

Adversarial Pixel Restoration as a Pretext Task for Transferable Perturbations

Hashmat Shadab Malik¹
hashmat.malik@mbzuai.ac.ae

Shahina K Kunhimon¹
shahina.kunhimon@mbzuai.ac.ae

Muzammal Naseer^{1,2}
muzammal.naseer@mbzuai.ac.ae

Salman Khan^{1,2}
salman.khan@mbzuai.ac.ae

Fahad Shahbaz Khan^{1,3}
fahad.khan@mbzuai.ac.ae

¹ Mohamed bin Zayed University of AI

² Australian National University

³ Linköping University

Abstract

Transferable adversarial attacks optimize adversaries from a pretrained surrogate model and known label space to fool the unknown black-box models. Therefore, these attacks are restricted by the availability of an effective surrogate model. In this work, we relax this assumption and propose *Adversarial Pixel Restoration* as a self-supervised alternative to train an effective surrogate model from scratch under the condition of *no* labels and *few* data samples. Our training approach is based on a min-max scheme which reduces overfitting via an adversarial objective and thus optimizes for a more generalizable surrogate model. Our proposed attack is complimentary to the adversarial pixel restoration and is independent of any task specific objective as it can be launched in a self-supervised manner. We successfully demonstrate the adversarial transferability of our approach to Vision Transformers as well as Convolutional Neural Networks for the tasks of classification, object detection, and video segmentation. Our training approach improves the transferability of the baseline unsupervised training method by 16.4% on ImageNet val. set. Our codes & pre-trained surrogate models are available at: <https://github.com/HashmatShadab/APR>.

1 Introduction

Adversarial attacks [0, 6, 16, 28, 29, 33, 39] add small, imperceptible but well optimized noise to the clean image which can elicit an incorrect decision from the model. These attack methods [0, 24, 21, 31, 53, 54] craft adversarial examples that can broadly be categorized based on how much information is available about the target model. In a *white-box* attack setting, the attacker has complete knowledge of the target model and can directly optimize adversarial perturbations for the given model. In a more realistic *black-box* attack setting, the attacker does not have access to the target model, its architectural details or targeted task

(*e.g.*, classification, segmentation or object detection). In such a case, adversarial examples are created on a surrogate model and then transferred to the black-box model. Adversarial examples generated from surrogate models trained on a large-scale dataset in a supervised manner have better transferability [12, 81]. Transferability of such attacks improves further by fine-tuning the surrogate model to enhance their representation capacity *e.g.*, by finding better self-ensemble from a given pretrained model [65]. The adversarial transferability hence depends on the generalizability of the surrogate model. Such attacks are also restricted by the availability of a pretrained surrogate and information about the label space.

In this work, we call into question this assumption and consider a stronger threat model where an attack is launched from few unannotated or cross-domain samples (*e.g.*, painting to ImageNet samples) without any knowledge of target networks or tasks. This threat model poses a challenge on how to learn an effective surrogate model from the limited unlabelled data and then how to generate self-supervised transferable adversarial examples.

With limited availability of data, neural networks can easily memorize the data [56] even with strong augmentations [26]. Therefore, Li *et al.* [26] propose to reconstruct transformed input images to learn a surrogate model. However, their effective surrogate training and attack approach requires supervision though annotated data samples. In order to reduce overfitting over few data samples and to find robust features, we take inspiration from adversarial training [16, 47] and propose self-supervised Adversarial Pixel Restoration to train a surrogate model. The min-max objective of our proposed training allows to find a flatter minima with robust features which compliments our self-supervised adversarial attack to achieve higher adversarial transferability. Our main contributions are as follows:

- We propose self-supervised Adversarial Pixel Restoration to find highly transferable patterns by learning over flatter loss surfaces. Our training approach allows launching cross-domain attacks without access to large-scale labeled data or pretrained models.
- Our adversarial attack is self-supervised in nature and independent of any task-specific objective. For instance, our approach optimizes the robust transformed loss surface of the surrogate via fooling its reconstruction ability. This allows to generate task independent adversaries. Therefore, our approach can transfer perturbations to a variety of tasks as we demonstrate for classification, object detection, and segmentation.
- We provide a thorough analysis of our proposed method to establish its effectiveness. We observe that our approach leads to smoother loss landscapes (Fig. 2), helping in crafting more generalizable adversarial examples. Our method remains effective even in extreme data scarcity *e.g.* when trained on two data samples only. (see Sec. 4.2).

