

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TESSER: TRANSFER-ENHANCING ADVERSARIAL ATTACKS FROM VISION TRANSFORMERS VIA SPECTRAL AND SEMANTIC REGULARIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Adversarial transferability remains a critical challenge in evaluating the robustness of deep neural networks. In security-critical applications, transferability enables black-box attacks without access to model internals, making it a key concern for real-world adversarial threat assessment. While Vision Transformers (ViTs) have demonstrated strong adversarial performance, existing attacks often fail to transfer effectively across architectures, especially from ViTs to Convolutional Neural Networks (CNNs) or hybrid models. In this paper, we introduce **TESSER**, a novel adversarial attack framework that enhances transferability via two key strategies: (1) *Feature-Sensitive Gradient Scaling (FSGS)*, which modulates gradients based on token-wise importance derived from intermediate feature activations, and (2) *Spectral Smoothness Regularization (SSR)*, which suppresses high-frequency noise in perturbations using a differentiable Gaussian prior. These components work in tandem to generate perturbations that are both semantically meaningful and spectrally smooth. Extensive experiments on ImageNet across 14 diverse architectures demonstrate that TESSER achieves +10.9% higher attack success rate (ASR) on CNNs and +7.2% on ViTs compared to the state-of-the-art Adaptive Token Tuning (ATT) method. Moreover, TESSER significantly improves robustness against defended models, achieving 53.55% ASR on adversarially trained CNNs and +15% higher ASR on robust ViTs. Qualitative analysis shows strong alignment between TESSER’s perturbations and salient visual regions identified via Grad-CAM, while frequency-domain analysis reveals a 12% reduction in high-frequency energy, confirming the effectiveness of spectral regularization.

1 INTRODUCTION

Deep learning models, particularly Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), have achieved state-of-the-art performance across a broad spectrum of computer vision tasks (Carion et al., 2020; Zhu et al., 2021; Ma et al., 2022). Despite this progress, these models remain highly vulnerable to adversarial examples—carefully crafted perturbations that are imperceptible to humans but cause misclassification (Goodfellow et al., 2014; Guesmi et al., 2023; 2024a;b). In safety-critical applications such as autonomous driving and medical imaging, this fragility raises significant security concerns.

Although white-box attacks, where attackers have full access to model parameters, have been extensively studied, black-box settings are more realistic in practice. These are based on the principle of *transferability*, where adversarial examples generated on a surrogate model are expected to fool unseen target models. However, transferability across architectures, especially from ViTs to CNNs or hybrid models, remains limited due to two key challenges: (1) the lack of **semantic selectivity**, where all tokens are perturbed uniformly without considering their relevance to the model’s prediction, and (2) the presence of **high-frequency noise** in perturbations, which tends to encode brittle, model-specific artifacts that do not generalize well.

Several recent works, such as ATT (Ming et al., 2024) and TGR (Zhang et al., 2023), have explored ViT-specific mechanisms for improving transferability by truncating or regularizing gradient flows. However, these approaches either use fixed gradient masks or overlook token-level semantics, leading

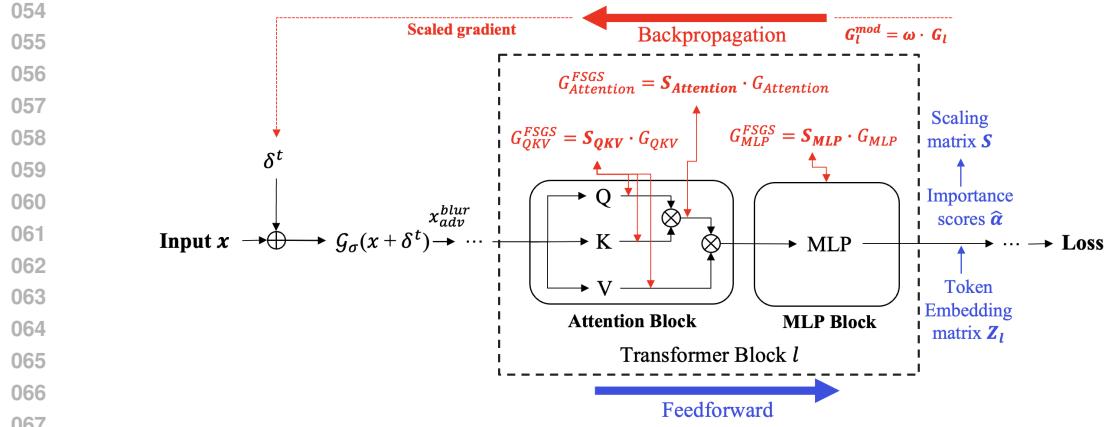


Figure 1: **Overview of the TESSER attack framework.** At each iteration, an adversarial perturbation δ^t is applied to the input image and smoothed via differentiable Gaussian blur $\mathcal{G}_\sigma(\cdot)$ to enforce spectral smoothness (SSR). The perturbed input is passed through the transformer, where token embeddings Z_l from each layer are used to compute token-wise importance scores $\hat{\alpha}$, which in turn define gradient scaling masks S . During backpropagation, gradients for the Attention, QKV, and MLP modules are reweighted according to their respective scaling masks ($S_{\text{Attention}}$, S_{QKV} , S_{MLP}) using Feature-Sensitive Gradient Scaling (FSGS). This encourages perturbations to align with semantically meaningful and transferable features while suppressing noise and irrelevant gradients.

