

Towards Transferable Adversarial Attacks on Vision Transformers

Zhipeng Wei^{1,2*}, Jingjing Chen^{1,2*}, Micah Goldblum³, Zuxuan Wu^{1,2†}
Tom Goldstein³, Yu-Gang Jiang^{1,2†}

¹Shanghai Key Lab of Intelligent Information Processing, School of Computer Science, Fudan University

²Shanghai Collaborative Innovation Center on Intelligent Visual Computing

³Department of Computer Science, University of Maryland

zpwei21@m.fudan.edu.cn, chenjingjing@fudan.edu.cn, goldblumcello@gmail.com,

zwxu@fudan.edu.cn, tomg@cs.umd.edu, ygj@fudan.edu.cn

Abstract

Vision transformers (ViTs) have demonstrated impressive performance on a series of computer vision tasks, yet they still suffer from adversarial examples. In this paper, we posit that adversarial attacks on transformers should be specially tailored for their architecture, jointly considering both patches and self-attention, in order to achieve high transferability. More specifically, we introduce a dual attack framework, which contains a Pay No Attention (PNA) attack and a PatchOut attack, to improve the transferability of adversarial samples across different ViTs. We show that skipping the gradients of attention during backpropagation can generate adversarial examples with high transferability. In addition, adversarial perturbations generated by optimizing randomly sampled subsets of patches at each iteration achieve higher attack success rates than attacks using all patches. We evaluate the transferability of attacks on state-of-the-art ViTs, CNNs and robustly trained CNNs. The results of these experiments demonstrate that the proposed dual attack can greatly boost transferability between ViTs and from ViTs to CNNs. In addition, the proposed method can easily be combined with existing transfer methods to boost performance.

Introduction

Vision Transformers (ViTs) (Dosovitskiy et al. 2020) have attained excellent performance compared to state-of-the-art CNNs on standard image recognition tasks. However, ViTs are still vulnerable to security threats via adversarial examples (Goodfellow, Shlens, and Szegedy 2014; Wang et al. 2021a; Chen et al. 2021a; Wei et al. 2020), which are nearly indistinguishable from natural images while containing perturbations that result in incorrect predictions. It is known that a model of unknown structure can be attacked using adversarial images crafted with a different “surrogate” model (Liu et al. 2016). This cross-model transferability property of adversarial examples makes it feasible to attack a “black-box” model without knowing its architecture or other properties.

Cross-model transferability is a well-studied phenomenon for CNNs (Xie et al. 2019; Dong et al. 2018; Wei et al. 2021). High-performance transfer attacks typically employ

*These authors contributed equally.

†Correspondence to: Zuxuan Wu, Yu-Gang Jiang.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

data augmentation and advanced gradient calculations to prevent perturbations from over-fitting to the model, as this would diminish their success when transferring to black-box models. In contrast, relatively little is known about the transferability properties of attacks crafted on ViTs, and extending existing approaches that work well on CNNs to transformers is non-trivial due to significant structural differences. More specifically, ViTs receive a sequence of flattened patches from images as inputs and use a series of multi-headed self-attention (MSA) layers to learn relationships between patches. The use of standard attacks without considering these unique architectural features will result in suboptimal results, and give the user an inaccurate sense of the adversarial vulnerability of the transformer model class.

In light of these structural differences, we aim to generate highly transferable adversarial examples using white-box ViTs as proxy models to attack different black-box ViTs, normally trained CNNs, and robustly trained CNNs. In particular, we introduce a dual attack method tailored for the architecture of ViTs—attacking the attention mechanism and patches simultaneously, both of which are the core components of popular transformer architectures. In particular, we use a Pay No Attention (PNA) attack and a PatchOut attack to manipulate the attention mechanism and image features in parallel. The PNA attack improves adversarial transferability by treating the attention weights computed on the forward pass as constants. In other words, it does not propagate through the branch of the computation graph that produces the attention weights, as illustrated in Figure 1(a). This prevents patches from strongly interacting, as measured by the Shapley interaction index (Shapley 1988), which is known to help boost adversarial transferability (Wang et al. 2020). In addition, the PatchOut attack randomly samples a subset of patches to receive updates on each iteration of the attack crafting process. This is akin to using dropout (Hinton et al. 2012) on perturbation patches, and helps to combat over-fitting. It also shares similarities with random feature selection in Random Forests (Breiman 2004), Dropout and Diversity Input (DI) (Xie et al. 2019). Besides, we integrate the L_2 norm (Chen et al. 2021b) into our dual attack to prefer a large distance from the benign sample.