2 Related Work

Several gradient-based methods [16, 23, 24, 52] have been proposed for crafting adversarial examples directly on the target classifiers. However, when the access to the target model is limited to just a finite amount of queries, current methods either rely on the transferability of surrogate models [57, 59, 57] or estimate the gradients/boundary of the target model [8, 9, 60]. Both of the above approaches either require a non-trivial number of queries from the target model or access to the training distribution, making it highly impractical in real-case scenarios. A practical threat model was introduced in [60], where the attacker does not have access to the training distribution of the target model, as well as querying is

prohibited. Authors train a generator-based surrogate model with the help of training data and a pretrained classifier obtained from a different domain than the target models. However, the pretrained classifier as well as the generator are trained on a large annotated dataset.

In [26], a stronger threat model with access to limited data samples (order of tens) was proposed. Inspired from self-supervised learning methods, autoencoder-based surrogate models are trained with limited data. However, the transferability of the autoencoders trained in an unsupervised manner is still moderate. Most of the previous works have explored the transferability of surrogate models trained on the target model’s training set. In [48], the authors observe that the features of robust classifiers can be used to generate adversarial examples that are more transferable. Building on this, [47] observe that having a classifier adversarially trained with a small perturbation budget (“slightly robust”), leads to highly transferable adversarial examples. Unlike prior works, we focus on constructing robust surrogate models in a fully unsupervised manner. We consider the case of training robust model with limited data to improve adversarial transferability. Furthermore, we also consider a practical scenario of availability of large unlabelled dataset for training of robust models which can be used for constructing cross-domain and cross-task adversaries.

3 Adversarial Pixel Restoration for Transferable Perturbations

Our goal is to learn a surrogate model, \mathcal{F} , from a given unlabelled data distribution, P_s , with a set of only few image samples (≤ 20). This setting is in contrast to the existing transferable adversarial attacks [12, 33, 35, 54] that assume access to a surrogate model trained on a large-scale annotated data (e.g., ImageNet [10]) in a supervised fashion. In the presence of unlabeled data, however, the surrogate model can be trained by defining a self-supervised task, \mathcal{T}_s , such as predicting rotation [15], solving jigsaw puzzle [36] or by matching different views of the same input sample [8]. A major challenge is that deep neural networks can easily memorise the data [56] and quickly overfit the few available samples even after applying strong data augmentation techniques [26]. This results in a surrogate model with less generalizable representations and consequently, adversarial attack launched from such a model has weak transferability (Sec. 4.1.1).

We propose adversarial pixel restoration as a prior to train the surrogate model \mathcal{F} that boosts transferability of the adversarial attacks (see Fig. 1). We create adversarial examples from the original input images by attacking the source model in the pixel space using our proposed adversarial attack (Sec. 3.2.1). Our approach also shifts the position of input pixels via transformations such as rotation or jigsaw shuffle. We then train a model to denoise and restore pixels to their original positions via our proposed adversarial training (Algo. 1). The surrogate model \mathcal{F} in our case consists of an autoencoder as explained below.

3.1 Surrogate Architecture

The surrogate model is based on an autoencoder [26]. The encoder consists of a stack of convolution layers. At the beginning of the architecture, convolution layers with a larger kernel size and stride help reduce the spatial resolution of the feature maps, followed by multiple residual blocks where the size of the feature maps is kept constant. The decoder is a lightweight model consisting of two transpose convolution layers to upsample the fea-

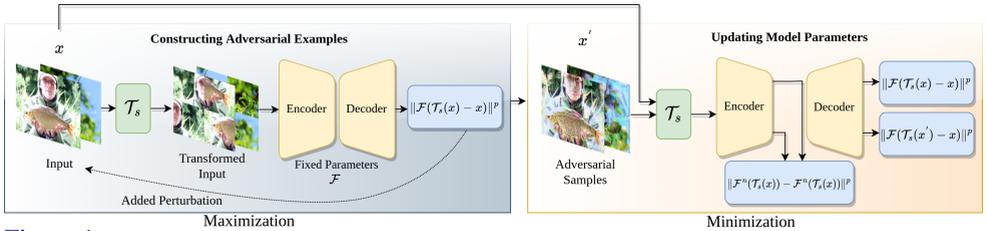


Figure 1: Our approach trains an autoencoder based surrogate model via self-supervised adversarial pixel restoration to learn generalizable representations from a limited number of data samples (≤ 20). Our training is based on a min-max strategy. We first generate adversarial examples by fooling model’s reconstructive ability (*maximization*), followed by updating the model parameters based on restoration of the transformed adversarial and clean sample (*minimization*). Our approach allows launching transferable self-supervised adversarial attacks without any knowledge of target (black-box) model.

ture maps. The adversarial training of these autoencoders is based on denoising and pixel restoration as explained next in Sec. 3.2.