to suboptimal alignment with transferable visual features. In this paper, we introduce **TESSER** (*Transfer-Enhancing Semantic and Spectral Regularization*) a novel adversarial attack framework specifically designed to improve black-box transferability from ViT-based models to a diverse set of architectures. TESSER integrates two complementary strategies:

- **Feature-Sensitive Gradient Scaling (FSGS):** a token-level gradient modulation method that scales gradients based on token importance derived from intermediate embeddings. Inspired by recent findings correlating token activation magnitudes with semantic relevance (Kobayashi et al., 2020; Wu et al., 2024; Modarressi et al., 2022), FSGS steers the attack toward semantically meaningful regions and away from background or non-informative tokens, enhancing cross-model generalization.
- **Spectral Smoothness Regularization (SSR):** a lightweight regularization mechanism that applies a differentiable Gaussian blur during each optimization step. SSR suppresses high-frequency noise, promoting low-frequency perturbations that are more resilient across architectures, particularly beneficial when transferring to CNNs and adversarially trained models.

Together, these modules enable TESSER to produce perturbations that are semantically aligned and spectrally smooth, two characteristics that we empirically demonstrate to be critical for enhancing transferability in adversarial attacks. Our main contributions are summarized as follows:

- We propose **TESSER**, a novel adversarial attack framework that combines semantic- and spectral-aware regularization to improve transferability from ViTs.
- We introduce Feature-Sensitive Gradient Scaling (FSGS), which reweights gradients for Attention, QKV, and MLP modules based on token-level importance, encouraging semantically aligned perturbations.
- We incorporate Spectral Smoothness Regularization (SSR) to reduce high-frequency noise and enhance cross-architecture generalization.
- We conduct extensive experiments on ImageNet across 14 diverse models (including ViTs, CNNs, and adversarially defended ViTs and CNNs), demonstrating that TESSER achieves up to **+10.9%** higher ASR over state-of-the-art baselines and consistently outperforms existing attacks in both black-box and robust scenarios.
- We conduct comprehensive ablation studies, Grad-CAM-based semantic alignment evaluations (Section 4.4), and frequency-domain analyses (Section 4.5) to demonstrate both the effectiveness and interpretability of our approach.

108
109
110
2 RELATED WORK111
112
113
114
115
116
117
118
119
120
121
Adversarial Attacks on CNNs and ViTs. Adversarial attacks are small, human-imperceptible
perturbations intentionally added to input data to mislead deep learning models (Goodfellow et al.,
2014). For Convolutional Neural Networks (CNNs), numerous gradient-based attacks have been
proposed to improve transferability, including momentum-based methods (Dong et al., 2018a),
variance tuning (Huang et al., 2019), and gradient skipping techniques (Wu et al., 2020). These
methods aim to stabilize perturbation updates and avoid local optima in the input space. However,
attack techniques designed for CNNs do not transfer well to Vision Transformers (ViTs), which have
fundamentally different architectures and information flow patterns. Recent works have proposed
ViT-specific attacks that exploit token structure and attention mechanisms (Naseer et al., 2022; Wei
et al., 2022). For example, Token Gradient Regularization (TGR) (Zhang et al., 2023) modifies
intermediate-layer gradients to reduce token-wise variance, improving transferability within ViT
families.122
123
124
125
126
127
128
129
130
131
132
133
134
135
Regularizing gradients is an effective way to suppress model-specific patterns and improve cross-
model generalization. In CNNs, methods such as SGM (Wu et al., 2020) and BPA (Xiaosen et al.,
2023) aim to manipulate the gradient flow through skip connections or rectify distortions introduced by
nonlinearities. Others have employed gradient variance reduction (Huang et al., 2019) and ensemble-
based tuning (Xiong et al., 2022). Attacks based on feature information (Wang et al., 2021; Ganeshan
et al., 2019) focus on disrupting salient internal representations. However, improperly guided feature-
based attacks risk discarding useful information and reducing transferability. To mitigate this, neuron
attribution methods (Zhang et al., 2023) and attention map diversification (Ren et al., 2025) have been
explored, particularly in ViTs. DiffAttack Chen et al. (2025) leverages generative diffusion models to
craft adversarial examples, exploiting their ability to model natural image distributions. By iteratively
guiding the diffusion process with adversarial objectives, it produces perturbations that are both
transferable and perceptually realistic. Compared to gradient-based methods, DiffAttack introduces
higher computational cost but demonstrates stronger performance in black-box and cross-architecture
scenarios.136
137
138
139
140
141
142
143
Forward Propagation Refinement (FPR) (Ren et al., 2025) is a recent surrogate-refinement strategy
tailored for improving adversarial transferability on Vision Transformers (ViTs). Unlike prior
methods that modify only the backward pass (e.g., PNAPO, TGR, GNS), FPR explicitly refines
forward activations by diversifying attention maps and stabilizing token embeddings. Specifically,
Attention Map Diversification (AMD) introduces controlled stochasticity into selected attention heads
to mitigate surrogate overfitting and implicitly induce beneficial gradient vanishing, while Momentum
Token Embedding (MTE) accumulates historical token embeddings to avoid local-optimum instability
during attack iterations.144
145
146
147
148
149
150
151
ATT (Ming et al., 2024) introduces hybrid token gradient truncation by weakening gradients in
attention and QKV blocks across layers of a ViT model. It leverages empirical observations of
gradient variance to suppress high-magnitude gradients associated with overfitting, thereby improving
transferability. However, ATT applies static truncation and does not explicitly consider token-level
semantic relevance, which may limit its effectiveness when generalizing across diverse architectures.
In contrast, our method introduces *Feature-Sensitive Gradient Scaling (FSGS)*, which adaptively
reweights gradients at a token level based on feature norms. This allows us to preserve semanti-
cally important gradients while suppressing noisy or architecture-specific ones, achieving improved
transferability across ViTs, hybrids, and CNNs.152
153
154
155
156
157
158
159
160
161
Input Diversity and Spectral Regularization. Input diversity has been widely adopted to improve
adversarial transferability. DI-FGSM (Xie et al., 2019) applies random resizing and padding, while
PatchOut (Wei et al., 2022) discards patch-wise perturbations to prevent overfitting. Recent self-paced
extensions further refine patch discarding based on semantic guidance (Ming et al., 2024). While
these approaches diversify the spatial patterns of inputs, few works address the frequency structure
of perturbations. Our method incorporates *Spectral Smoothness Regularization (SSR)* by applying
differentiable Gaussian blur during optimization. SSR suppresses high-frequency noise and promotes
smooth perturbation patterns that generalize better across model architectures, particularly important
for CNNs and early ViT layers that rely on localized features. **Importantly, input diversity is**
orthogonal to our method, and can be combined with TESSER for further gains. We provide
additional results and analysis combining input diversity with our framework in the Appendix C.

