

FACL-Attack: Frequency-Aware Contrastive Learning for Transferable Adversarial Attacks

Hunmin Yang^{1,2,*}, Jongoh Jeong^{1,*}, Kuk-Jin Yoon¹

¹Visual Intelligence Lab., KAIST

²Agency for Defense Development

{hmyang, jeong2, kjyoon}@kaist.ac.kr

Abstract

Deep neural networks are known to be vulnerable to security risks due to the inherent transferable nature of adversarial examples. Despite the success of recent generative model-based attacks demonstrating strong transferability, it still remains a challenge to design an efficient attack strategy in a real-world strict black-box setting, where both the target domain and model architectures are unknown. In this paper, we seek to explore a feature contrastive approach in the frequency domain to generate adversarial examples that are robust in both cross-domain and cross-model settings. With that goal in mind, we propose two modules that are only employed during the training phase: a **Frequency-Aware Domain Randomization (FADR)** module to randomize domain-variant low- and high-range frequency components and a **Frequency-Augmented Contrastive Learning (FACL)** module to effectively separate domain-invariant mid-frequency features of clean and perturbed image. We demonstrate strong transferability of our generated adversarial perturbations through extensive cross-domain and cross-model experiments, while keeping the inference time complexity.

Introduction

Deep neural networks have brought forth tremendous improvements in visual recognition tasks. However, the inherent transferable nature of adversarial examples still exposes the security vulnerability to malicious attackers targeting such susceptible classifiers, causing serious threats and undesirable outcomes in real-world applications. The majority of current attack methods can be primarily classified into two main categories: iterative or optimization-based approaches, and generative model-based approaches. Over the past years, iterative attack methods (Goodfellow, Shlens, and Szegedy 2015; Madry et al. 2018, 2017; Croce and Hein 2020; Lorenz et al. 2021; Dong et al. 2018; Xie et al. 2019; Lu et al. 2020; Naseer et al. 2020) have been the standard attack protocol for its simplicity and effectiveness. However, this iterative approach is frequently constrained by inefficient time complexity and the potential risk of over-fitting to the training data and models. Moreover, it has shown limited applicability in practical situations due to the low transferability to unknown models and domains.

*These authors contributed equally.

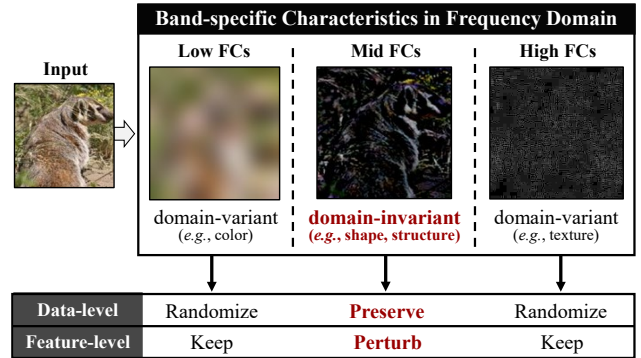


Figure 1: To boost the transferability of adversarial examples, we exploit band-specific characteristics of natural images in the frequency domain. Our method randomizes domain-variant low- and high-band frequency components (FCs) in the data space, and contrasts domain-invariant mid-range clean and perturbed feature pairs in the feature space.

Regarding the transferability of adversarial attacks, threat model is typically carried out in three different settings (*i.e.*, white-box, black-box, and strict black-box) depending on the prior knowledge of the model architecture and data distributions by the adversary. In each respective setting, the adversary has either complete knowledge of the target model profile (*i.e.*, architecture and weights) and data distributions reflecting the target domain, query access to the limited black-box only, or no information at all. In this work, we specifically consider the *strict black-box* case in which the victim attributes are completely *unknown* to the attacker since such a scenario is commonly encountered in practical real-world settings. We believe that crafting adversarial examples in this strict black box setting has practical values towards stronger transferability, as well as safe and reliable deployment of deep learning models.

In this light, generative attacks (Poursaeed et al. 2018; Naseer et al. 2019; Salzman et al. 2021; Naseer et al. 2021; Zhang et al. 2022) have recently gained attention, demonstrating the high transferability across unknown models and domains. Moreover, generator-based attacks yield lower time complexity than iterative or optimization-based methods in the inference stage, which is also a crucial part

for real-world attacks. While the current chain of generative attack methods (Poursaeed et al. 2018; Naseer et al. 2019, 2021; Salzmänn et al. 2021; Zhang et al. 2022; Wu et al. 2020) are time-efficient and effective against various black-box settings, we remark that their methods do not actively leverage domain-related characteristics to facilitate more transferable attacks.

In that sense, our work is inspired by frequency domain manipulations (Yin et al. 2019; Wang et al. 2020a,b) in domain adaptation (DA) (Yang and Soatto 2020) and generalization (DG) (Huang et al. 2021; Xu et al. 2021), demonstrating the superior generalization capabilities of the trained model. As we target transferable attack on unknown target domains and victim models to boost the transferability in a similar setting, we seek to exploit domain-related characteristics from simple yet effective frequency manipulations.

Several recent studies have focused on frequency-based adversarial attacks to manipulate adversarial examples, aimed at deeper understanding of their dataset dependency (Maiya et al. 2021), adversarial robustness (Duan et al. 2021), and the security vulnerability (Dziugaite, Ghahramani, and Roy 2016). With a slightly different motive, SSAH (Luo et al. 2022) aims to improve the perceptual quality, whereas (Guo, Frank, and Weinberger 2019) designs low-frequency perturbations to enhance the efficiency of black-box queries. Although low-frequency perturbations are efficient, they are known to provide less effective transfer between models (Sharma, Ding, and Brubaker 2019). As such, we delve deeper into frequency-driven approaches that effectively enhance the transferability of adversarial examples, especially crafted in a generative framework.