3.2 Adversarial Training via Denoising

The feature space of a slightly robust classifier produces highly transferable adversarial examples [14]. However, adversarial training of such models is computationally demanding due to iterative training and also requires a large-scale labeled dataset that might not be accessible to the attacker in real-world scenarios. We assume a more practical threat model, where the attacker has access to a data distribution P_S with limited number of samples without any annotations. We adversarially learn a surrogate model on the data distribution P_S to transfer adversarial perturbations to the target (black-box) models trained for different tasks (e.g., classifications, object detection and segmentation) on possibly different target distribution $Q_t \neq P_S$. In such challenging attack settings, we adversarially train an autoencoder \mathcal{F} via min-max training strategy [16]. At maximization step, we create adversarial examples by fooling the model \mathcal{F} via adversarial pixel transformations with a single step attack (Algo. 1). At minimization step, we denoise and restore the feature and pixel space to achieve generalizable loss surfaces (Fig. 2) which leads to more transferable attack.

3.2.1 Adversarial Pixel Transformations

For given input samples $\mathbf{x} \sim P_S$, we first find adversarial example, \mathbf{x}' subjected to: $\|\mathbf{x} - \mathbf{x}'\|_\infty \leq \epsilon$, by maximizing the following objective (\mathcal{L}_{max}):

$$\underset{\mathbf{x}'}{\text{maximize}} \quad \mathcal{L}_{max} = \|\mathcal{F}(\mathcal{T}_s(\mathbf{x}')) - \mathbf{x}\|^p, \quad (1)$$

where \mathcal{T}_s represents the pixel transformation (e.g., rotation or jigsaw shuffle) that shifts the pixel positions. Therefore, our attack fools the model’s ability to restore the transformed pixel space by maximizing the loss presented in Eq. 1 and ultimately help to robustify self-supervised features. Our attack approach can also be extended to benefit from supervisory signals e.g., by fooling prototypes [26, 46] as follows:

$$\underset{\mathbf{x}'}{\text{maximize}} \quad \sum_{c=1}^C \left(y_c \|\mathcal{F}(\mathcal{T}_s(\mathbf{x}')) - \mathbf{x}^{(c)}\|^p \right), \quad (2)$$

Algorithm 1 Adversarial Denoising and Restoration

Require: Source data distribution P_s , pixel transformation \mathcal{T}_s , attack step size δ , perceptual budget ε , balancing parameter λ , and maximum training iterations T

Ensure: Randomly initialize surrogate model \mathcal{F}

- 1: **for** $t \in [1, 2, \dots, T]$ **do**
- 2: Randomly sample from P_s : $\mathbf{x} \sim P_s$
- 3: Initialize adversary $\mathbf{x}' \leftarrow \mathbf{x}$
- 4: Optimize \mathbf{x}' using Eq. 1 or 2: ▷ Adversarial Pixel Transformation

$$\mathbf{x}' \leftarrow \mathbf{x} + \delta \times \text{sign}(\nabla \mathcal{L}_{max})$$

- 5: Project adversaries within allowed perceptual budget: $\mathbf{x}' \leftarrow \text{clip}(\mathbf{x}', \mathbf{x} - \varepsilon, \mathbf{x} + \varepsilon)$
- 6: Forward-pass \mathbf{x} and \mathbf{x}' through model \mathcal{F} and update its parameters θ by minimizing the loss given in Eq. 5: ▷ Pixel Restoration

$$\theta \leftarrow \theta - \alpha \times \nabla \mathcal{L}_{min},$$

where α is the learning rate.

7: **end for**

where C represents the number of categories, y_c represents one-hot encoded labels, and $\mathbf{x}^{(c)}$ is the chosen prototype for a particular class. Our proposed attack objective in Eq. 2 helps to optimize for robust discriminative features with better adversarial transferability.