162

3 METHODOLOGY

163

3.1 PRELIMINARIES

167 Let $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ denote an input image with ground-truth label $y \in \{1, 2, \dots, K\}$, and let $f(\cdot)$
 168 be a deep neural network classifier. The goal of an untargeted adversarial attack is to generate a
 169 perturbation δ such that the perturbed input $\mathbf{x}^{\text{adv}} = \mathbf{x} + \delta$ is misclassified by the model, i.e., $f(\mathbf{x}^{\text{adv}}) \neq$
 170 y , while ensuring that $\|\delta\|_\infty \leq \epsilon$. Unlike CNNs that process local image regions hierarchically,
 171 Vision Transformers (Dosovitskiy et al., 2021) operate on a sequence of non-overlapping image
 172 patches. Given an input image \mathbf{x} , it is partitioned into $N = \frac{HW}{P^2}$ patches of size $P \times P$, each linearly
 173 projected to a D -dimensional embedding, resulting in tokens $\{\mathbf{z}_1, \dots, \mathbf{z}_N\} \subset \mathbb{R}^D$. A learnable
 174 classification token \mathbf{z}_{cls} is prepended, yielding a token sequence $\mathbf{Z}^{(0)} \in \mathbb{R}^{(N+1) \times D}$, which is enriched
 175 with positional encodings. ViTs consist of a stack of L transformer blocks. Each block contains a
 176 Multi-Head Self-Attention (MHSA) module and a Multi-Layer Perceptron (MLP) module, connected
 177 via residual connections and layer normalization (LN).

178

3.2 FEATURE-SENSITIVE GRADIENT SCALING (FSGS)

181 To improve transferability, we propose *Feature-Sensitive Gradient Scaling (FSGS)*, a fine-grained
 182 gradient modulation strategy that steers adversarial updates toward semantically relevant tokens while
 183 suppressing gradients associated with model-specific or noisy patterns. Unlike prior methods such as
 184 ATT (Ming et al., 2024) and TGR (Zhang et al., 2023), which rely on fixed truncation or uniform
 185 regularization, FSGS leverages intermediate transformer features to dynamically adjust gradient flow
 186 on a per-token basis.

187 **Limitations of Prior Gradient Modulation Approaches.** ATT weakens gradients across transformer
 188 modules based on empirical variance, but applies static masks that may disregard salient tokens. TGR
 189 promotes token-wise gradient uniformity without regard for token semantics, leading to potentially
 190 ineffective or redundant updates. In contrast, FSGS introduces adaptive scaling conditioned on the
 191 importance of each token, measured directly from the model’s internal activations. This content-aware
 192 reweighting enhances the alignment of perturbations with generalizable visual features and improves
 193 cross-architecture transfer.

194 **Why Token Activation Norm and Feature-Sensitive Gradient Scaling (FSGS)?** Token activation
 195 norms in Vision Transformers have been empirically shown to correlate with semantic saliency, with
 196 higher-norm tokens often corresponding to class-relevant features or foreground objects Kobayashi
 197 et al. (2020); Modarressi et al. (2022); Wu et al. (2024). Our Grad-CAM visualizations (Section 4.4)
 198 confirm this trend, showing strong alignment between high-norm tokens and semantically meaningful
 199 regions. This motivates using token norm as a saliency prior to guide adversarial perturbations. FSGS
 200 operationalizes this intuition by amplifying gradients from semantically important tokens while
 201 suppressing less informative ones. Importantly, not all layers benefit equally: early ViT layers capture
 202 low-level, architecture-dependent patterns (textures, positional cues) that hinder transfer, whereas
 203 deeper layers encode more robust, class-discriminative features Raghu et al. (2021); Bhojanapalli
 204 et al. (2021); Kim et al. (2024). To account for this, FSGS adopts a dual-stage strategy: in early
 205 layers, gradients are scaled by $(1 - \alpha)$ to downweight noisy signals, while in deeper layers, α is used
 206 to strengthen semantically aligned features. This design ensures perturbations are both semantically
 207 grounded and transferable across architectures, improving attack effectiveness in black-box settings
 (see Appendix A).