To this end, we propose a novel generative attack method, **FACL-Attack**, to facilitate transferable attacks across various domains and models from the frequency domain perspective. In our training, we introduce frequency-aware domain randomization and feature contrastive learning, explicitly leveraging band-specific characteristics of image attributes such as color, shape, and texture, as illustrated in Figure 1. We highlight our contributions as follows:

- We propose two modules to boost the adversarial transferability, *FADR* and *FACL*, in which *FADR* randomizes *domain-variant* data components while *FACL* contrasts *domain-invariant* feature pairs in the frequency domain.
- We achieve the state-of-the-art attack transferability across various domains and model architectures, demonstrating the effectiveness of our method.
- Our plug-and-play modules can be easily integrated into existing generative attack frameworks, further boosting the transferability while keeping the time complexity.

Related Work

Generator-based Adversarial Attack

Generative attack (Poursaeed et al. 2018; Naseer et al. 2019, 2021; Salzmänn et al. 2021; Zhang et al. 2022) employs the concept of adversarial training (Goodfellow et al. 2020) to create perturbations across entire data distributions. This is achieved by regarding a pre-trained surrogate model as

a discriminator, and it is advantageous due to the ability of generating diverse forms of perturbations across multiple images simultaneously. Existing methods aim to enhance the generator training by leveraging both the cross-entropy (CE) loss (Poursaeed et al. 2018) and the relativistic CE loss (Naseer et al. 2019), improving the transferability across domains and models. Recent studies (Salzmänn et al. 2021; Zhang et al. 2022) utilize features extracted from the mid-level layers of the surrogate model, which encompass a higher degree of shared information among different model architectures. We follow the traces of the recent works and explore a method to further enhance the transferability by introducing a novel perspective from the frequency domain.

Frequency-based Approach for Generalization

Convolutional neural networks are known to exhibit intriguing attributes within the frequency domain (Yin et al. 2019; Tsuzuku and Sato 2019; Yin et al. 2019; Wang et al. 2020a,b), demonstrating proficient generalization capability by effectively harnessing the band-specific information derived from Fourier filtering (Dziugaite, Ghahramani, and Roy 2016; Guo et al. 2017; Long et al. 2022). Spectral manipulations for enhancing the generalization capability can be achieved through simple yet powerful transformations like the Fast Fourier Transform (FFT), which dissects an image into amplitude components that vary across domains and phase components that remain consistent across different domains (Xu et al. 2021). The Discrete Cosine Transform (DCT) also serves as an efficient technique to decompose spectral elements into domain-agnostic mid-frequency components (mid-FCs) and domain-specific low- and high-FCs, which contributed to the effective spectral domain randomization in FSDR (Huang et al. 2021). In our work, we also employ the DCT to decompose images into domain-agnostic and domain-specific frequency components, facilitating the effective domain randomization and feature-level contrastive learning for transferable attacks.

Feature Contrastive Learning

Manipulating image representations in the feature space has demonstrated significant performance improvement in real-world scenarios characterized by domain shifts. In the field of DA and DG, common approaches such as feature alignment (Yang et al. 2022) and intra-class feature distance minimization with inter-class maximization (Kang et al. 2019; Luo et al. 2022; Jeong and Kim 2022) are successful in mitigating the domain discrepancies. Specifically, several studies (Wang et al. 2022; Kim et al. 2021) have directly addressed the domain gap issue by manipulating pairs of domain-invariant representations from various domains that correspond to samples of the same class. Continuing in the realm of generative attacks, recent studies have employed CLIP-based (Aich et al. 2022) and object-centric (Aich et al. 2023) features for effective training of the perturbation generator. In our work, we leverage frequency-augmented feature contrastive learning on domain-agnostic mid-band feature pairs. Simultaneously, we reduce the significance of domain-specific features in the low- and high-bands to improve the adversarial transferability.