To validate the effectiveness of our approach, and to illustrate how the gradients of attention weights impair ad-




Figure 1: (a): An illustration of the self-attention block. The attention map calculated by Q and K is dropped in gradient backpropagation (red dashed box). (b): Adversarial examples are generated by using gradients that bypass attention (no gradients pass through the dotted red box in Figure 1(a)). We split the 12 attention blocks of ViT-B/16 into three chunks: MSA 1-4, 5-8 and 9-12. We consider the 8 different propagation paths that include/exclude every combination of these three blocks. The rightmost path skips all attention units and achieves the best black-box attack success rate (labeled in green) while the red path through all attention blocks achieves the worst black-box attack success rate (labeled in red). The attacks are crafted by BIM on 2,000 images under perturbations with L_∞ -norm bounded above by $\epsilon = 16$. The number of iterations is 10.

versarial transferability, we conduct two toy experiments using the BIM (Kurakin, Goodfellow, and Bengio 2016) attack on the white-box model ViT-B/16 (Dosovitskiy et al. 2020) using the ImageNet validation dataset. Figure 1(b) shows the results of PNA. We observe that the attack success rate (ASR) decreases as more attention gradients are used during backpropagation. Bypassing all gradients of attention (the green path) improves ASR from 29.92% to 42.47%. Figure 2 shows the results of PatchOut, where we randomly select ten patches as one input pattern. We call such a sparse perturbation a “ten-patch.” We see that stacking multiple ten-patches achieves a higher ASR than using perturbations produced by optimizing on the whole image at once. This observation demonstrates that stacking perturbations from diverse input patterns can help alleviate the over-fitting problem.

We further conduct extensive experiments on the ImageNet dataset (Russakovsky et al. 2015). We demonstrate that the proposed dual attack method significantly improves ASR against black-box ViTs, normally trained CNNs, and robustly trained CNNs. Compared to MI (Dong et al. 2018) and SGM (Wu et al. 2020a), the dual attack improves ASR




Figure 2: Attack success rate against black-box models with varying number of perturbation updates. We consider using either ten-patch updates or whole-image updates. The blue line generates adversarial perturbations by accumulating multiple ten-patch perturbations, which are generated by optimizing a single ten-patch at each stage. The orange line represents perturbations generated using whole-image updates. We average the attack success rate over numerous black-box models: DeiT-B, LeViT-256, CaiT-S-24, PiT-B, ConViT-B, TNT-S and Visformer-S.

by 15.86%, 27.68%, 23.52% on average for ViTs, normally trained CNNs and robustly trained CNNs respectively based on the adversarial transferability. Our results demonstrate the feasibility of using white-box ViTs to attack other black-box models, and they further demonstrate the vulnerability of seemingly robustly trained CNNs. We briefly summarize our primary contributions as follows:

- We study the attention mechanism in ViTs and propose the Pay No Attention (PNA) attack to craft adversarial examples without backpropping through attention. PNA is applicable to any gradient-based attack method and any attention-based neural network.
- We study how stacking perturbations using random subsets of patches can improve attack transferability, and we propose the PatchOut Attack to generate adversarial examples using different patches as input at each iteration.
- We conduct comprehensive transfer attack experiments using 4 different white-box ViTs against 8 black-box ViTs, 4 CNNs, and 3 robustly trained CNNs, showing that the proposed dual attack can greatly improve adversarial transferability and indicating that it can be combined with existing methods to further improve performance.

Related Work

Adversarial attacks (Yan, Wei, and Li 2020) are often studied under white- and black-box threat models. The white-box setting allows the adversary fully access to victim models, while the black-box setting only permits access to the output of a victim. Hence, the black-box setting is substantially more challenging. Transfer-based attacks generate adversarial examples on white-box models with the intent that the attacked samples will also be effective against black-box models. In this section, we review existing work on transfer-based attacks as well as on ViTs.

Transfer-based Attacks on CNNs

Fast Gradient Sign Method (FGSM) (Goodfellow, Shlens, and Szegedy 2014) and Basic Iterative Method (BIM) (Kurakin, Goodfellow, and Bengio 2017) are fundamental white-box attack methods which are commonly used for transfer-based attacks. FGSM performs a one-step update in the direction of the sign of the gradient to maximize loss. BIM applies FGSM several times iteratively with a small step size. Due to the fine-grained updates, BIM often generates more powerful adversarial examples than FGSM for white-box attacks, but BIM simultaneously achieves lower transferability than FGSM because of the over-fitting problem (Kurakin, Goodfellow, and Bengio 2016). To overcome this problem, several efforts have been made to improve adversarial transferability. These efforts can be roughly categorized into two groups: data augmentation and advanced gradient calculation. Data augmentation ensures that the attack’s effectiveness is invariant to input transformations. For example, the Diversity Input (DI) attack (Xie et al. 2019) applies random resizing and padding to the inputs with a fixed probability at each iteration. The Translation-Invariant (TI) attack (Dong et al. 2019) optimizes a perturbation over an ensemble of translated images by convolving the gradient with a linear or Gaussian kernel, motivated by the near translation-invariance of CNNs. Scale-Invariant Method (SIM) (Lin et al. 2019) optimizes a perturbation over an ensemble of scaled images at each iteration.