3.2.2 Pixel Restoration

For a given adversarial sample, \mathbf{x}' created using Eqs. 1 or 2, we train the surrogate model, \mathcal{F} , by pixel restoration. Our loss function penalizes the model \mathcal{F} by minimizing the reconstruction error between the original sample, \mathbf{x} , and the model's output for the transformed adversarial as well as the transformed original sample as follows:

$$\mathcal{L}_{out} = \|\mathcal{F}(\mathcal{T}_s(\mathbf{x}')) - \mathbf{x}\|^p + \|\mathcal{F}(\mathcal{T}_s(\mathbf{x})) - \mathbf{x}\|^p. \quad (3)$$

We further regulate the model's feature space during adversarial training by enforcing alignment between the original and adversarial feature distributions as follows:

$$\mathcal{L}_{feature} = \|\mathcal{F}^n(\mathcal{T}_s(\mathbf{x}')) - \mathcal{F}^n(\mathcal{T}_s(\mathbf{x}))\|^p, \quad (4)$$

where \mathcal{F}^n represents the intermediate (encoder) layer output. Overall training objective is,

$$\mathcal{L}_{min} = \mathcal{L}_{out} + \lambda \mathcal{L}_{feature}, \quad (5)$$

where λ is the balancing parameter.

3.2.3 Behavior of Robust Loss Surfaces

We visualize the loss landscapes of our trained autoencoders (Fig. 2). We use the filter normalization method proposed in [25], which shows the structure of the loss surface along with random directions near the optimal pretrained parameters. We observe that our approach leads to more flatter minima as compared to the baseline method [26]. This has significant effect on finding generalizable adversarial example with better transferability (Sec. 4.1.1).

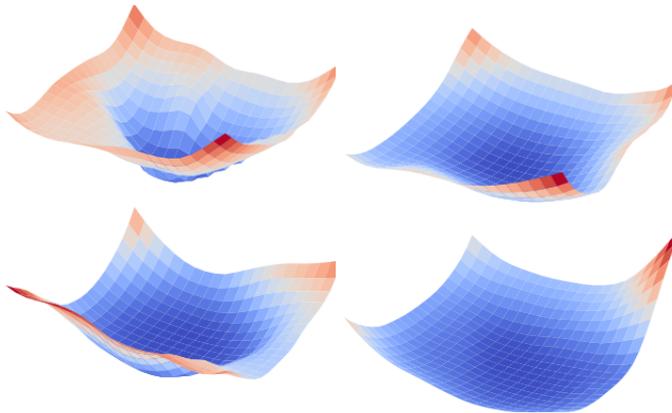


Figure 2: Reconstruction loss landscape of autoencoders. The first column shows the loss landscapes of autoencoders trained using the baseline method on the task of rotation (*top*) and jigsaw (*bottom*). The second column shows the corresponding flatter loss surfaces obtained by training the autoencoders using our approach.

4 Experimental Protocols

For a pixel range $[0, 1]$, we create l_∞ adversarial examples under perceptual budget of $\epsilon \leq 0.1$ following [26]. We show adversarial transferability of our approach against ImageNet trained models (Convolutional and ViTs), object detection and segmentation from in-domain and cross-domain (paintings, medical scans) data samples. We used Adam [22] optimizer with a learning rate of 0.001 for our proposed adversarial pixel restoration (Algo. 1) at $\epsilon \leq \frac{2}{255}$, $\delta = \frac{2}{255}$, $\lambda = 1$, $p = 2$. We provide detailed ablative analysis on the effect of these hyperparameters in Sec. 4.2.

Surrogate Training with Few Samples: Similar to [26], we assume access to only few data samples (*e.g.*, 20 samples to train a single autoencoder). We apply Eq. 1 as adversarial pixel transformation (\mathcal{T}) based on rotation or jigsaw in an unsupervised setting. We also incorporate our method in the supervised prototypical training of autoencoders mentioned in [26], where the reconstruction objective function used during the minimization and maximization step is formulated by Eq. 2. We compare the best prototypical setting reported in [26], comprising of multiple decoder networks. All models are trained for the same no. of iterations as [26]. We provide pseudo code of Algo. 1 in Appendix D.

Surrogate Training with Large Dataset: There is abundance of unannotated data available via online sources. Therefore, we also scale our proposed self-supervised adversarial training to large-scale datasets. Specifically, we train a surrogate model (single autoencoder) on paintings (79k samples) [9], CoCo (40k samples) [27], and Comics (50k samples) [2] to test unsupervised cross-domain adversarial transferability of our method.

Target Models: Adversarial perturbations from our trained autoencoders (surrogate models) are transferred to classification models including convolutional: VGG-19 [45], Inception-v3 (Inc-V3) [60], ResNet-152 (Res152) [47], Dense161 [49], SeNet [48], Wide-ResNet-50 (WRN) [65], and MobileNet-V2 (MNet-V2) [43], and Vision Transformers: ViT-T and ViT-S [43], DeiT-T and DeiT-S [61]. We also evaluated adversarial vulnerability of robust ResNet-50 models [42]. Further, we transfer attack to DETR [8] and DINO [8] to evaluate on object detection and video segmentation tasks.