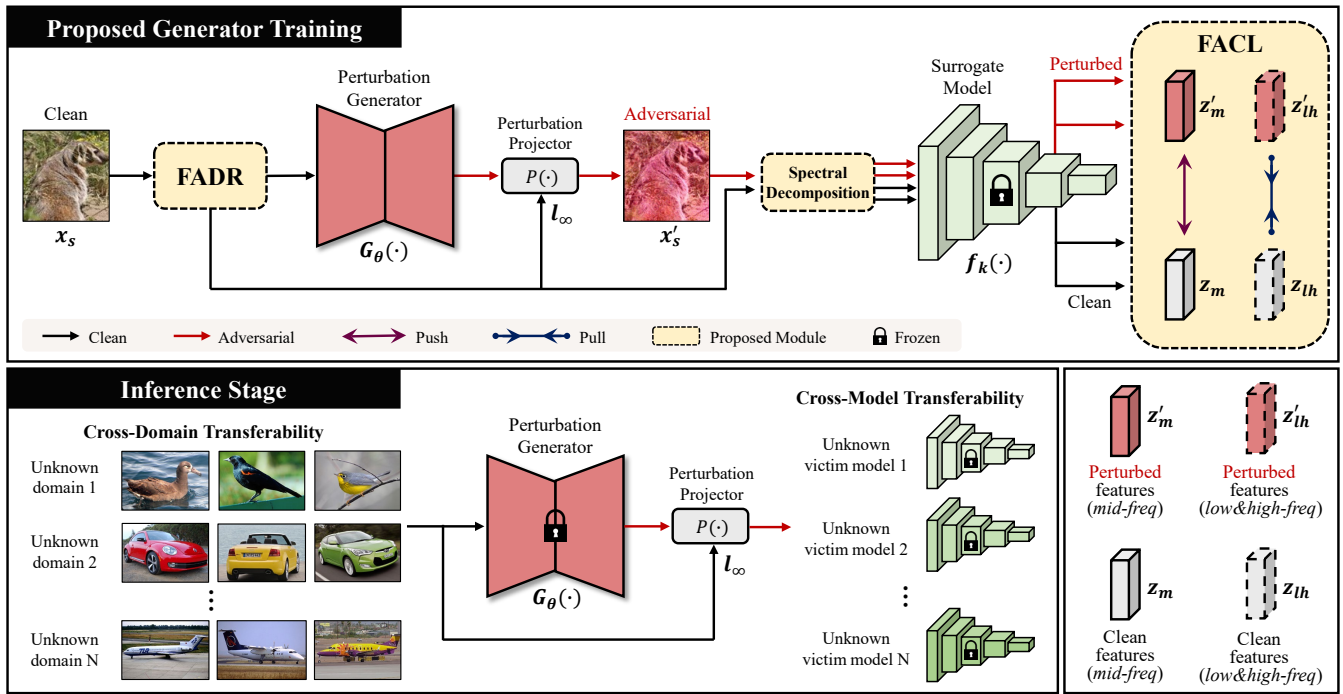


Figure 2: Overview of FACL-Attack. From the clean input image, our FADR module outputs the augmented image after spectral transformation, which is targeted to randomize only the domain-variant low/high FCs. The perturbation generator $G_{\theta}(\cdot)$ then produces the l_{∞} -budget bounded adversarial image x'_s with perturbation projector $P(\cdot)$ from the randomized image. The resulting clean and adversarial image pairs are decomposed into mid-band (domain-agnostic) and low/high-band (domain-specific) FCs, whose features $f_k(\cdot)$ extracted from the k -th layer of the surrogate model are contrasted in our FACL module to boost the adversarial transferability. The adversarial image x'_s is colored only for visualization.

Proposed Attack Method: FACL-Attack

Problem definition. Generating adversarial examples revolves around solving an optimization problem, whereas generating transferable adversarial examples addresses the challenge of generalization. Our goal is to train a generative model $G_{\theta}(\cdot)$ to craft adversarial perturbations δ that are well transferable to arbitrary domains and victim models aimed to trigger mispredictions on the image classifier $f(\cdot)$. Specifically, the generator maps the clean image x to its corresponding adversarial example $x' = G_{\theta}(x)$ containing perturbations constrained by $\|\delta\|_{\infty} \leq \epsilon$.

Overview of FACL-Attack. Our method aims to train a robust perturbation generator that yields effective adversarial examples given arbitrary images from black-box domains to induce the unknown victim model to output misclassification. It consists of two key modular operations in the frequency domain, each applied to the input image data and features extracted from the surrogate model only during the training stage, as illustrated in Figure 2.

As inspired by the power of frequency domain augmentation in domain generalization (Huang et al. 2021; Xu et al. 2021), our first module, *Frequency-Aware Domain Randomization (FADR)*, transforms a pixel-domain image to the frequency-domain components using DCT. It randomizes *domain-variant* low- and high-frequency band components and preserves *domain-invariant* mid-frequency com-

ponents in the input image. Then a perturbation generator is trained to craft bounded adversarial images x'_s , i.e., perturbation δ added to the clean image x_s and constrained by perturbation projector $P(\cdot)$. We then spectrally decompose the randomized x_s and x'_s into each low- and high-band, and mid-band frequency component, which are inversely transformed to the image domain by IDCT and passed through the pre-defined surrogate model for feature extraction. Following the recent line of works (Salzmann et al. 2021; Zhang et al. 2022) on transferable generative attacks, we leverage the mid-layer features $f_k(\cdot)$ for feature contrastive learning. Each band-specific clean and perturbed feature pair is contrasted in our *Frequency-Augmented Contrastive Learning (FACL)* module, whereby *domain-agnostic* mid-band FC pair is to repel and *domain-specific* low- and high-band FC pair is to attract each other. This straightforward but effective data- and feature-level guidance in the frequency domain significantly contributes to boost the adversarial transferability as demonstrated in the following sections.

Frequency-Aware Domain Randomization

This subsection describes our FADR module designed to boost the robustness of perturbation generator $G_{\theta}(\cdot)$ against arbitrary domain shifts in practical real-world scenarios. Inspired by recent works that convert the training image from pixel space into frequency space for boosting the do-




Figure 3: Visualization of spectral transformation in FADR. From the clean input image (column 1), our FADR decomposes the image into mid-band (column 2) and low/high-band (column 3) FCs. The FADR only randomizes the low/high-band FCs, yielding the augmented output in column 4. Here we demonstrate transformations with large hyper-parameters of $\rho = 0.5$ and $\sigma = 8$ for visualization.

main generalization capabilities (Huang et al. 2021; Xu et al. 2021), we decompose the input training images into multiple-range FCs by leveraging DCT, and apply random masked filtering operation on *domain-specific* image attributes that lie in the low- and high-frequency bands. While FSDR (Huang et al. 2021) and FACT (Xu et al. 2021) each employs histogram matching and Fourier-based amplitude mix-up, our proposed FADR module explicitly manipulates the DCT coefficients to diversify input images, aligning with a recent work (Long et al. 2022) that narrows the gap between the surrogate model and possible victim models via spectrum transformation. We remark that our approach applies domain randomization exclusively to *domain-specific* FCs that are subject to change from various domains, whereas the existing work (Long et al. 2022) applies spectral transformation over the whole frequency bands containing not only domain-specific information, but also domain-agnostic semantic details.