208 **Token-Level Importance Estimation.** Given a token embedding matrix $\mathbf{Z} \in \mathbb{R}^{T \times D}$, we estimate
 209 the importance of token i using the activation norm $\alpha_i = \|\mathbf{z}_i\|_2$, which serves as a proxy for semantic
 210 saliency. This assumption is supported by prior work in both NLP and vision (Kobayashi et al., 2020;
 211 Wu et al., 2024; Modarressi et al., 2022), which shows that activation magnitudes often correlate with
 212 token informativeness or attention saliency. For instance, Kobayashi et al. (2020) and Modarressi
 213 et al. (2022) argue that vector norms contribute substantially to a token’s influence, while Wu et al.
 214 (2024) highlight the role of transformed token magnitudes in ViT explanations. These scores are
 215 min-max normalized: $\hat{\alpha}_i = \frac{\alpha_i - \min_j \alpha_j}{\max_j \alpha_j - \min_j \alpha_j + \epsilon}$, where ϵ ensures numerical stability.

216 **Gradient Reweighting.** Each token’s gradient is modulated by a scaling factor: Let $l \in \{1, \dots, L\}$
 217 denote the index of the current transformer block, and let $\mathcal{E} \subset \{1, \dots, L\}$ be the set of early layers

(e.g., $\mathcal{E} = \{1, \dots, k\}$). Define an indicator function:

$$\beta^{(l)} = \begin{cases} 1 & \text{if } l \in \mathcal{E} \quad (\text{early layer}) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The final scaling factor for token i at layer l is then computed as: $s_i^{(l)} = \gamma_{\text{base}} + \lambda \cdot [(1 - \beta^{(l)}) \cdot \hat{\alpha}_i + \beta^{(l)} \cdot (1 - \hat{\alpha}_i)]$. And the FSGS-modulated gradient is: $\mathbf{g}_i^{(l), \text{FSGS}} = s_i^{(l)} \cdot \mathbf{g}_i^{(l)}$. Here, $\gamma_{\text{base}} \in (0, 1]$ ensures minimum gradient flow, while λ controls the suppression strength for less important tokens. This reweighting selectively amplifies gradients linked to semantically meaningful content. FSGS is applied independently to the QKV projections, attention weights, and MLP layers, using module-specific hyperparameters $\lambda_{\text{qkv}}, \lambda_{\text{attn}}, \lambda_{\text{mlp}}$, allowing tailored control over each component. FSGS is implemented via backward hooks, imposes negligible overhead, and integrates seamlessly with iterative attack frameworks. By aligning perturbations with high-importance regions, it enhances the semantic coherence and transferability of adversarial examples across both homogeneous and heterogeneous architectures.

3.3 SPECTRAL SMOOTHNESS REGULARIZATION (SSR)

We propose *Spectral Smoothness Regularization (SSR)* to suppress high-frequency perturbation artifacts that hinder cross-architecture transferability. At each PGD iteration, SSR applies a differentiable Gaussian blur to the adversarial input, enforcing a low-pass constraint on the evolving perturbation: $\mathbf{x}_{\text{adv}}^{\text{blur}} = \mathcal{G}_\sigma(\mathbf{x} + \delta)$, where δ is the perturbation and $\mathcal{G}_\sigma(\cdot)$ denotes Gaussian blur with standard deviation σ . The motivation follows from both signal processing and adversarial transferability studies: high-frequency perturbations often overfit surrogate-specific features and fail to generalize (Tsipras et al., 2019; Yin et al., 2019), whereas lower-frequency structures better align with perceptually salient, transferable patterns. Unlike input diversity approaches (Xie et al., 2019), which randomize input transformations, SSR directly regularizes the spectral content of the perturbation itself. It also differs from smoothing-based defenses, since the blur is applied *during optimization*, shaping the perturbation rather than post-processing it. SSR is lightweight, parameter-free, and compatible with any gradient-based attack. In practice, it synergizes with FSGS by reducing high-frequency noise while preserving semantically aligned gradients, leading to stronger transferability in both black-box and cross-architecture scenarios.

3.4 MODULE-WISE GRADIENT MODULATION

Vision Transformers differ from CNNs not only in architecture but also in how features and gradients evolve with depth. Prior studies (Ming et al., 2024; Yosinski et al., 2014; Naseer et al., 2022) have shown that deeper transformer layers tend to encode more specialized, model-specific patterns (particularly in the attention maps) which can harm the transferability of adversarial perturbations. To address this, we introduce a *Module-wise gradient modulation* strategy that suppresses unstable gradients in deep attention layers and softly attenuates the gradient flow in all modules (Attention, QKV, MLP) based on their layer depth. Inspired by ATT (Ming et al., 2024), our approach consists of two key components:

Selective Attention Truncation. We truncate the gradients flowing through the *Attention module* for deep transformer blocks beyond a fixed threshold l_{cut} , by setting their attention gradients to zero: $\mathbf{g}_l^{\text{attn}} \leftarrow \mathbb{1}_{[l < l_{\text{cut}}]} \cdot \mathbf{g}_l^{\text{attn}}$. This effectively disables attention backpropagation in deeper layers, mitigating overfitting to model-specific global patterns.

Module-Wise Gradient Weakening. For all layers $l \in \{1, \dots, L\}$ and modules $m \in \{\text{attn}, \text{qkv}, \text{mlp}\}$, we scale the gradients using a module-specific weakening factor $\omega^{(m)} \in (0, 1]$: $\mathbf{g}_l^{(m)} \leftarrow \omega^{(m)} \cdot \mathbf{g}_l^{(m)}$. This softly adjusts the contribution of each module based on its depth and functional role, before applying further refinement via FSGS. The weakening factors $\omega_m^{(l)}$ and the truncation layer threshold l_{cut} are predefined based on empirical sensitivity, further hyperparameter sensitivity studies are provided in Appendix D.

