv:2102.00029*, 2021. 2
- [36] J. Lin, C. Song, K. He, L. Wang, and J. E. Hopcroft, “Nesterov accelerated gradient and scale invariance for adversarial attacks,” in *International Conference on Learning Representations*, 2020. 2
- [37] D. Wu, Y. Wang, S.-T. Xia, J. Bailey, and X. Ma, “Skip connections matter: On the transferability of adversarial examples generated with resnets,” in *International Conference on Learning Representations*, 2020. 2
- [38] X. Wang, X. He, J. Wang, and K. He, “Admix: Enhancing the transferability of adversarial attacks,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16158–16167. 2, 6
- [39] X. Wang, J. Ren, S. Lin, X. Zhu, Y. Wang, and Q. Zhang, “A unified approach to interpreting and boosting adversarial transferability,” in *International Conference on Learning Representations*, 2021. 2
- [40] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *ECCV*. Springer, 2016, pp. 694–711. 2
- [41] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in *CVPR*, 2017, pp. 1125–1134. 2
- [42] T. B. Brown, N. Carlini, C. Zhang, C. Olsson, P. Christiano, and I. Goodfellow, “Unrestricted adversarial examples,” *arXiv preprint arXiv:1809.08352*, 2018. 2
- [43] R. Alaifari, G. S. Alberti, and T. Gauksson, “ADef: an iterative algorithm to construct adversarial deformations,” in *ICLR*, 2019. 2
- [44] L. Engstrom, B. Tran, D. Tsipras, L. Schmidt, and A. Madry, “Exploring the landscape of spatial robustness,” in *ICML*. PMLR, 2019, pp. 1802–1811. 2, 7
- [45] H. Hosseini and R. Poovendran, “Semantic adversarial examples,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 1614–1619. 2
- [46] C. Laidlaw and S. Feizi, “Functional adversarial attacks,” in *NeurIPS*, 2019. 2, 7
- [47] Z. Zhao, Z. Liu, and M. Larson, “Towards large yet imperceptible adversarial image perturbations with perceptual color distance,” in *CVPR*, 2020, pp. 1039–1048. 2
- [48] Z. Zhao, Z. Liu, and M. A. Larson, “Adversarial color enhancement: Generating unrestricted adversarial images by optimizing a color filter,” in *BMVC*, 2020. 2
- [49] A. S. Shamsabadi, R. Sanchez-Matilla, and A. Cavallaro, “Colorfool: Semantic adversarial colorization,” in *CVPR*, 2020, pp. 1151–1160. 2
- [50] A. Bhattad, M. J. Chong, K. Liang, B. Li, and D. A. Forsyth, “Unrestricted adversarial examples via semantic manipulation,” in *ICLR*, 2020. 2, 3
- [51] Y. Song, R. Shu, N. Kushman, and S. Ermon, “Constructing unrestricted adversarial examples with generative models,” in *NeurIPS*, vol. 31. Curran Associates, Inc., 2018. 2
- [52] S. Goyal, C. Qin, P.-S. Huang, T. Cemgil, K. Dvijotham, T. Mann, and P. Kohli, “Achieving robustness in the wild via adversarial mixing with disentangled representations,” in *CVPR*, 2020, pp. 1211–1220. 2
- [53] H. Qiu, C. Xiao, L. Yang, X. Yan, H. Lee, and B. Li, “Semanticadv: Generating adversarial examples via attribute-conditioned image editing,” in *ECCV*. Springer, 2020, pp. 19–37. 2
- [54] E. Wong and J. Z. Kolter, “Learning perturbation sets for robust machine learning,” in *ICLR*, 2021. 2
- [55] V. U. Prabhu, J. Whaley, and S. Francisco, “Art-attack! on style transfers with textures, label categories and adversarial examples,” 2018. 3
- [56] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, “Theoretically principled trade-off between robustness and accuracy,” in *ICML*, 2019. 3
- [57] J. Zhang, J. Zhu, G. Niu, B. Han, M. Sugiyama, and M. Kankanhalli, “Geometry-aware instance-reweighted adversarial training,” in *ICLR*, 2021. 3
- [58] D. Hitaj, G. Pagnotta, I. Masi, and L. V. Mancini, “Evaluating the robustness of geometry-aware instance-reweighted adversarial training,” *arXiv preprint arXiv:2103.01914*, 2021. 3
- [59] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *CVPR*, 2018, pp. 586–595. 3, 7
- [60] J. Cohen, E. Rosenfeld, and Z. Kolter, “Certified adversarial robustness via randomized smoothing,” in *ICML*. PMLR, 2019, pp. 1310–1320. 3
- [61] H. Salman, G. Yang, J. Li, P. Zhang, H. Zhang, I. Razenshteyn, and S. Bubeck, “Provably robust deep learning via adversarially trained