In converting the input images into the frequency domain, we apply DCT to each channel separately. We then apply random masked filtering to diversify the input images for boosting the cross-domain transferability. Our spectral transformation operation $\mathcal{T}_{\text{FADR}}(\cdot)$ for source images \mathbf{x}_s can be mathematically expressed as follows:

$$\mathcal{T}_{\text{FADR}}(\mathbf{x}_s) = \phi^{-1}\left((\phi(\mathbf{x}_s + \boldsymbol{\xi}) \odot \mathbf{M})\right), \quad (1)$$

with the mask \mathbf{M} defined as follows:

$$\mathbf{M} = \begin{cases} \mathcal{U}(1 - \rho, 1 + \rho), & \text{if } f < f_l, \\ 1, & \text{if } f_l \leq f < f_h, \\ \mathcal{U}(1 - \rho, 1 + \rho), & \text{if } f \geq f_h, \end{cases} \quad (2)$$

where \odot , ϕ , ϕ^{-1} denote Hadamard product, DCT, and inverse DCT (IDCT) operation, respectively. The random noise $\boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is sampled from a Gaussian distribution, and the mask values are randomly sampled from Uniform distribution, denoted as \mathcal{U} . For the random mask matrix \mathbf{M} which has same dimension with the DCT output, we assign its matrix component values as defined in Equation 2,

where we set the low and high thresholds as f_l , and f_h , respectively, to distinguish low-, mid-, and high-frequency bands. Note that we can parameterize our FADR module with hyper-parameters ρ and σ . The spectral transformation in our FADR module is conceptually illustrated in Figure 3.

The augmented image output from FADR is then fed as input to the generator $G_\theta(\cdot)$ to yield the adversarial image $\mathbf{x}'_s = P(G_\theta(\mathcal{T}_{\text{FADR}}(\mathbf{x}_s)))$, after the perturbation projection within the pre-defined budget of $\|\delta\|_\infty \leq \epsilon$.

Frequency-Augmented Contrastive Learning

Recent works on multi-object scene attacks have highlighted the importance of feature-level contrast for transferable generative attacks. In a similar approach to their ideas of exploiting local patch differences (Aich et al. 2023) or CLIP features (Aich et al. 2022), our FACL module seeks to apply feature contrast specifically in the *domain-agnostic* mid-frequency range for improving the generalization capability of the trained perturbation generator $G_\theta(\cdot)$.

Spectral decomposition. According to the training pipeline of our FACL-Attack in Figure 2, the generated adversarial image \mathbf{x}'_s undergoes spectral decomposition before feature extraction from the surrogate model. This process is carried out by using a band-pass filter \mathbf{M}_{bp} and a band-reject filter \mathbf{M}_{br} , to decompose the surrogate model inputs into mid- and low/high-band FCs, respectively. The spectral decomposition operator is defined as follows:

$$\mathbf{M}_{\text{bp}} = \begin{cases} 1, & \text{if } f_l \leq f < f_h, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where \mathbf{M}_{br} is the opposite of \mathbf{M}_{bp} , holding its values in reverse. Then the spectrally decomposed features from the surrogate model \mathbf{f} are defined as:

$$\mathbf{z}_{\text{band}} = f_k\left(\phi^{-1}\left(\phi(\mathbf{x}_{\text{input}}) \odot \mathbf{M}_{\text{band}}\right)\right), \quad (4)$$

where \mathbf{M}_{band} is set to either \mathbf{M}_{bp} or \mathbf{M}_{br} , and $f_k(\cdot)$ denotes the k -th layer of f . Given \mathbf{x}_s and \mathbf{x}'_s as input, we finally obtain two pairs of band-specific frequency-augmented features to contrast, *i.e.*, $(\mathbf{z}_m, \mathbf{z}'_m)$ for repelling, and $(\mathbf{z}_{lh}, \mathbf{z}'_{lh})$ for attracting each other.

Loss function. The baseline loss $\mathcal{L}_{\text{orig}}$ for attacking the surrogate model via contrasting clean and adversarial feature pairs is defined as follows:

$$\mathcal{L}_{\text{orig}} = \text{sim}(f_k(\mathbf{x}_s), f_k(\mathbf{x}'_s)), \quad (5)$$

where sim refers to the standard cosine similarity metric. To boost the attack transferability, our FACL module effectively exploits the spectrally decomposed feature pairs in our proposed FACL loss function defined as follows:

$$\mathcal{L}_{\text{FACL}} = \text{sim}(\mathbf{z}_m, \mathbf{z}'_m) - \text{sim}(\mathbf{z}_{lh}, \mathbf{z}'_{lh}), \quad (6)$$

where the goal of $\mathcal{L}_{\text{FACL}}$ is to reinforce the effectiveness of *domain-agnostic* mid-band feature contrast $(\mathbf{z}_m, \mathbf{z}'_m)$, while minimizing the importance of *domain-specific* low- and high-band feature difference $(\mathbf{z}_{lh}, \mathbf{z}'_{lh})$. In this approach, our $\mathcal{L}_{\text{FACL}}$ facilitates the push-pull action among the band-specific feature pairs, further guiding the perturbation generation towards more robust regime, as shown in Figure 4.