Advanced gradient calculations stabilize the update direction or add new terms to the loss function in order to better optimize adversarial perturbations or to reduce over-fitting. For example, Momentum Iterative (MI) attack (Dong et al. 2018) stabilizes update directions and escapes from poor local maxima by integrating a momentum term. Nesterov Accelerated Gradient (Lin et al. 2019) can also be viewed as another momentum term for improving transferability. Transferable Adversarial Perturbations (TAP) (Zhou et al. 2018) maximize the distance between natural images and their adversarial examples in intermediate feature maps and also uses regularization to smooth the resulting perturbations. The Attention-Guided Transfer Attack (ATA) (Wu et al. 2020b) also maximizes distance in the attention maps obtained by Grad-CAM (Selvaraju et al. 2017) for critical feature destruction. Skip Gradient Method (SGM) (Wu et al. 2020a) utilizes a decay factor to reduce gradients from the residual modules and encourages the attack to focus on more transferable low-level information. The Interaction-Reduced (IR) attack (Wang et al. 2020) discovers the negative correlation between adversarial transferability and the interactions inside perturbations. Based on the findings of IR, Wang et al. propose to decrease interactions inside perturbations, which can be implicitly represented by Shapley values. Nevertheless, the above approaches are designed for CNNs, and the performance degrades significantly when the attacks are directly applied to ViTs due to key architectural differences.

Vision Transformers

Inspired by the great success of transformers on Natural Language Processing, the vision transformer (ViT) was first

introduced by (Dosovitskiy et al. 2020). This model receives raw image patches as input and is pre-trained with a large image dataset. Subsequently, many works have proposed to further improve the accuracy and efficiency of ViTs. In order to overcome the necessity of pre-training ViTs on massive datasets, DeiT (Touvron et al. 2021a) introduces a teacher-student strategy specific to transformers which utilizes a new distillation token to learn knowledge from CNNs. T2T-ViT (Yuan et al. 2021) introduces a T2T module to model the local structure of the image and adopts a deep-narrow structure for the backbone of Transformers. TNT (Han et al. 2021) utilizes an outer transformer block and an inner transformer block to learn relationships between patches and within patches separately. CaiT (Touvron et al. 2021b) builds deeper transformers and inserts class tokens only in later layers. CvT (Wu et al. 2021) introduces convolutions into ViTs with the benefits of CNNs, which inserts the class token only in the last layer. Other ViTs such as LeViT (Graham et al. 2021), PiT (Heo et al. 2021), ConViT (d’Ascoli et al. 2021), Visformer (Chen et al. 2021c), M2TR (Wang et al. 2021b) and Deepvit (Zhou et al. 2021) further improve ViTs from different angles as well. Apart from the above works, there are also some other papers (Tang et al. 2021; Paul and Chen 2021; Shi and Han 2021) that discuss the robustness of ViTs.

Transfer-based Attacks on ViTs

Compared to transfer-based attacks on CNNs, less effort has been made in investigating the transferability of adversarial examples across different ViTs. One related work (Naseer et al. 2021) proposes the Self-Ensemble (SE) method, which boosts adversarial transferability by optimizing perturbations on an ensemble of models. This method utilizes the class token at each layer with a shared classification head to build an ensemble of models. As the perturbations are optimized over an ensemble of models, the generated example remains adversarial with respect to multiple models, and this property may encourage generalization to other black-box models. Additionally, (Naseer et al. 2021) also introduces a Token Refinement (TR) module to fine-tune the class tokens for further enhancing transferability. Despite exhibiting promising performance, the broad applicability of this method is limited since many ViT models do not have enough class tokens to build an ensemble. Furthermore, TR needs to access the ImageNet training set during a fine-tuning process, which is time-consuming. In contrast, our method is generalizable across different ViTs models and is easy to implement.

Methodology

Preliminary

Consider an image sample $x \in \mathcal{X} \subset \mathbb{R}^{H \times W \times C}$ and its ground-truth label $y \in \mathcal{Y} = \{1, \dots, K\}$, where H , W , C denote height, width, and the number of channels respectively, and K represents the number of classes. We reshape the image x into a sequence of flattened patches $x_p = \{x_p^1, \dots, x_p^i, \dots, x_p^N\} \in \mathbb{R}^{N \times (P^2 \cdot C)}$, where x_p^i denotes the i -th patch of x , (P, P) is the resolution of x_p^i , and

$N = H \cdot W/P^2$ is the number of patches. Then, x_p is the input to a vision transformer. We use $f(x) : \mathcal{X} \rightarrow \mathcal{Y}$ to denote the prediction function for a white-box ViT surrogate model. We use g to denote a black-box victim model, which can be a vanilla CNN, robustly trained CNN, and ViT. Following (Dong et al. 2018; Xie et al. 2019; Wu et al. 2020a), we focus on untargeted adversarial attacks and enforce an L_∞ -norm constraint on perturbations. Hence, the goal of the transfer-based attack is to add a perturbation δ to x yielding an adversarial example x_{adv} using information from f in order to maximize the probability that $g(x_{adv}) \neq y$ subject to the constraint that $\|\delta\|_\infty \leq \epsilon$. The optimization problem for generating adversarial examples on white-box models is formulated as follows:

$$\arg \max_{\delta} J(f(x + \delta), y), \text{ s.t. } \|\delta\|_\infty < \epsilon, \quad (1)$$

where $J(\cdot, \cdot)$ is the loss function (e.g. cross-entropy).