- smoothed classifiers,” in *NeurIPS*. Curran Associates Inc., 2019. 3
- [62] H. Zhang, H. Chen, C. Xiao, S. Goyal, R. Stanforth, B. Li, D. Boning, and C.-J. Hsieh, “Towards stable and efficient training of verifiably robust neural networks,” in *ICLR*, 2020. 3
- [63] K. Leino, Z. Wang, and M. Fredrikson, “Globally-robust neural networks,” in *ICML*, 2021. 3
- [64] X. Huang and S. Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” in *ICCV*, 2017, pp. 1501–1510. 3
- [65] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004. 4
- [66] M. Sharif, L. Bauer, and M. K. Reiter, “On the suitability of lp-norms for creating and preventing adversarial examples,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 1605–1613. 4
- [67] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *CVPR*. IEEE, 2009, pp. 248–255. 5, 6
- [68] R. Wightman, “Pytorch image models,” <https://github.com/rwightman/pytorch-image-models>, 2019. 5
- [69] E. Wong, L. Rice, and J. Z. Kolter, “Fast is better than free: Revisiting adversarial training,” in *ICLR*, 2020. 7
- [70] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, “Adversarial training for free!” in *NeurIPS*, vol. 32. Curran Associates, Inc., 2019. 7
- [71] C. Xie, Y. Wu, L. v. d. Maaten, A. L. Yuille, and K. He, “Feature denoising for improving adversarial robustness,” in *CVPR*, 2019, pp. 501–509. 7, 8, 12
- [72] X. Mao, G. Qi, Y. Chen, X. Li, S. Ye, Y. He, and H. Xue, “Rethinking the design principles of robust vision transformer,” *arXiv preprint arXiv:2105.07926*, 2021. 7
- [73] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. Zhu, S. Parajuli, M. Guo, D. Song, J. Steinhardt, and J. Gilmer, “The many faces of robustness: A critical analysis of out-of-distribution generalization,” *ICCV*, 2021. 7
- [74] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, “Self-training with noisy student improves ImageNet classification,” in *CVPR*, 2020, pp. 10 687–10 698. 7
- [75] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” *arXiv preprint arXiv:2103.14030*, 2021. 7
- [76] F. Tramer and D. Boneh, “Adversarial training and robustness for multiple perturbations,” in *NeurIPS*, vol. 32. Curran Associates, Inc., 2019. 7
- [77] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do ImageNet classifiers generalize to ImageNet?” in *ICML*. PMLR, 2019, pp. 5389–5400. 7
- [78] R. Taori, A. Dave, V. Shankar, N. Carlini, B. Recht, and L. Schmidt, “Measuring robustness to natural distribution shifts in image classification,” in *NeurIPS*, 2020. 7
- [79] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, “ImageNet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness,” *arXiv preprint arXiv:1811.12231*, 2018. 8