Figure 4: Clean image, unbounded adversarial images from baseline and FACL, and the final difference map (Diff(baseline, baseline+FACL)), from left to right. Our generated adversarial perturbations are more focused on domain-agnostic semantic region such as shape, facilitating more transferable attack.

Final learning objective. We train our perturbation generator by minimizing the total loss function as follows:

$$\min_{\theta} (\lambda_{\text{orig}} \cdot \mathcal{L}_{\text{orig}} + \lambda_{\text{FACL}} \cdot \mathcal{L}_{\text{FACL}}), \quad (7)$$

where λ_{orig} and λ_{FACL} are loss coefficients. The objective guides our generator $G_{\theta}(\cdot)$ to generate more robust perturbations to domain shifts as well as model variances.

Experiments

Experimental Setup

Datasets and attack settings. We evaluate our method over challenging strict black-box settings (*i.e.*, *cross-domain* and *cross-model*) in the image classification task. We set the target domain and victim model to be different from the source domain and surrogate model. The perturbation generator is trained on ImageNet-1K (Russakovsky et al. 2015) and evaluated on CUB-201-2011 (Wah et al. 2011), Stanford Cars (Krause et al. 2013), and FGVC Aircraft (Maji et al. 2013). As BIA (Zhang et al. 2022) highlights the importance of using a large-scale dataset for training, we train on ImageNet-1K accordingly. For the cross-model setting, we evaluate our method over black-box models but white-box domain (*i.e.*, ImageNet-1K) setting. The details for the datasets are described in Table 1.

Surrogate and victim models. Our perturbation generator is trained against ImageNet-1K pre-trained surrogate models (*e.g.*, VGG-16 (Simonyan and Zisserman 2015)). For the cross-model evaluation, we investigate other architectures including VGG-19 (Simonyan and Zisserman 2015), ResNet50 (Res-50), ResNet152 (Res-152) (He et al. 2016), DenseNet121 (Dense-121), DenseNet169 (Dense-169) (Huang et al. 2017) and Inception-v3 (Inc-v3) (Szegedy et al. 2016). For the cross-domain setting (*i.e.*, CUB-201-2011, Stanford Cars, and FGVC Aircraft), we use fine-grained classification models trained under DCL framework (Chen et al. 2019) with three different backbones, which include Res-50, SENet154 and SE-ResNet101 (SE-Res101) (Hu, Shen, and Sun 2018).

Dataset	# Class	# Train / Val.	Resolution
ImageNet-1K	1,000	1.28 M / 50,000	224×224
CUB-200-2011	200	5,994 / 5,794	448×448
Stanford Cars	196	8,144 / 8,041	448×448
FGVC Aircraft	100	6,667 / 3,333	448×448

Table 1: Description of datasets.

Implementation details. We closely follow the implementation of recent works on generative attacks (Naseer et al. 2019; Salzmann et al. 2021; Zhang et al. 2022; Aich et al. 2022) for fair comparison. Our perturbation generator consists of down-sampling, residual, and up-sampling blocks that translate clean images into adversarial examples. The surrogate model layer from which we extract features is *Maxpool.3* for VGG-16. We train with an Adam optimizer ($\beta_1 = 0.5$, $\beta_2 = 0.999$) (Kingma and Ba 2015) with the learning rate of 2×10^{-4} , and the batch size of 16 for 1 epoch. The perturbation budget for crafting the adversarial image is $l_{\infty} \leq 10$. For the FADR hyper-parameters, we follow a prior work (Huang et al. 2021) to set the low and high frequency threshold to $f_l = 7$ and $f_h = 112$, respectively. We use $\rho = 0.01$ and $\sigma = 8$ for spectral transformation and describe more details in *Supplementary*.

Evaluation metric and competitors. We choose the top-1 classification accuracy after attacks as our main evaluation metric, unless otherwise stated. The reported results are the average values obtained from three random seed runs. The competitors include the state-of-the-art generative attacks such as GAP (Poursaeed et al. 2018), CDA (Naseer et al. 2019), LTP (Salzmann et al. 2021), and BIA (Zhang et al. 2022). We set BIA as our baseline.

Main Results

Cross-domain evaluation results. We compare our FACL-Attack with the state-of-the-art generative-model based attacks on various black-box domains with black-box models. During the training stage, we leverage the ImageNet-1K as the source domain to train a perturbation generator against a pre-trained surrogate model. In the inference stage, the trained perturbation generator is evaluated on various black-box domains (*i.e.*, CUB-200-2011, Stanford Cars, and FGVC Aircraft) with black-box victim models. The victim models include pre-trained models which were trained via DCL framework with three different backbones (*i.e.*, Res-50, SENet154, and SE-Res101).