Evaluation Metrics: We report drop in Top-1 (%) accuracy, Mean Average Precision (mAP), and Jaccard Index for classification, object detection and segmentation, respectively.

Datasets: We evaluate classification models on 5k samples from ImageNet validation set in the same setting as [26]. DETR and DINO are evaluated on CoCo (5k samples) and DAVIS

Transformation	Method	VGG-19	Inc-V3	Res152	Dense121	SeNet	WRN	MNet-V2	Average
Jigsaw	[26]	31.54	50.28	46.24	42.38	59.06	51.24	25.24	43.71
	Ours	16.82	25.54	31.18	22.64	38.06	25.76	13.70	24.81 _(-18.9)
Rotation	[26]	31.14	48.14	47.40	41.26	58.20	50.72	26.00	43.27
	Ours	19.02	25.76	33.60	25.60	38.92	29.78	15.38	26.87 _(-16.4)
Prototypical	[26]	18.74	33.68	34.72	26.06	42.36	33.14	16.34	29.29
	Ours	17.02	21.48	28.66	21.06	35.04	23.56	13.06	22.84 _(-6.45)

Table 1: Our attack consistently boost adversarial transferability across ImageNet trained models. Our training approach (Algo. 1) hence prove to be complimentary to autoencoders trained with different self-supervised (SS) tasks. Results (Top-1 (%), *lower is better*) are reported on 5k images from ImageNet validation set introduced by [26] under the same perceptual budget ($\epsilon \leq 0.1$).

Transformation	Method	Deit-T	Deit-S	ViT-T	ViT-S	Average
Jigsaw	[26]	45.3	62.42	36.62	62.1	51.61
	Ours	43.50	59.50	18.0	52.48	43.37 _(-8.24)
Rotation	[26]	46.1	62.0	37.8	60.38	51.57
	Ours	40.84	55.22	19.64	48.62	41.08 _(-10.49)
Prototypical	[26]	38.18	54.96	21.5	50.28	41.23
	Ours	34.16	51.54	16.74	45.3	36.94 _(-4.29)

Table 2: Comparative analysis of adversarial transferability for Vision Transformers models on ImageNet validation set. The Top-1 (%) under l_∞ bound $\epsilon = 0.1$ is shown (*lower is better*). Our method performs favorably well.

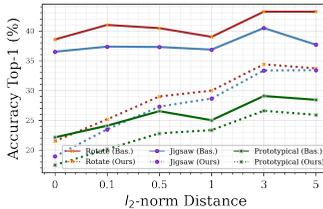
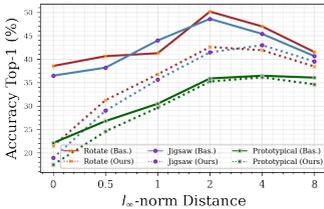


Figure 3: Models trained on l_∞ examples and large norm distance are less vulnerable to our attack, however, such models also lose accuracy on clean images [26].

(2k samples) validation sets respectively.

Baseline Adversarial Attack: For adversarial prototypical training, we use the same supervised attack objective as proposed by [26] for direct comparison. Specifically, we use 200 iterations of I-FGSM [24] followed by 100 iterations of ILA [20]. The attack objective for the surrogate models trained in self-supervised manner is simply based on maximizing the reconstruction error as described in Eq. 1.

4.1 Results

4.1.1 Transferability from Few In-Domain Samples

Our surrogate models trained on few data samples show significantly higher adversarial transferability as compared to the baseline [26]. Attack generated on our self-supervised models (rotation, and jigsaw) performs even better than supervised (prototypical) models of [26] (Tables 1, 2). Our approach further boosts the transferability rates when combined with supervised adversarial prototypical training. This compliments the benefits of our method in both supervised and self-supervised settings. We observe that vision transformers are more robust as compared to convolutional networks against such attacks [54], however, our approach provides non-trivial gains in fooling the Vision Transformers (Table 2). Similarly, adversarially robust models [22] are less vulnerable to such attacks (Fig. 3).

Transformation	Method (→)	[27]	Ours
Rotation	CoCo	28.56	23.31
	Paintings	27.83	17.75
	Comics	58.38	24.19
Jigsaw	CoCo	43.93	31.28
	Paintings	44.07	33.42
	Comics	67.70	41.54

Table 3: Adversarial perturbations are transferred from a single auto-encoder trained on CoCo, Paintings or Comics. We report average of Top-1 (%) accuracy against targeted convolutional networks (Table 1). Our proposed approach provide favorable results (*lower is better*).