Pay No Attention (PNA)

In ViTs, a single MSA uses multiple self-attention structures, where each such structure, *i.e.* a head, learns its own unique features. By projecting concatenated outputs from multiple heads, MSA combines information from different representation subspaces (Vaswani et al. 2017). Based on the observation that the gradients of attention in each head impair the generation of highly transferable adversarial examples, we drop the attention gradients to alleviate the interaction between patches.

Figure 1(a) illustrates the proposed Pay No Attention (PNA) attack. Given an input patch embedding $Z \in \mathbb{R}^{N \times D}$, query, key, and value weights $W^Q, W^K, W^V \in \mathbb{R}^{D \times D_h}$, the attention can be formulated as follows:

$$A = \text{softmax}(ZW^Q(ZW^K)^T / \sqrt{D_h}), \quad (2)$$

where $A \in \mathbb{R}^{N \times N}$ denotes the attention weight. Then the output of this head can be defined as follows:

$$Z' = A(ZW^V). \quad (3)$$

Thus the gradient of the output Z' with respect to input Z can be decomposed as:

$$\frac{\partial Z'}{\partial Z} = (\mathbb{I} \otimes A) \frac{\partial(ZW^V)}{\partial Z} + ((ZW^V)^T \otimes \mathbb{I}) \frac{\partial A}{\partial Z}. \quad (4)$$

Our proposed PNA attack ignores backpropagation through the attention branch, *i.e.*, it sets $\frac{\partial A}{\partial Z} = 0$. This is equivalent to treating the attention weights as a constant, or using a ‘‘stop gradient’’ on the attention weights. This results in the approximation

$$\frac{\partial Z'}{\partial Z} \approx (\mathbb{I} \otimes A) \frac{\partial(ZW^V)}{\partial Z} = (\mathbb{I} \otimes A)((W^V)^T \otimes \mathbb{I}), \quad (5)$$

where \otimes denotes the Kronecker product. PNA forces the perturbation to exploit the network only by using feature representations, and not by exploited highly model-specific properties of attention. This results in adversarial examples with high transferability. Skipping attention also allows gradients to focus on each patch individually, rather than relying on complex interactions between patches. IR (Wang

Algorithm 1: The dual attack on ViTs

Input: The loss function J of Equation 7, a white-box model f , a clean image x with its ground-truth class y .

Parameter: The perturbation budget ϵ , iteration number I , used patch number T .

Output: The adversarial example.

```

1:  $\delta_0 \leftarrow \mathbf{0}$ 
2:  $\alpha \leftarrow \frac{\epsilon}{I}$ 
3: for  $i = 0$  to  $I - 1$  do
4:    $x_s \leftarrow \text{PatchOut}(x_p, T)$ 
5:    $M \leftarrow \text{Equation 6}$ 
6:    $g \leftarrow \text{PNA}(\nabla_{\delta} J \text{ with the } L_2 \text{ norm})$ 
7:    $\delta_i \leftarrow \text{clip}_{\epsilon}(\delta_{i-1} + \alpha \cdot g)$ 
8: end for
9:  $x_{adv} = x + \delta_I$ 
10: return  $x_{adv}$ 

```

et al. 2020) showed a negative correlation between adversarial transferability and these multi-patch interactions. We believe this is another reason why PNA improves transferability.

PatchOut Attack

To motivate our work, we consider the findings of (Xie et al. 2019), which suggests that diverse input patterns can improve the transferability of adversarial examples by alleviating the over-fitting phenomenon. Therefore, we introduce the PatchOut Attack, which randomly attacks a subset of patches at each iteration to alleviate over-fitting.

We use T to control the number of used patches at each iteration, and $x_s = \{x_s^1, \dots, x_s^t, \dots, x_s^T\}$ to denote the selected patches. Thus, the attack mask $M \in \{0, 1\}^{H \times W \times C}$ can be formulated as:

$$M_p^i = \begin{cases} 1, & \text{if } x_p^i \text{ in } x_s \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

where $M_p^i \in \{0, 1\}^{P \times P \times C}$ is the region corresponding to x_p^i in the image. Therefore, PatchOut replaces Equation 1 with:

$$\arg \max_{\delta} J(f(x + M \odot \delta), y) + \lambda \|\delta\|_2, \text{ s.t. } \|\delta\|_\infty < \epsilon, \quad (7)$$

where \odot denotes element-wise multiplication. The added second term encourages perturbations to have a large L_2 norm, preferring a large distance from x . λ controls the balance between the loss function and the regularization term.

We finally summarize the proposed dual attack to craft adversarial examples in Algorithm 1, where the function $\text{PatchOut}(\cdot, \cdot)$ randomly select T patches from x_p to generate x_s , the function $\text{PNA}(\cdot)$ bypasses the gradients of attention, and $\text{clip}_{\epsilon}(\cdot)$ restricts each entry of generated perturbations to be within $[-\epsilon, \epsilon]$.