As shown in Table 2, our FACL-Attack outperforms on most *cross-domain* benchmarks, among which are also *cross-model*, by significant margins. This demonstrates the strong and robust transferable capability of the generator trained by our novel approach with data- and feature-level guidance in the frequency domain. We posit that the remarkable generalization ability of FACL-Attack owes to the synergy between our two proposed modules that effectively guide feature-level separation in the *domain-agnostic* mid-frequency band (*i.e.*, FACL), complemented by data-level randomization only applied to the *domain-specific* fre-

Method	CUB-200-2011			Stanford Cars			FGVC Aircraft			AVG.
	Res-50	SENet154	SE-Res101	Res-50	SENet154	SE-Res101	Res-50	SENet154	SE-Res101	
Clean	87.35	86.81	86.56	94.35	93.36	92.97	92.23	92.08	91.90	90.85
GAP (Poursaeed et al. 2018)	68.85	74.11	72.73	85.64	84.34	87.84	81.40	81.88	76.90	79.30
CDA (Naseer et al. 2019)	69.69	62.51	71.00	75.94	72.45	84.64	71.53	58.33	63.39	69.94
LTP (Salzmann et al. 2021)	<u>30.86</u>	<u>52.50</u>	62.86	<u>34.54</u>	65.53	73.88	15.90	60.37	52.75	<u>49.91</u>
BIA (Zhang et al. 2022)	32.74	52.99	<u>58.04</u>	39.61	69.90	<u>70.17</u>	28.92	60.31	46.92	51.07
FACL-Attack (Ours)	24.74	44.06	53.75	26.58	<u>65.71</u>	61.40	<u>19.72</u>	52.01	<u>48.51</u>	44.05

Table 2: Cross-domain evaluation results. The perturbation generator is trained on ImageNet-1K with VGG-16 as the surrogate model and evaluated on black-box domains with black-box models. We compare the top-1 classification accuracy after attacks with the perturbation budget of $l_\infty \leq 10$ (the lower, the better).

Method	Venue	VGG-16	VGG-19	Res-50	Res-152	Dense-121	Dense-169	Inc-v3	AVG.
Clean	-	70.14	70.95	74.61	77.34	74.22	75.75	76.19	74.17
GAP (Poursaeed et al. 2018)	CVPR'18	23.63	28.56	57.87	65.50	57.94	61.37	63.30	55.76
CDA (Naseer et al. 2019)	NeurIPS'19	0.40	0.77	36.27	51.05	38.89	42.67	54.02	32.01
LTP (Salzmann et al. 2021)	NeurIPS'21	1.61	<u>2.74</u>	<u>21.70</u>	<u>39.88</u>	<u>23.42</u>	25.46	41.27	<u>22.30</u>
BIA (Zhang et al. 2022)	ICLR'22	1.55	3.61	25.36	42.98	26.97	32.35	<u>41.20</u>	<u>24.86</u>
FACL-Attack (Ours)	-	<u>1.45</u>	2.92	19.72	36.61	21.34	<u>25.61</u>	29.97	19.66

Table 3: Cross-model evaluation results. The perturbation generator is trained on ImageNet-1K with VGG-16 as the surrogate model and evaluated on black-box models including white-box model (*i.e.*, VGG-16). We compare the top-1 classification accuracy after attacks with the perturbation budget of $l_\infty \leq 10$ (the lower, the better).

quency components (*i.e.*, FADR). In other words, our spectral approach does help improve the generalization capability of the perturbation generator to other black-box domains as well as unknown network architectures. Moreover, our proposed training modules are complementary with existing generative-model based attack frameworks and can further improve the attack transferability, as demonstrated in *Supplementary*.

Cross-model evaluation results. Although we demonstrated the effectiveness of FACL-attack on boosting the transferability in strict black-box settings (*i.e.*, *cross-domain* as well as *cross-model*) as shown in Table 2, we further investigated on the black-box model scenario in a controlled white-box domain (*i.e.*, ImageNet-1K). In other words, the generator is trained against a surrogate model (*i.e.*, VGG-16) and evaluated on various victim models which include VGG-16 (white-box), VGG-19, Res-50, Res-152, Dense-121, Dense-169, and Inc-v3.

As shown in Table 3, ours also outperforms on most generative attacks where they seem to partially overfit to the white-box model (*i.e.*, VGG-16). Our outperforming results validate the strong transferability in *cross-model* attacks, in addition to *cross-domain*. We posit that the frequency-augmented feature learning could help the perturbation generator craft more robust perturbations, which exhibit better generalization capability to unknown feature space. This aligns with a recent finding (Long et al. 2022) that spectral data randomization contributes to enhance the transferability via simulating diverse victim models.

Method	$\mathcal{L}_{\text{orig}}$	$\mathcal{T}_{\text{FADR}}$	$\mathcal{L}_{\text{FACL}}$	Cross-Domain	Cross-Model
Clean				90.85	74.17
Baseline	✓			51.07	24.86
FADR	✓	✓		46.24	20.28
FACL	✓		✓	45.36	20.70
Ours	✓	✓	✓	44.05	19.66

Table 4: Ablation study on our proposed modules. $\mathcal{T}_{\text{FADR}}$ and $\mathcal{L}_{\text{FACL}}$ are defined in Eq. 1 and 6, respectively.

More Analyses

Ablation study on our proposed modules. We examined different attack designs to find out how our proposed modules contribute to the attack transferability. As shown in Table 4, we trained the perturbation generator by employing each method and evaluated under realistic black-box settings. *Cross-Domain* is defined as ImageNet-1K \rightarrow {CUB-200-2011, Stanford Cars, FGVC Aircraft} and *Cross-Model* indicates VGG-16 \rightarrow {VGG-16, VGG-19, Res-50, Res-152, Dense-121, Dense-169, Inc-v3}. *Baseline* is trained with $\mathcal{L}_{\text{orig}}$ without any data randomization or band-specific feature contrast. *FADR* is trained with $\mathcal{L}_{\text{orig}}$ and frequency-aware domain randomization using $\mathcal{T}_{\text{FADR}}$. *FACL* is trained with $\mathcal{L}_{\text{orig}}$ and band-specific feature contrast using $\mathcal{L}_{\text{FACL}}$.