Transformation (↓)	Dataset (→)	CoCo		Paintings		Comics	
		No Attack	[27] Ours	[27] Ours	[27] Ours	[27] Ours	
Rotation		39.7	19.3 14.6	17.2 11.9	34.3 13.3		
Jigsaw		39.7	24.7 14.5	24.1 14	38 20.8		

Table 4: Adversarial transferability to object detector (DETR) based on mAP at [0.5:0.95] is evaluated on CoCo validation set [27].

Transformation (↓)	Dataset (→)	CoCo		Paintings		Comics	
		No Attack	[27] Ours	[27] Ours	[27] Ours	[27] Ours	
Rotation		61.8	53.2 48.9	52.6 46.9	57.78 47.81		
Jigsaw		61.8	53.9 46.6	53.2 48.5	58.29 51.65		

Table 5: Adversarial transferability to object segmentation (DINO) based on Jacard index metric is evaluated on DAVIS validation set [27].

4.1.2 Transferability from Cross-Domain Samples

Extra unsupervised data can boost the performance of adversarial training [7]. We extend our approach to large-scale, unlabelled datasets to observe its effect on adversarial transferability of our approach as explained below. All surrogate models are trained for 50 epochs for cross-task adversarial transferability experiments.

Classification: In this task, we train a single autoencoder and transfer its adversarial perturbations to ImageNet trained convolutional models as described in Sec. 3. We observe that our proposed adversarial training significantly improves upon the baseline [27] (Table 3). We further note that the surrogate models trained on ‘paintings’ dataset show higher adversarial transferability while ‘Rotation’ as pixel transformation performs better. Further analysis on adversarial transferability of our attack is provided in Appendix A.

Object Detection and Segmentation: Our adversarial attack based on simple transformed reconstruction error (Eq. 1) compliments our proposed adversarial training and successfully fools DETR for object detection and DINO for video segmentation (Tables 4 & 5). This signifies that our attack can be launched in real-world setting without any knowledge about the deployed vision system.

4.2 Ablative Analysis

We thoroughly analyze and develop better understanding about the behavior of our approach by studying the effect of its different components including, **a) Effect of Adversarial Pixel Restoration Prior on Training**, **b) Effect of Perceptual Budget ϵ for our Single step Attack (Eq. 1)**, **c) Effect of Iterative Attack during Training**, **d) Effect of Training Iterations and Data size**, and **e) Contribution of Losses**. All ablative experiments are conducted in limited data setting (Sec. 4.1.1).

Effect of Adversarial Pixel Restoration Prior: The number of parameters of the surrogate model are significantly higher than the number of input samples, a simple objective of reconstructing the original image can lead to identical mapping. While adding pixel transformations (such as rotation or jigsaw) somewhat alleviates this problem, our adversarial denoising with pixel transformation further resolves it by lowering overfitting, resulting in better generalizability. Fig. 5 (left) shows training loss comparison between simple reconstruction objective (*Naive*), [27] and our method. The surrogate auto-encoder quickly collapses

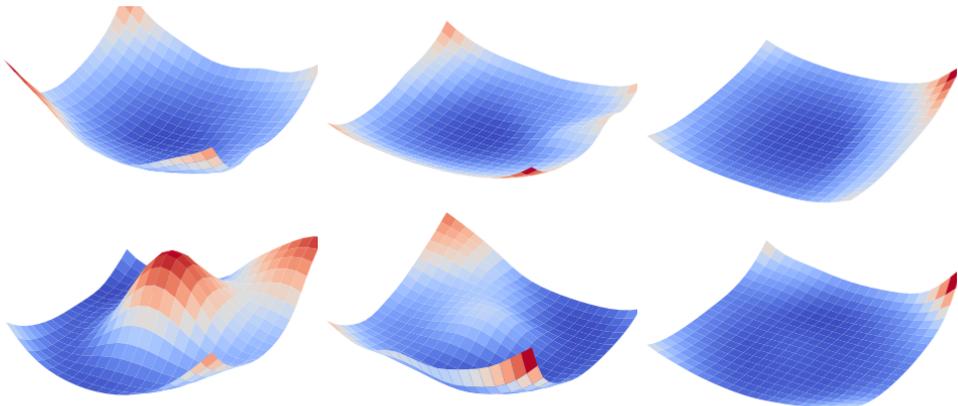


Figure 4: Loss landscape of surrogate models with increasing perceptual budget from left to right. The first row shows the loss surface on the clean samples, while as the second row plots the loss surface with respect to the corresponding adversaries.