All gradient weakening and truncation operations are applied via backward hooks before the application of FSGS. This ordering ensures that noisy gradients are first suppressed or removed, and only the semantically meaningful signals are preserved and amplified by FSGS. Importantly, our method remains fully differentiable and does not alter the model’s forward pass, preserving compatibility

270 with any transformer backbone. The overall optimization algorithm and different hyper-parameters
 271 for training adversarial example are provided in Appendix B.
 272

273 **4 EXPERIMENTS**
 274

275 **4.1 EXPERIMENT SETUP**
 276

277 **Dataset.** Following prior works (Wei et al., 2022; Zhang et al., 2023; Ming et al., 2024), we randomly
 278 selected 1,000 clean images from the ILSVRC2012 validation set (Russakovsky et al., 2015), ensuring
 279 that all surrogate models correctly classify each image with high confidence. This selection facilitates
 280 a consistent and fair evaluation of transferability between models.

281 **Models.** We employ four representative Vision Transformer models as surrogate architectures:
 282 ViT-B/16 (Dosovitskiy et al., 2021), PiT-B (Heo et al., 2021), CaiT-S24 (Touvron et al., 2021b), and
 283 Visformer-S (Chen et al., 2021). To assess cross-architecture generalization, we group evaluation
 284 into two categories: ViT-to-ViT and ViT-to-CNN transfer. For ViT-to-ViT, we use four unseen target
 285 ViTs: DeiT-B (Touvron et al., 2021a), TNT-S (Han et al., 2021), LeViT-256 (Graham et al., 2021),
 286 and ConViT-B (d’Ascoli et al., 2021). For ViT-to-CNN, we evaluate against four deep CNN models:
 287 Inception-v3 (Inc-v3), Inception-v4 (Inc-v4), Inception-ResNet-v2 (IncRes-v2), and ResNet-v2-152
 288 (Res-v2) (Szegedy et al., 2016; 2017; He et al., 2016). Additionally, to evaluate robustness against
 289 adversarial defenses, we include three adversarially trained CNN models: Inc-v3-ens3, Inc-v4-ens4,
 290 and IncRes-v2-adv (Madry et al., 2018; Xu et al., 2022) and two adversarially trained ViTs: Swin-B
 291 (Mo et al., 2022) and XCiT-S (Debenedetti et al., 2023).

292 **Baselines.** We compare our method against a suite of strong baseline attacks. These include
 293 momentum- and variance-based methods such as MI-FGSM (MIM) (Dong et al., 2018b), VMI-
 294 FGSM (VMI) (Wang & He, 2021), and Skip Gradient Method (SGM) (Wu et al., 2020). We also
 295 include three state-of-the-art transformer-specific attacks: PNA (Wei et al., 2022), TGR (Zhang et al.,
 296 2023), and ATT (Ming et al., 2024), which incorporate attention structure or token-level heuristics
 297 into their gradient manipulation strategies. We also compare against diffusion-based attacks such as
 298 Diffattack (Chen et al., 2025).

299 **Evaluation Metrics.** We evaluate attack performance using the standard *Attack Success Rate* (ASR),
 300 defined as the proportion of adversarial examples that successfully fool the target model. Higher
 301 ASR (\uparrow) indicates stronger transferability.

302 **Parameter Settings.** All experiments use a maximum perturbation bound of $\epsilon = 16/255$, consistent
 303 with prior work (Zhang et al., 2023). The number of PGD iterations is set to $T = 10$, with a step size
 304 of $\eta = \epsilon/T = 1.6/255$. Momentum is used for stabilization with decay factor $\mu = 1.0$. Model- and
 305 method-specific hyperparameters follow their original settings unless otherwise stated. Input images
 306 are resized to 224×224 , and the patch size for transformer models is fixed at 16×16 . For spectral
 307 smoothness regularization, we apply Gaussian blur with fixed kernel size (3×3) and $\sigma = 0.5$. We
 308 set $\gamma_{\text{base}} = 0.5$. The weakening factors ω , layer truncation threshold l_{cut} , and the adaptive scaling
 309 factor to λ are tuned per model to balance the influence of QKV, Attention, and MLP gradients within
 310 the backward pass. The specific values of these hyperparameters are provided in Appendix B.

311 **4.2 EVALUATING THE TRANSFERABILITY**
 312

313 We evaluate the black-box transferability of adversarial examples generated by TESSER across ViTs,
 314 CNNs, and adversarially defended CNNs. Table 1 shows results when attacking ViTs using ViT-based
 315 surrogates. TESSER achieves an average ASR of **83.2%**, outperforming the strongest baseline
 316 (ATT) by **+5.8%** and DiffAttack by **+12.2%**. On CNN targets, where ViT-based attacks typically
 317 degrade, TESSER maintains strong performance with **74.4% ASR +10.9%** higher than ATT. This
 318 indicates that our semantic and frequency-aware perturbations generalize beyond transformer-specific
 319 structures. TESSER’s improvements are particularly notable on hybrid architectures like LeViT and
 320 ConViT, where both spatial alignment and cross-attention modeling are critical.