Experiments

Experimental Settings

Dataset Following the setting from Dong et al. (2018, 2019); Xie et al. (2019); Lin et al. (2019), we randomly sample one image, which is correctly classified by all models,

Method	ViT-B/16	PiT-B	CaiT-S-24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B
FGSM	15.57	19.80	20.43	19.37	22.08	22.78	18.80	25.58
BIM	20.77	22.17	22.63	22.70	33.53	32.13	20.45	35.30
MI	41.23	45.23	47.13	45.97	56.03	55.23	43.75	58.25
DI	32.57	45.13	43.07	47.77	48.08	55.18	43.25	49.35
TI	19.33	17.67	16.50	19.00	25.13	28.18	13.70	27.53
SIM	34.97	32.73	35.17	31.13	44.13	46.73	36.43	45.68
SGM	38.87	41.60	52.30	48.80	60.53	64.33	51.13	60.68
IR	21.33	22.70	24.00	23.43	34.00	33.43	21.30	36.38
TAP	25.27	24.73	33.40	32.20	43.20	39.78	30.03	42.20
ATA	3.47	1.13	0.97	2.67	3.68	3.37	2.02	3.72
SE	29.05	21.25	31.40	24.90	45.23	37.87	21.73	46.03
Ours	46.10	52.40	59.87	58.60	63.85	67.25	57.62	63.70

Table 1: MASR (%) against ViTs with various attack methods. We use ViT-B/16, PiT-B, CaiT-S-24 and Visformer-S as white-box models respectively to generate adversarial examples. Each model is evaluated on adversarial examples generated by white-box ViTs, and the ASRs are averaged over surrogate models to obtain MASR. The best results are in bold.

from each class from the ImageNet 2012 validation dataset (Russakovsky et al. 2015), to conduct our experiments.

Models We evaluate the performance and transferability of the proposed method under two different settings: 1) The surrogate and victim models are both ViTs; 2) The surrogate is a ViT, and the victim model is a CNN. In these experiments, we evaluate performance on multiple ViT variants as well as multiple CNNs architectures, including both normally trained models and robustly trained models. For ViTs, we conduct experiments on ViT-B/16 (Dosovitskiy et al. 2020), DeiT-B (Touvron et al. 2021a), TNT-S (Han et al. 2021), LeViT-256 (Graham et al. 2021), PiT-B (Heo et al. 2021), CaiT-S-24 (Touvron et al. 2021b), ConViT-B (d’Ascoli et al. 2021), and Visformer-S (Chen et al. 2021c). While for CNNs, the experiments are conducted on Inception v3 (Inc-v3) (Szegedy et al. 2016), Inception v4 (Inc-v4), Inception ResNet v2 (IncRes-v2) (Szegedy et al. 2017), and ResNet v2-152 (Res-v2) (He et al. 2016). We chose these models to conduct our experiments since they are publicly available in the timm library (Wightman 2019). We randomly select ViT-B/16, PiT-B, CaiT-S-24, Visformer-S as the white-box models to generate adversarial examples. When selecting one model as the white-box surrogate, we use the remaining models as black-box victims when evaluating performance.

Evaluation We use Attack Success Rate (ASR) to quantify performance on black-box models. Here, ASR denotes the proportion of generated adversarial examples from the surrogate model that are successfully misclassified by the black-box victim model. A higher success rate means better adversarial transferability. We also use Mean ASR (MASR) to denote the mean value of ASR for one black-box victim, averages across different white-box surrogate models. Following Dong et al. (2018, 2019); Xie et al. (2019), we set the norm constraint $\epsilon = 16$ and J as the cross-entropy loss function. For the iterative attack, we set $I = 10$ and thus the step size $\alpha = 1.6$. We resize all images to 224×224 to conduct experiments. For the inputs of ViTs, we set the patch

size $P = 16$, thus the number of the patches is $N = 196$.

Performance Comparison

We compare our method with several baseline attacks, including MI (Dong et al. 2018), DI (Xie et al. 2019), TI (Dong et al. 2019), SIM (Lin et al. 2019), SGM (Wu et al. 2020a), IR (Wang et al. 2020), TAP (Zhou et al. 2018), ATA (Wu et al. 2020b), and SE (Naseer et al. 2021), which are integrated into BIM (Kurakin, Goodfellow, and Bengio 2016). Note that we do not compare with TR (Naseer et al. 2021) because it requires fine-tuning on ImageNet, and thus would yield an unfair comparison in terms of both data access and compute cost. For each baseline method, we follow their original settings in our experiments. We set our patch number $T = 130$, and we set the balancing parameter $\lambda = 0.1$ in PatchOut according to the experimental results. For performance comparison, we report the MASR obtained from white-box models including ViT-B/16, PiT-B, CaiT-S-24 and Visformer-S.