As shown in Table 4, *Baseline* trained with naive mid-layer feature contrast (*i.e.*, $\mathcal{L}_{\text{orig}}$) does not perform well due to the domain bias and model over-fitting. *FADR* and *FACL* each outperforms *Baseline* by a large margin, demonstrat-

Method	Clean	Baseline	All-Rand	Ours
Cross-Domain	90.85	51.07	47.24	44.05
Cross-Model	74.17	24.86	21.68	19.66

Table 5: Comparison with domain randomization on the entire frequency band.

Method	Accuracy ↓	SSIM ↑	PSNR ↑	LPIPS ↓
BIA ($l_\infty \leq 10$)	24.86	0.73	28.71	0.49
Ours ($l_\infty \leq 10$)	19.66	0.72	28.61	0.49
Ours ($l_\infty \leq 9$)	23.85	0.75	29.48	0.47

Table 6: Comparison on image quality of adversarial examples with cross-model accuracy on ImageNet-1K.




Figure 5: Average cross-domain evaluation results across various frequency thresholds.

ing the importance of selectively randomizing the *domain-variant* data components and contrasting *domain-invariant* feature pairs for boosting the black-box transferability, respectively. Furthermore, *Ours* performs the best consistently. We speculate that *FADR* and *FACL* are complementary since data augmentation through our *FADR* facilitates the stable feature contrastive learning.

Comparison with full-band randomization. We further investigated on the effectiveness of our domain randomization scheme, comparing with the full-band frequency randomization as practiced before (Long et al. 2022). As shown in Table 5, our novel domain-aware approach is superior to the naive full-range randomization method (*i.e.*, *All-Rand*). Remarkably, *All-Rand* is closely related to a recent work, namely *SSA* (Long et al. 2022), which improves the iterative attack transferability by full-range spectral augmentation. Compared to *SSA*, our method exclusively randomizes the *domain-specific* low/high-FCs and exploits the frequency-augmented feature contrast. Ours outperforms *All-Rand* by 3.19%p and 2.02%p in each cross-domain and cross-model evaluation. Without identifying and preserving *domain-agnostic* information, even the state-of-the-art method could excessively randomize images, resulting in the degradation of image semantics and leading to the sub-optimal adversarial perturbation generation.

Sensitivity on frequency thresholds. We investigated the sensitivity of the chosen frequency thresholds to verify the robustness of our approach. As shown in Figure 5, our method shows robust attack performance across adjacent threshold values, surpassing the baseline performance.




Figure 6: Qualitative results. Clean images (row 1), unbounded adversarial images (row 2), and bounded ($l_\infty \leq 10$) adversarial images (row 3) are shown for various domains. All of the final unbounded adversarial image samples cause victim classifier models to make incorrect predictions.

This implies that *mid-frequency* range, in general, contains *domain-agnostic* information that is effective in generating more transferable perturbations against arbitrary domains and model architectures.

Analysis on image quality. Although our work is focused on generating more powerful adversarial perturbations, the image quality of the crafted adversarial examples should also be carefully examined. As shown in Figure 6, *FACL* Attack can craft effective and high-quality adversarial images with imperceptible perturbations. We also conducted a quantitative evaluation of image dissimilarity metrics between clean and adversarial image pairs, including SSIM, PSNR, and LPIPS. As shown in Table 6, we found that ours with a lower perturbation of $l_\infty \leq 9$ demonstrates superior image quality than the baseline with $l_\infty \leq 10$ while achieving better attack performance. In other words, it can yield better attack transferability with lower perturbation power and better image quality, which are very remarkable assets for real-world black-box attacks. We refer to *Supplementary* for more qualitative and quantitative evaluation results.

Conclusion

In this paper, we have introduced a novel generator-based transferable attack method, leveraging spectral transformation and feature contrast in the frequency domain. Our work drew inspiration from domain generalization approaches that utilize frequency domain techniques, adapting and enhancing them for the attack framework. In our method, we target spectral data randomization on *domain-specific* image components, and *domain-agnostic* feature contrast for training a more robust perturbation generator. Extensive evaluation results validate the effectiveness in practical black-box scenarios with domain shifts and model variances. It can also be integrated into existing attack frameworks, further boosting the transferability while keeping the inference time.

Acknowledgments

This work was partially supported by the Agency for Defense Development grant funded by the Korean Government, and by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF2022R1A2B5B03002636).