Perceptual Budget ϵ (\rightarrow)	$\frac{2}{255}$	$\frac{4}{255}$	$\frac{8}{255}$
Rotation	26.87	27.86	34.01
Jigsaw	24.81	27.77	35.94
Prototypical	22.84	22.30	24.04

Training Loss (\rightarrow)	\mathcal{L}_{out}	$\mathcal{L}_{out} + \mathcal{L}_{feature}$
Rotation	29.62	26.87 (-3.03)
Jigsaw	27.13	24.81 (-4.59)
Prototypical	23.98	22.84 (-0.84)

Table 6: Effect of Perceptual Budget (ϵ). Top-1 average accuracy is reported on convolution networks under limited data constraint (Table 1).

Table 7: Contribution of Loss components. Top-1 average accuracy is reported on convolution networks under limited data constraint (Table 1).

to identity as the training loss overfits early on in training without using pixel restoration as prior. On the other hand, our approach allows meaningful representation learning with higher iterations.

Effect of Perceptual Budget (ϵ): Our adversarial training (Algo. 1) is computationally efficient as it is based on a single attack (maximization) step. The effect of attack step size on model generalizability and hence its transferability is presented in Table 6. We observe that adversarial perturbation computed on models trained with a smaller ϵ results in significant improvements in terms of attack transferability. In Fig. 4, we plot the loss surface around clean and adversarial examples on surrogate models trained with increasing perceptual budget. While the loss surface around clean examples becomes smoother as perceptual budget increases, it becomes harder to maximize the reconstruction error or flip decisions on the excessively smooth loss surface during attack. This behaviour is in line with [47] which shows that slightly robust classifier models generate highly transferable adversarial examples. This is further evident by the shift in attention caused by our method (Fig. 7).

Effect of Iterative Attack: During our adversarial training, increasing the attack iterations does not help to further boost the adversarial transferability; rather, performance often degrades significantly. In Fig. 5 (right), attack loss is plotted for surrogate models (at perceptual budget $\epsilon = 4$) trained with different number of attack iteration (or maximization steps). As the surrogate model becomes more robust with increasing attack iterations, we notice that maximizing the attack objective becomes more challenging.

Contribution of Losses: We explore the effect of our proposed pixel and feature reconstruction losses (Eq. 5) in Table 7. We observe that feature reconstruction complements the pixel reconstruction and leads to better surrogate model with more transferable adversarial space.

Training iterations and Data Size: We explore the effect on transferability of surrogate

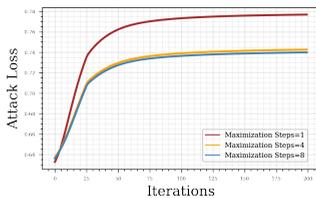
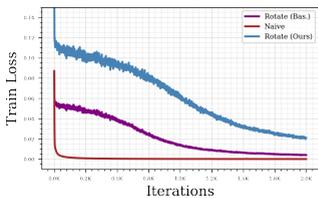


Figure 5: Comparing effect of Adversarial Pixel Restoration Prior on the training loss of surrogate models(*left*) and effect of increasing maximization steps during training on the attack loss(*left*).

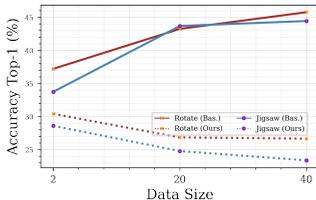
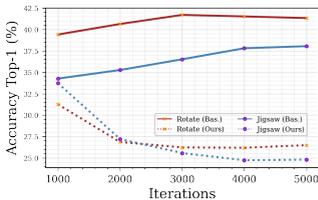


Figure 6: Top-1 (%) average accuracy against convolutional networks. The performance of our method improves with more training iterations and data size in contrast to the baseline [26].

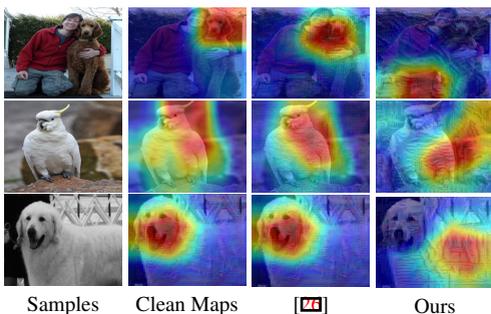


Figure 7: GradCAM [24] explanation of adversarial examples. The activation maps were generated on ImageNet pretrained ResNet50 model. Adversarial examples are transferred from surrogate trained via our adversarial rotation restoration task at the perceptual budget of $\epsilon \leq 0.1$ [26]. Our approach significantly shifts the attention of the model, boosting the mis-classification rates on the adversarial examples (see Appendix C for more qualitative examples).