Performance on ViTs We first evaluate the adversarial transferability of our method across different ViTs. Table 1 summarizes the results on different black-box ViT models. From the results, we have the following observations. First, TI and ATA, which improve the adversarial transferability significantly on CNNs, exhibit poor performances on ViTs. As TI is proposed based on the property of translation invariance in CNNs, it is not applicable to ViTs since the architecture of ViTs leads to much less image-specific inductive bias than CNNs (e.g., translation invariance) (Yuan et al. 2021). Furthermore, while the key idea of ATA is to corrupt the features of discriminative regions shared among different CNNs, it is also not applicable to ViTs for the reason that the discriminative region varies significantly across different ViTs. Second, compared to other baseline methods designed for CNNs, SGM performs much better on ViTs. This is because the key idea of SGM is to boost transferability by encouraging the model to focus more on low-level information, which is also applicable to ViTs since SGM is

Method	Inc-v3	Inc-v4	IncRes-v2	Res-v2	Inc-v3 _{ens3}	Inc-v3 _{ens4}	IncRes-v2 _{ens}
FGSM	20.78	18.80	16.38	19.05	14.62	13.98	10.38
BIM	17.88	14.77	12.40	11.82	8.02	6.20	4.45
MI	39.65	37.43	32.17	33.28	24.80	22.10	17.68
DI	32.78	31.75	26.40	25.00	17.73	15.22	11.12
TI	23.27	23.60	15.28	20.88	18.80	20.80	13.92
SIM	30.55	27.63	24.17	24.48	19.45	17.67	13.40
SGM	38.42	34.00	27.25	27.25	18.70	16.53	11.68
IR	17.65	15.83	12.08	12.48	8.05	6.50	4.78
TAP	29.58	26.10	20.67	19.23	12.58	10.62	6.85
ATA	3.25	2.53	2.03	2.07	1.20	0.85	0.92
SE	18.40	16.47	12.33	12.63	9.13	6.73	5.10
Ours	47.95	45.12	38.45	38.93	26.20	22.85	18.10

Table 2: MASR (%) against CNNs with various attack methods. We use ViT-B/16, PiT-B, CaiT-S-24 and Visformer-S as white-box models to generate adversarial examples. Each CNN is evaluated on adversarial examples generated by white-box ViTs, and the ASRs are averaged over surrogate models to obtain MASR. The best results are in bold.

not designed based on the unique characteristics of CNNs. Third, SE, which is proposed for improving the transferability of ViTs, does not work well across different structures of ViTs. This is because the performance of SE is closely related to the number of class tokens for building the ensemble models. As there are few or even no class tokens in some ViTs (e.g. PiT-B, Visformer-S), SE works poorly for such situations. Lastly, our dual attack consistently outperforms the baseline attacking methods and achieves a 58.67 % average MASR by averaging MASR across different black-box models. Our method, which considers the core components of ViTs, achieves the best results in ViTs with different structures. These experiments validate the effectiveness of the proposed dual attack.

Performance on CNNs We further evaluate the performance of transferring the adversarial examples generated on white-box ViTs to attack black-box CNNs. The results are summarized in Table 2. We have the following observations. First, the attack success rates of all methods decrease significantly, which indicates that adversarial samples generated on ViTs are less transferable on CNNs than ViTs, presumably due to structural differences between ViTs and CNNs. Second, MI improves the transferability of the generated adversarial examples by utilizing the momentum term to stabilize update directions. Third, compared to other baseline methods, our proposed dual attack significantly improves MASR despite the different structures between ViTs and CNNs. The dual attack has a 42.61% average MASR on normally trained CNNs and a 22.38% average MASR on robustly trained CNNs. These results indicate the feasibility of using ViTs to attack robustly trained CNNs.

Combining with Existing Methods The proposed dual attack tailored for ViTs can be easily combined with existing methods to further improve performance. We demonstrate this by combining our method with MI and SGM. The reason we choose MI and SGM is that these two methods perform much better than other baseline methods. Using these combined methods, we conduct experiments with ViT-



Figure 3: Visualization of randomly picked clean images and their corresponding adversarial images, crafted on the ViT-B/16 model using our dual attack.

B/16 as the white-box model. Specifically, SGM + Ours uses fewer gradients than MultiLayer Perceptron (MLP) modules according to a decay factor that is set as 0.5. Table 3 shows the results of the integrated methods. We observe that the proposed dual attack significantly improves the transferability of MI and SGM across ViTs, normally trained CNNs, and robustly trained CNNs. In particular, the dual attack increases the average MASR of MI and SGM by 28.54 % and 102.49 %, respectively, when attacking robustly trained CNNs. The results demonstrate that the dual attack can easily be combined with existing methods to further boost performance.

Visualization of Adversarial Examples Figure 3 depicts 8 randomly selected clean images and their corresponding adversarial examples. Here, we observe that the adversarial perturbations are hardly perceptible.

Method	ViTs	Normal CNNs	Robust CNNs
MI	49.10	35.63	21.53
MI + Ours	59.70	48.79	30.13
SGM	52.28	31.73	15.64
SGM + Ours	63.48	52.31	31.37

Table 3: Average ASR (%) of models under attacks that combine the proposed method and existing algorithms. The experiment is conducted by using ViT-B/16 as the white-box model. Average ASR is obtained by averaging ASRs of different types of black-box models. "Normal CNNs" and "Robust CNNs" denote normally trained CNNs and robustly trained CNNs respectively.