321 When facing adversarially trained CNNs (Table 3), TESSER achieves **53.55% ASR**, surpassing
 322 all baselines by a large margin. This suggests that TESSER generates perturbations that are not
 323 only transferable but also robust against strong defenses, an essential property for real-world attack
 324 scenarios. We also observe that the relative gains of TESSER vary across target types. For ViTs, the
 325 gains are moderate, likely because transformer-specific methods already perform reasonably well in
 326 this setting. However, the improvement is more pronounced on CNNs and defended CNNs, where

324
325
326
327 Table 1: The attack success rate (%) of various transfer-based attacks against eight ViT models and
328 the average attack success rate (%) of all black-box models. The best results are highlighted in **bold**.
329
330
331
332
333

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	MIM	100.0*	34.5	64.1	36.5	64.3	50.2	33.8	66.0	49.9
	VMI	99.6*	48.8	74.4	49.5	73.0	64.8	50.3	75.9	62.4
	SGM	100.0*	36.9	77.1	40.1	77.9	61.6	40.2	78.4	58.9
	PNA	100.0*	45.2	78.6	47.7	78.6	62.8	47.1	79.5	62.8
	TGR	100.0*	49.5	85.0	53.8	85.6	73.1	56.5	85.4	69.8
	DiffAttack	96.3*	60.1	70.4	63.3	75.4	71	57.5	74.36	71
	FPR	100*	37.7	77.9	40.0	77.2	74.3	42.1	78.0	65.9
	FPR + GRA	99.3*	61.7	88.0	65.2	87.3	86.8	65.5	89.5	80.41
PiT-B	ATT	99.9*	57.5	90.3	63.9	90.8	82.0	66.8	90.8	77.4
	Ours	100*	61.7	94	68.3	92.5	85.6	72.2	91.4	83.2↑
CaiT-S/24	MIM	24.7	100.0*	34.7	44.5	33.9	43.0	38.3	37.8	36.7
	VMI	38.9	99.7*	51.0	56.6	50.1	57.0	52.6	51.7	51.1
	SGM	41.8	100.0*	57.3	73.9	57.9	72.6	68.1	59.9	61.6
	PNA	47.9	100.0*	62.6	74.6	62.4	70.6	67.3	61.7	63.9
	TGR	60.3	100.0*	80.2	87.3	78.0	87.1	81.6	76.5	78.7
	ATT	69.6	100.0*	86.1	91.9	85.5	93.5	89.0	85.5	85.9
	Ours	74.9	100.0*	91.6	93.2	92.1	95	92.4	91.7	91.4↑
	MIM	70.9	54.8	99.8*	55.1	90.2	76.4	54.8	88.5	70.1
Visformer-S	VMI	76.3	63.6	98.8*	67.3	88.5	82.3	67.0	88.1	76.2
	SGM	86.0	55.8	100.0*	68.2	97.7	91.1	74.9	96.7	81.5
	PNA	82.4	60.7	99.7*	67.7	95.7	86.9	67.1	94.0	79.2
	TGR	88.2	66.1	100.0*	75.4	98.8	92.8	74.7	97.9	84.8
	ATT	93.6	76.4	100.0*	85.9	99.4	96.9	87.4	98.8	91.2
	Ours	95.2	81.4	100*	90.3	99.6	97.5	90.7	98.9	94.2↑
	MIM	28.1	50.4	41.0	99.9*	36.9	51.9	49.4	39.6	42.5
	VMI	39.2	60.0	56.6	100.0*	54.1	62.8	59.1	54.4	55.2
441	SGM	18.8	41.8	34.9	100.0*	31.2	52.1	52.7	29.5	37.3
	PNA	35.4	61.5	54.7	100.0*	51.0	66.3	64.5	50.7	54.9
	TGR	41.2	70.3	62.0	100.0*	59.5	74.7	74.8	56.2	62.7
	ATT	44.7	70.9	68.7	100.0*	66.4	78.8	80.9	58.4	67.0
	Ours	57.6	79.4	78.4	100.0*	75.9	83.2	85.3	69.6	78.7↑

347
348
349 ATT and TGR degrade significantly. This asymmetry suggests that our method is particularly effective
350 at bridging the architectural gap between transformer and non-transformer models. Furthermore,
351 TESSER’s performance is more stable across all target types, showing lower variance than competing
352 methods, which reinforces the robustness of our approach. Additional results and extended analysis
353 are presented in Appendix C, in addition to a comparison with AutoAttack (Appendix F) and targeted
354 attack evaluations (Appendix E).

355 We conducted additional experiments on robust ViT models trained via adversarial training with
356 $\epsilon = 4$, including Swin-B (Mo et al., 2022) and XCiT-S (Debenedetti et al., 2023). We compared
357 TESSER against state-of-the-art attacks (PNA+PO, TGR+PO, and ATT+SPPO) using their optimal
358 hyperparameters. As shown in Table 2, TESSER consistently achieves the highest ASR on both robust
359 and corresponding standard ViT models, confirming its strong effectiveness even under adversarial
360 defense settings. These results demonstrate that TESSER’s transferability extends to robust ViTs, not
361 just CNNs and hybrids.

363 4.3 ABLATION ON MODULE-WISE GRADIENT MODULATION

364
365
366 Table 2: The attack success rate (%) of various
367 transfer-based attacks against robust ViTs. The
368 best results are highlighted in **bold**.
369

Model	Attack	Robust ViTs		Normal ViTs	
		Swin-B	Xcit-S	Swin-B	Xcit-S
ViT-B/16	clean	5.4	46.8	0.4	0.2
	PNA+PO	8.8	51.7	47.5	45.5
	TGR+PO	15.8	56.5	54.4	54.5
	ATT+SPPO	16.9	56.7	70.4	68.6
	TESSER	29.7↑	70.8↑	99.9↑	77.9↑
PiT-B	PNA+PO	9.2	51.8	67.0	71.2
	TGR+PO	17.9	58.2	77.3	80.7
	ATT+SPPO	18.7	58.3	90.4	92.8
	TESSER	31.9↑	71.6↑	100↑	95.4↑