References

- Aich, A.; Khang-Ta, C.; Gupta, A.; Song, C.; Krishnamurthy, S. V.; Asif, M. S.; and Roy-Chowdhury, A. K. 2022. GAMA: Generative Adversarial Multi-Object Scene Attacks. *arXiv preprint arXiv:2209.09502*.
- Aich, A.; Li, S.; Song, C.; Asif, M. S.; Krishnamurthy, S. V.; and Roy-Chowdhury, A. K. 2023. Leveraging Local Patch Differences in Multi-Object Scenes for Generative Adversarial Attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1308–1318.
- Chen, Y.; Bai, Y.; Zhang, W.; and Mei, T. 2019. Destruction and Construction Learning for Fine-Grained Image Recognition. In *CVPR*.
- Croce, F.; and Hein, M. 2020. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, 2206–2216. PMLR.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting Adversarial Attacks with Momentum. In *CVPR*.
- Duan, R.; Chen, Y.; Niu, D.; Yang, Y.; Qin, A. K.; and He, Y. 2021. Advdrop: Adversarial attack to dnns by dropping information. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7506–7515.
- Dziugaite, G. K.; Ghahramani, Z.; and Roy, D. M. 2016. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Goodfellow, I.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *International Conference on Learning Representations*.
- Guo, C.; Frank, J. S.; and Weinberger, K. Q. 2019. Low Frequency Adversarial Perturbation. In Globerson, A.; and Silva, R., eds., *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*, volume 115 of *Proceedings of Machine Learning Research*, 1127–1137. AUAI Press.
- Guo, C.; Rana, M.; Cisse, M.; and Van Der Maaten, L. 2017. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2017. Densely Connected Convolutional Networks. In *CVPR*.
- Huang, J.; Guan, D.; Xiao, A.; and Lu, S. 2021. Fsd: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6891–6902.
- Jeong, J.; and Kim, J.-H. 2022. Doubly Contrastive End-to-End Semantic Segmentation for Autonomous Driving under Adverse Weather. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press.
- Kang, G.; Jiang, L.; Yang, Y.; and Hauptmann, A. G. 2019. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4893–4902.
- Kim, D.; Yoo, Y.; Park, S.; Kim, J.; and Lee, J. 2021. Self-freg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9619–9628.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Krause, J.; Deng, J.; Stark, M.; and Fei-Fei, L. 2013. Collecting a Large-Scale Dataset of Fine-Grained Cars.
- Long, Y.; Zhang, Q.; Zeng, B.; Gao, L.; Liu, X.; Zhang, J.; and Song, J. 2022. Frequency domain model augmentation for adversarial attack. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, 549–566. Springer.
- Lorenz, P.; Harder, P.; Straßel, D.; Keuper, M.; and Keuper, J. 2021. Detecting autoattack perturbations in the frequency domain. *arXiv preprint arXiv:2111.08785*.
- Lu, Y.; Jia, Y.; Wang, J.; Li, B.; Chai, W.; Carin, L.; and Velipasalar, S. 2020. Enhancing Cross-Task Black-Box Transferability of Adversarial Examples With Dispersion Reduction. In *CVPR*.
- Luo, C.; Lin, Q.; Xie, W.; Wu, B.; Xie, J.; and Shen, L. 2022. Frequency-driven imperceptible adversarial attack on semantic similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15315–15324.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *ICLR*.
- Maiya, S. R.; Ehrlich, M.; Agarwal, V.; Lim, S.-N.; Goldstein, T.; and Shrivastava, A. 2021. A frequency perspective of adversarial robustness. *arXiv preprint arXiv:2111.00861*.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M. B.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. volume abs/1306.5151.
- Naseer, M.; Khan, S.; Hayat, M.; Khan, F. S.; and Porikli, F. 2021. On generating transferable targeted perturbations. In

- Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7708–7717.
- Naseer, M.; Khan, S. H.; Hayat, M.; Khan, F. S.; and Porikli, F. 2020. A Self-supervised Approach for Adversarial Robustness. In *CVPR*.
- Naseer, M. M.; Khan, S. H.; Khan, M. H.; Shahbaz Khan, F.; and Porikli, F. 2019. Cross-domain transferability of adversarial perturbations. *Advances in Neural Information Processing Systems*, 32.
- Poursaeed, O.; Katsman, I.; Gao, B.; and Belongie, S. 2018. Generative adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4422–4431.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Li, F. 2015. ImageNet Large Scale Visual Recognition Challenge. *IJCV*.
- Salzmann, M.; et al. 2021. Learning transferable adversarial perturbations. *Advances in Neural Information Processing Systems*, 34: 13950–13962.
- Sharma, Y.; Ding, G. W.; and Brubaker, M. A. 2019. On the Effectiveness of Low Frequency Perturbations.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In Bengio, Y.; and LeCun, Y., eds., *ICLR*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the Inception Architecture for Computer Vision. In *CVPR*.
- Tsuzuku, Y.; and Sato, I. 2019. On the structural sensitivity of deep convolutional networks to the directions of fourier basis functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 51–60.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, California Institute of Technology.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020a. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8684–8694.
- Wang, R.; Wu, Z.; Weng, Z.; Chen, J.; Qi, G.-J.; and Jiang, Y.-G. 2022. Cross-domain contrastive learning for unsupervised domain adaptation. *IEEE Transactions on Multimedia*.
- Wang, Z.; Yang, Y.; Shrivastava, A.; Rawal, V.; and Ding, Z. 2020b. Towards frequency-based explanation for robust cnn. *arXiv preprint arXiv:2005.03141*.
- Wu, W.; Su, Y.; Chen, X.; Zhao, S.; King, I.; Lyu, M. R.; and Tai, Y.-W. 2020. 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 *CVPR*.
- Xu, Q.; Zhang, R.; Zhang, Y.; Wang, Y.; and Tian, Q. 2021. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14383–14392.
- Yang, C.; Cheung, Y.-M.; Ding, J.; Tan, K. C.; Xue, B.; and Zhang, M. 2022. Contrastive learning assisted-alignment for partial domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*.
- Yang, Y.; and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4085–4095.
- Yin, D.; Gontijo Lopes, R.; Shlens, J.; Cubuk, E. D.; and Gilmer, J. 2019. A fourier perspective on model robustness in computer vision. *Advances in Neural Information Processing Systems*, 32.
- Zhang, Q.; Li, X.; Chen, Y.; Song, J.; Gao, L.; He, Y.; and Xue, H. 2022. Beyond imagenet attack: Towards crafting adversarial examples for black-box domains. *arXiv preprint arXiv:2201.11528*.