models w.r.t **a)** training iterations, and **b)** the number of data samples in Fig. 6. Note that [26] trains a single autoencoder on 20 samples and thus needs 250 autoencoder to attack 5k ImageNet validation samples. The number of autoencoders increases to 2.5k when only 2 samples are available for training of single model. In the same setting, the performance of our method improves with more iterations (Fig. 6 *left plot*) in contrast to [26]. We report average Top-1 (%) accuracy across all convolutional models (*lower is better*). Similarly, as we increase the number of data samples during training and reduce the number of autoencoders, the performance of our approach significantly improves as compared to the baseline [26]. This indicates that our adversarial objective successfully reduces overfitting while increases generalizability of crafted adversarial perturbations.

5 Conclusion

In this work, we show the benefits of adversarial training to learn transferable adversarial perturbations. Our approach trains an effective surrogate by learning to restore adversarial pixel transformations created via our proposed attack. Our adversarial training reduces overfitting during training and can exploit very few data samples to learn meaningful adversarial features while it can also scale to large unsupervised datasets. Our attack is task independent and allows cross-domain attacks (*e.g.*, learning surrogate on comics and transferring its perturbations to models trained on natural images). Our results bring attention to the use of self-supervised adversarial training for transferable adversarial attacks.

References

- [1] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning (ICML)*, 2018.
- [2] Cenk Bircanoğlu. Comic books classification dataset : <https://www.kaggle.com/datasets/cenkbircanoglu/comic-books-classification>. Kaggle, 2017.
- [3] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *International Conference on Learning Representations (ICLR)*, 2018.
- [4] Painter by Number. <https://www.kaggle.com/c/painter-by-numbers/data>. Kaggle, 2017.
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision (ECCV)*. Springer, 2020.
- [6] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*, 2017.
- [7] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled data improves adversarial robustness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [9] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui 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.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [11] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Discovering adversarial examples with momentum. *arXiv preprint arXiv:1710.06081*, 2017.
- [12] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

- [14] Lianli Gao, Qilong Zhang, Jingkuan Song, and Heng Tao Shen. Patch-wise++ perturbation for adversarial targeted attacks. *arXiv preprint arXiv:2012.15503*, 2020.
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [16] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [18] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [19] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [20] Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, 2019.
- [21] Nathan Inkawhich, Kevin Liang, Binghui Wang, Matthew Inkawhich, Lawrence Carin, and Yiran Chen. Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- [24] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pages 99–112. Chapman and Hall/CRC, 2018.
- [25] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems (NeurIPS)*, 2018.
- [26] Qizhang Li, Yiwen Guo, and Hao Chen. Practical no-box adversarial attacks against dnns. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision (ECCV)*. Springer, 2014.

- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [29] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. DeepFool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] Nina Narodytska and Shiva Kasiviswanathan. Simple black-box adversarial attacks on deep neural networks. In *Computer Vision and Pattern Recognition (CVPR) Workshop*, 2017.
- [31] Muhammad Muzammal Naseer, Salman H Khan, Muhammad Haris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-domain transferability of adversarial perturbations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [32] Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. A self-supervised approach for adversarial robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [33] Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. On generating transferable targeted perturbations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- [34] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [35] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169*, 2021.
- [36] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision (ECCV)*, 2016.
- [37] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [38] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *Euro S&P*, 2016.
- [39] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Asia Conference on Computer and Communications Security*, 2017.
- [40] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.

- [41] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [42] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [43] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [44] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017.
- [45] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [46] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems (NeurIPS)*, 2017.
- [47] Jacob Springer, Melanie Mitchell, and Garrett Kenyon. A little robustness goes a long way: Leveraging robust features for targeted transfer attacks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [48] Jacob M Springer, Melanie Mitchell, and Garrett T Kenyon. Adversarial perturbations are not so weird: Entanglement of robust and non-robust features in neural network classifiers. *arXiv preprint arXiv:2102.05110*, 2021.
- [49] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [50] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Conference (CVPR)*, 2016.
- [51] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning (ICML)*, 2021.
- [52] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- [53] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the transferability of adversarial attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

-
- [54] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [55] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference (BMVC)*, 2016.
- [56] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- [57] Mingyi Zhou, Jing Wu, Yipeng Liu, Shuaicheng Liu, and Ce Zhu. Dast: Data-free substitute training for adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.