PatchOut	L_2	PNA	ViTs	Normal	Robust
-	-	-	26.97	10.63	5.50
✓	-	-	34.92	12.55	7.07
-	✓	-	37.79	16.55	10.00
-	-	✓	42.47	18.75	10.97
✓	✓	-	48.82	24.55	14.67
-	✓	✓	49.69	22.36	13.54
✓	✓	✓	59.15	36.23	23.87


Table 4: Average ASR (%) for our proposed method with different component combinations. PatchOut indicates the operation of sampling a subset of patches, L_2 denotes the regularization term in Equation 7, PNA is the operation that bypasses the gradients of attention. '✓' indicates that the component is used while '-' indicates that it is not used. "Normal" and "Robust" denote normally trained CNNs and robustly trained CNNs respectively.

Ablation Studies


To validate the effect of each component proposed in our method, we evaluate the performance under various combinations of the components: PatchOut, the regularization term L_2 in Equation 7, and Pay No Attention (PNA). We also conduct experiments to analyze the effect of hyper-parameters, including the used patch number T and the balance parameter λ . For the ablation study, we use ViT-B/16 as the white-box model and generate adversarial examples on 2,000 randomly sampled images to attack other black-box models.

Effect of Components Table 4 shows the results of different component combinations. It is observed that using any of the components leads to performance improvements on both ViTs and CNNs. Among these components, PNA improves transferability most significantly on both ViTs and CNNs, demonstrating that gradients from ViTs, except for those from attention modules, are most vulnerable. The best performance is achieved when combining all components together, which suggests that three components contribute to the improvement of transferability in complementary ways.

Effect of Parameters We investigate the effect of parameters T and λ and summarize the results in Figure 4. The used patch number T determines how many patches



(a) The used patch number T



(b) λ

Figure 4: Average ASR (%) of ViTs under attacks using various values of (a) patch number T and (b) λ in Equation 7.

are used in each attack iteration. When all patches are included ($T = N$), PatchOut degenerates into BIM. When T is small, PatchOut is not able to generate strong adversarial examples in a limited number of iterations ($I = 10$). Thus, it is vital to study optimal settings of T . We experiment by tuning T without L_2 . Figure 4(a) shows the results, with $T \in [10, 40, 70, 100, 130, 160, 196]$. When $T = 130$, PatchOut achieves optimal results. With this value of T , each patch can be attacked multiple times in the context of different patch subsets. For the balance parameter λ , we set $\lambda \in [0.001, 0.01, 0.1, 1, 10]$ to conduct the experiments with $T = 130$. Figure 4(b) illustrates the effect of various λ . When $\lambda = 0.1$, PatchOut achieves the best result by balancing the contribution of each term in Equation 7.

Conclusion

In this paper, we identify several properties of ViTs that can be leveraged to produce more transferable adversarial examples. Specifically, we find that ignoring the gradients of attention units and only perturbing a subset of the patches at each iteration prevents overfitting and creates diverse input patterns, thus increasing transferability. We conduct two toy experiments to validate these assumptions and propose the dual attack for ViTs consisting of the Pay No Attention (PNA) attack and the PatchOut attack. We conduct a series of experiments to show that the proposed method can greatly improve transferability. Combining our attack with other existing methods, our techniques consistently enhance transferability. The results of this study show that cross-model transferability happens even across very wide architectural gaps between models. The transferability observed in this work further suggests that the feature extractors and implicit biases of transformer architectures and CNNs are not as different as one might suspect, and we think adversarial examples might serve as the basis for future studies of the similarities and differences between the two. Code is available at <https://github.com/zhipeng-wei/PNA-PatchOut>.

Acknowledgments

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. This work was supported in part by NSFC project (#62032006), Science and Technology Commission of Shanghai Municipality Project (20511101000), and in part by Shanghai Pujiang Program (20PJ1401900).