To understand the individual and combined contributions of our gradient modulation strategy across different transformer modules, we conduct an ablation study by selectively applying Feature-Sensitive Gradient Scaling to the Attention, QKV, and MLP components. Table 4 presents the attack success rates (ASR) on ViT-based models, CNNs, and defended CNNs under different configurations. When FSGS is applied to a single module, the Attention pathway contributes the most to transferability, particularly for ViTs, achieving an ASR of 80.1%. MLP-only and QKV-only configurations also yield strong improvements over the baseline, with notable gains on CNNs and defended mod-

378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
 432
 433
 434
 435
 436
 437
 438
 439
 440
 441
 442
 443
 444
 445
 446
 447
 448
 449
 450
 451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494
 495
 496
 497
 498
 499
 500
 501
 502
 503
 504
 505
 506
 507
 508
 509
 510
 511
 512
 513
 514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539
 540
 541
 542
 543
 544
 545
 546
 547
 548
 549
 550
 551
 552
 553
 554
 555
 556
 557
 558
 559
 560
 561
 562
 563
 564
 565
 566
 567
 568
 569
 570
 571
 572
 573
 574
 575
 576
 577
 578
 579
 580
 581
 582
 583
 584
 585
 586
 587
 588
 589
 590
 591
 592
 593
 594
 595
 596
 597
 598
 599
 600
 601
 602
 603
 604
 605
 606
 607
 608
 609
 610
 611
 612
 613
 614
 615
 616
 617
 618
 619
 620
 621
 622
 623
 624
 625
 626
 627
 628
 629
 630
 631
 632
 633
 634
 635
 636
 637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647
 648
 649
 650
 651
 652
 653
 654
 655
 656
 657
 658
 659
 660
 661
 662
 663
 664
 665
 666
 667
 668
 669
 670
 671
 672
 673
 674
 675
 676
 677
 678
 679
 680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701
 702
 703
 704
 705
 706
 707
 708
 709
 710
 711
 712
 713
 714
 715
 716
 717
 718
 719
 720
 721
 722
 723
 724
 725
 726
 727
 728
 729
 730
 731
 732
 733
 734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755
 756
 757
 758
 759
 760
 761
 762
 763
 764
 765
 766
 767
 768
 769
 770
 771
 772
 773
 774
 775
 776
 777
 778
 779
 780
 781
 782
 783
 784
 785
 786
 787
 788
 789
 790
 791
 792
 793
 794
 795
 796
 797
 798
 799
 800
 801
 802
 803
 804
 805
 806
 807
 808
 809
 810
 811
 812
 813
 814
 815
 816
 817
 818
 819
 820
 821
 822
 823
 824
 825
 826
 827
 828
 829
 830
 831
 832
 833
 834
 835
 836
 837
 838
 839
 840
 841
 842
 843
 844
 845
 846
 847
 848
 849
 850
 851
 852
 853
 854
 855
 856
 857
 858
 859
 860
 861
 862
 863
 864
 865
 866
 867
 868
 869
 870
 871
 872
 873
 874
 875
 876
 877
 878
 879
 880
 881
 882
 883
 884
 885
 886
 887
 888
 889
 890
 891
 892
 893
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917
 918
 919
 920
 921
 922
 923
 924
 925
 926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971
 972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025
 1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079
 1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187
 1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619
 1620
 1621
 1622
 1623
 1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639
 1640
 1641
 1642
 1643
 1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673
 1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727
 1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781
 1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835
 1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870