VII. APPENDIX

A. More Visualization Results

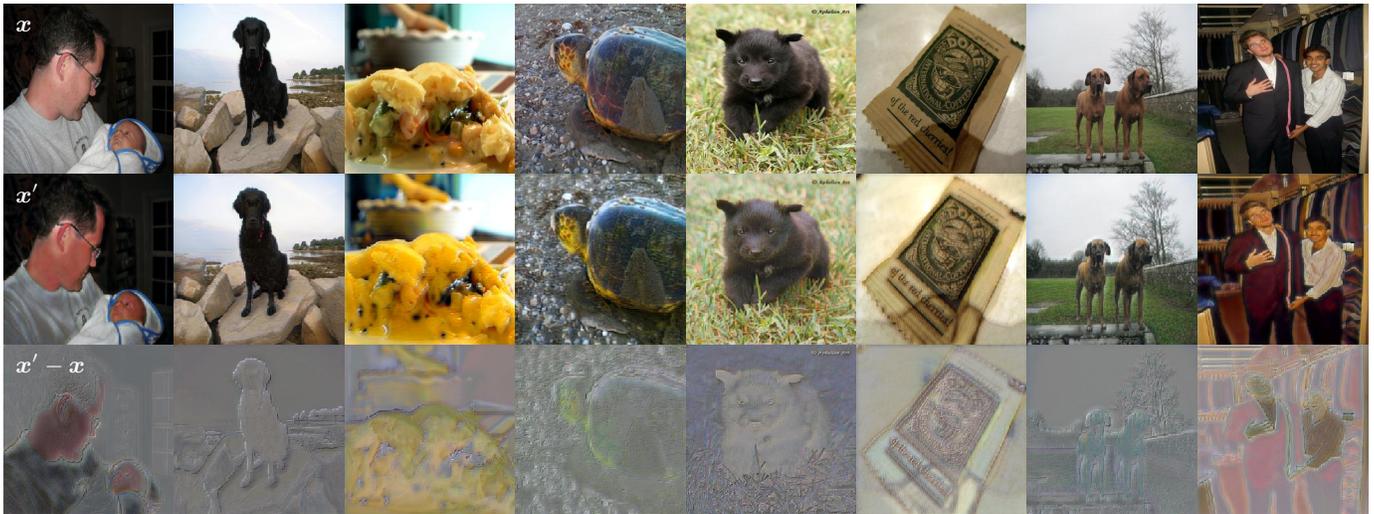


Fig. 10: **Adversarial examples of GA-DMI-FSA**, which are misclassified by *all* the models in Table. III. Top: benign examples x . Middle: unrestricted adversarial examples x' . Bottom: normalized adversarial perturbations $x' - x$.

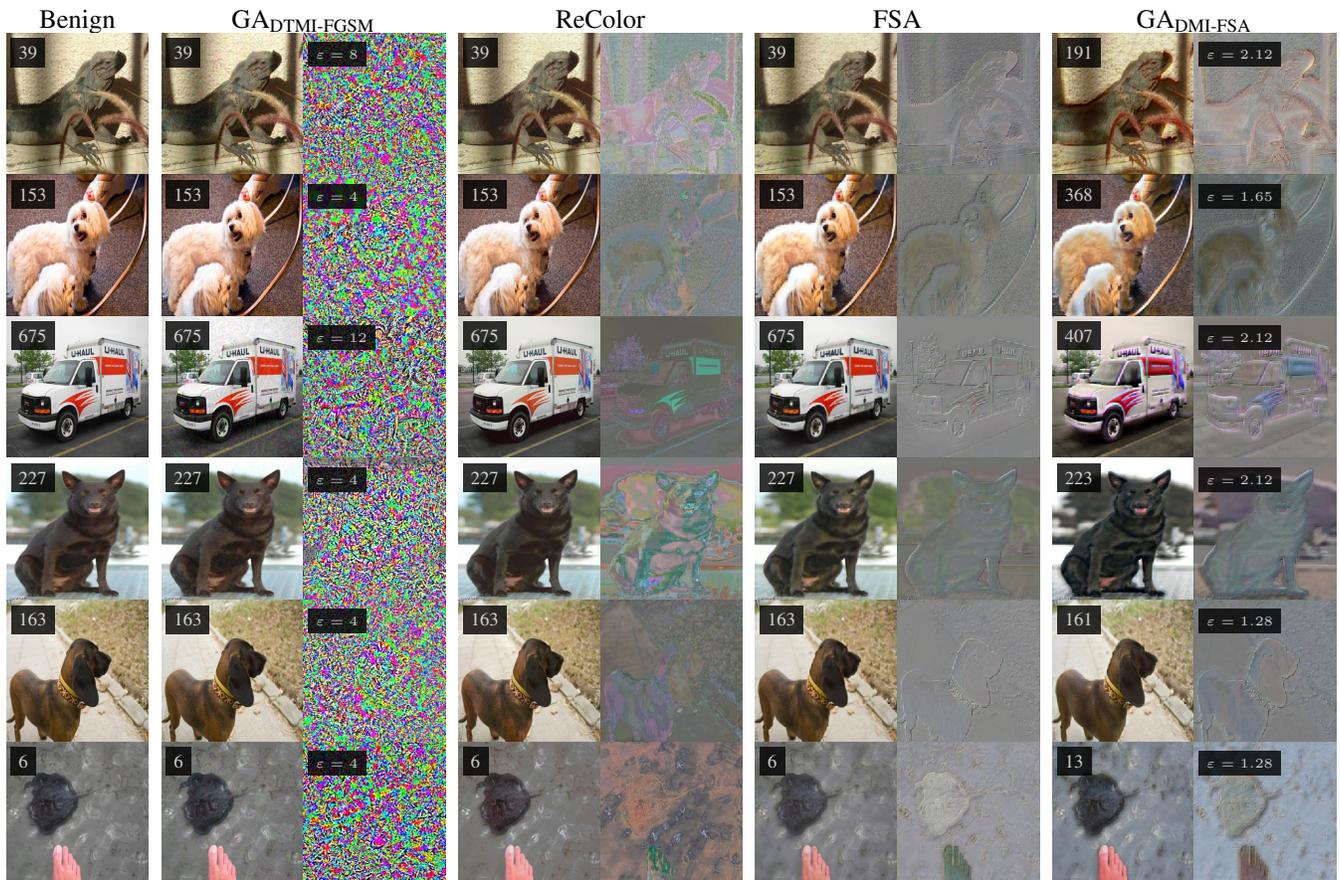


Fig. 11: **Visualization of transfer attack on Resnext101-DenoiseAll** [71].