References

- Breiman, L. 2004. Random Forests. *Machine Learning*, 45: 5–32.
- Chen, K.; Wei, Z.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2021a. Attacking Video Recognition Models with Bullet-Screen Comments. *arXiv preprint arXiv:2110.15629*.
- Chen, S.; Tao, Q.; Ye, Z.; and Huang, X. 2021b. Going Far Boosts Attack Transferability, but Do Not Do It. *arXiv preprint arXiv:2102.10343*.
- Chen, Z.; Xie, L.; Niu, J.; Liu, X.; Wei, L.; and Tian, Q. 2021c. Visformer: The vision-friendly transformer. *arXiv preprint arXiv:2104.12533*.
- d’Ascoli, S.; Touvron, H.; Leavitt, M.; Morcos, A.; Biroli, G.; and Sagun, L. 2021. Convit: Improving vision transformers with soft convolutional inductive biases. *arXiv preprint arXiv:2103.10697*.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9185–9193.
- Dong, Y.; Pang, T.; Su, H.; and Zhu, J. 2019. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4312–4321.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Graham, B.; El-Nouby, A.; Touvron, H.; Stock, P.; Joulin, A.; Jégou, H.; and Douze, M. 2021. LeViT: a Vision Transformer in ConvNet’s Clothing for Faster Inference. *arXiv preprint arXiv:2104.01136*.
- Han, K.; Xiao, A.; Wu, E.; Guo, J.; Xu, C.; and Wang, Y. 2021. Transformer in transformer. *arXiv preprint arXiv:2103.00112*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*, 630–645. Springer.
- Heo, B.; Yun, S.; Han, D.; Chun, S.; Choe, J.; and Oh, S. J. 2021. Rethinking spatial dimensions of vision transformers. *arXiv preprint arXiv:2103.16302*.
- Hinton, G. E.; Srivastava, N.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *ArXiv*, abs/1207.0580.
- Kurakin, A.; Goodfellow, I.; and Bengio, S. 2016. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2017. Adversarial examples in the physical world. *ArXiv*, abs/1607.02533.
- Lin, J.; Song, C.; He, K.; Wang, L.; and Hopcroft, J. E. 2019. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*.
- Liu, Y.; Chen, X.; Liu, C.; and Song, D. 2016. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*.
- Naseer, M.; Ranasinghe, K.; Khan, S.; Khan, F. S.; and Porikli, F. 2021. On Improving Adversarial Transferability of Vision Transformers. *arXiv preprint arXiv:2106.04169*.
- Paul, S.; and Chen, P.-Y. 2021. Vision transformers are robust learners. *arXiv preprint arXiv:2105.07581*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Shapley, L. S. 1988. *A value for n-person games*, 31–40. Cambridge University Press.
- Shi, Y.; and Han, Y. 2021. Decision-based Black-box Attack Against Vision Transformers via Patch-wise Adversarial Removal. *ArXiv*, abs/2112.03492.
- Szegedy, C.; Ioffe, S.; Vanhoucke, V.; and Alemi, A. A. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826.
- Tang, S.; Gong, R.; Wang, Y.; Liu, A.; Wang, J.; Chen, X.; Yu, F.; Liu, X.; Song, D.; Yuille, A.; et al. 2021. Robustart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021a. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, 10347–10357. PMLR.
- Touvron, H.; Cord, M.; Sablayrolles, A.; Synnaeve, G.; and Jégou, H. 2021b. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*.

Vaswani, A.; Shazeer, N. M.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. *ArXiv*, abs/1706.03762.

Wang, J.; Liu, A.; Yin, Z.; Liu, S.; Tang, S.; and Liu, X. 2021a. Dual Attention Suppression Attack: Generate Adversarial Camouflage in Physical World. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Wang, J.; Wu, Z.; Chen, J.; and Jiang, Y.-G. 2021b. M2TR: Multi-modal Multi-scale Transformers for Deepfake Detection. *arXiv preprint arXiv:2104.09770*.

Wang, X.; Ren, J.; Lin, S.; Zhu, X.; Wang, Y.; and Zhang, Q. 2020. A unified approach to interpreting and boosting adversarial transferability. *arXiv preprint arXiv:2010.04055*.

Wei, Z.; Chen, J.; Wei, X.; Jiang, L.; Chua, T.-S.; Zhou, F.; and Jiang, Y.-G. 2020. Heuristic black-box adversarial attacks on video recognition models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12338–12345.

Wei, Z.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2021. Cross-Modal Transferable Adversarial Attacks from Images to Videos. *arXiv:2112.05379*.

Wightman, R. 2019. PyTorch Image Models. <https://github.com/rwightman/pytorch-image-models>. Accessed: 2021-08-20.

Wu, D.; Wang, Y.; Xia, S.-T.; Bailey, J.; and Ma, X. 2020a. Skip connections matter: On the transferability of adversarial examples generated with resnets. *arXiv preprint arXiv:2002.05990*.

Wu, H.; Xiao, B.; Codella, N.; Liu, M.; Dai, X.; Yuan, L.; and Zhang, L. 2021. Cvt: Introducing convolutions to vision transformers. *arXiv preprint arXiv:2103.15808*.

Wu, W.; Su, Y.; Chen, X.; Zhao, S.; King, I.; Lyu, M. R.; and Tai, Y.-W. 2020b. Boosting the transferability of adversarial samples via attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1161–1170.

Xie, C.; Zhang, Z.; Zhou, Y.; Bai, S.; Wang, J.; Ren, Z.; and Yuille, A. L. 2019. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2730–2739.

Yan, H.; Wei, X.; and Li, B. 2020. Sparse black-box video attack with reinforcement learning. *arXiv preprint arXiv:2001.03754*.

Yuan, L.; Chen, Y.; Wang, T.; Yu, W.; Shi, Y.; Jiang, Z.; Tay, F. E.; Feng, J.; and Yan, S. 2021. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*.

Zhou, D.; Kang, B.; Jin, X.; Yang, L.; Lian, X.; Jiang, Z.; Hou, Q.; and Feng, J. 2021. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*.

Zhou, W.; Hou, X.; Chen, Y.; Tang, M.; Huang, X.; Gan, X.; and Yang, Y. 2018. Transferable adversarial perturbations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 452–467.