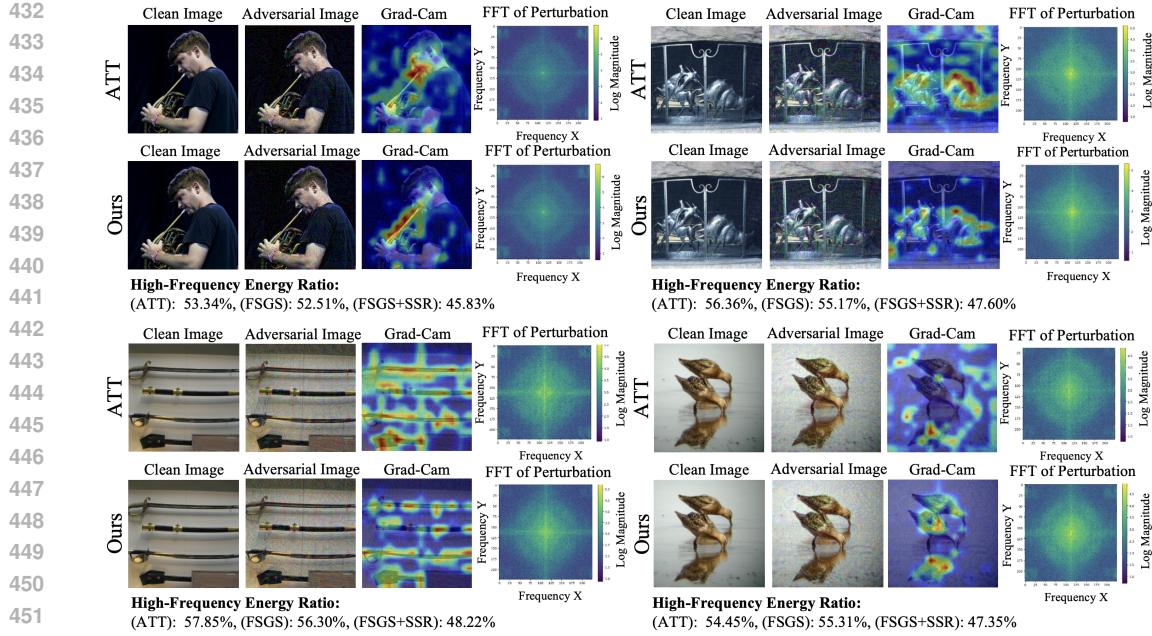


Figure 2: **Qualitative and frequency-domain comparison between ATT and our method (FSGS and FSGS + SSR).** Each row shows clean images, adversarial examples when using FSGS, Grad-CAM (guided by the adversarial label) overlays, and FFT log-magnitude spectra when using SSR. Our method produces perturbations that better align with semantically relevant regions and exhibit smoother frequency profiles. Further results and analysis are provided in Appendix D.

comparison using LPIPS, SSIM, and PSNR across TESSER and transfer-based attacks such as ATT and TGR. As shown in Table 5, TESSER achieves significantly higher imperceptibility, with 50% reduction in LPIPS, 33% improvement in SSIM, and +5 dB increase in PSNR, demonstrating strong stealthiness without sacrificing effectiveness.

4.5 SPECTRAL SMOOTHNESS EVALUATION VIA FREQUENCY-DOMAIN ANALYSIS

To quantitatively assess the effect of Spectral Smoothness Regularization (SSR), we conduct a frequency-domain analysis of the generated perturbations. Specifically, we compute the 2D Fast Fourier Transform (FFT) of each perturbation and evaluate the *high-frequency energy ratio*, defined as the proportion of energy outside the central low-frequency band in the log-magnitude spectrum (as shown in Figure 2). Given a perturbation $\delta \in \mathbb{R}^{3 \times H \times W}$, we compute its FFT, shift the spectrum to center the low frequencies, and apply a radial mask to isolate high-frequency components.

This experiment is repeated on a batch of adversarial samples to compare the spectral concentration of different attack variants. Our results demonstrate that SSR substantially reduces the high-frequency energy of perturbations. Across examples, ATT shows the highest high-frequency ratios (e.g., 53–56%), while FSGS reduces this moderately (~52–55%). When combined with SSR, the high-frequency ratio drops further (to ~45–47%), indicating smoother and more transferable perturbations. This confirms that SSR encourages low-frequency perturbation structure, complementing the token-aware gradient modulation of FSGS.

Table 5: Stealth Evaluation of Transfer-Based Attacks.

Metric	TGR	ATT	TESSER
LPIPS \downarrow	0.35	0.42	0.21
SSIM \uparrow	0.66	0.57	0.77
PSNR \uparrow	22.23 dB	19.70 dB	25.04 dB

486 **5 CONCLUSION**
 487

488 We proposed TESSER, a unified adversarial attack framework designed to improve transferability
 489 across diverse model architectures. By integrating Feature-Sensitive Gradient Scaling (FSGS) and
 490 Spectral Smoothness Regularization (SSR), TESSER guides adversarial gradients through seman-
 491 tically meaningful token activations and enforces smooth, low-frequency perturbation structures.
 492 Combined with layer and module-wise gradient modulation, our method effectively mitigates overfit-
 493 ting to model-specific representations and enhances generalization to unseen targets. Experimental
 494 results across a wide range of ViTs, hybrid models, and CNNs demonstrate that TESSER consistently
 495 outperforms state-of-the-art transfer attacks in both accuracy degradation and optimization efficiency.

496
 497 **REFERENCES**
 498

499 Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *Annual*
 500 *Meeting of the Association for Computational Linguistics*, 2020. URL <https://api.semanticscholar.org/CorpusID:218487351>.

502 Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. In *Proceedings*
 503 *of the IEEE/CVF international conference on computer vision*, pp. 10231–10241, 2021.

506 Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey
 507 Zagoruyko. End-to-end object detection with transformers. In *European conference on computer*
 508 *vision*, pp. 213–229. Springer, 2020.

509 Jianqi Chen, Hao Chen, Keyan Chen, Yilan Zhang, Zhengxia Zou, and Zhenwei Shi. Diffusion models
 510 for imperceptible and transferable adversarial attack. *IEEE Transactions on Pattern Analysis and*
 511 *Machine Intelligence*, 47(2):961–977, 2025. doi: 10.1109/TPAMI.2024.3480519.

513 Zhengsu Chen, Lingxi Xie, Jianwei Niu, Xuefeng Liu, Longhui Wei, and Qi Tian. Visformer: The
 514 vision-friendly transformer. In *Proceedings of the IEEE/CVF international conference on computer*
 515 *vision*, pp. 589–598, 2021.

516 Edoardo DeBenedetti, Vikash Sehwag, and Prateek Mittal. A light recipe to train robust vision
 517 transformers. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*,
 518 pp. 225–253. IEEE, 2023.

520 Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting
 521 adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision*
 522 *and pattern recognition*, pp. 9185–9193, 2018a.

524 Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting
 525 adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision*
 526 *and pattern recognition*, pp. 9185–9193, 2018b.

527 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 528 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,
 529 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale,
 530 2021. URL <https://arxiv.org/abs/2010.11929>.

531 Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun.
 532 Convit: Improving vision transformers with soft convolutional inductive biases. In *International*
 533 *conference on machine learning*, pp. 2286–2296. PMLR, 2021.

535 Aditya Ganeshan, Vivek BS, and R Venkatesh Babu. Fda: Feature disruptive attack. In *Proceedings*
 536 *of the IEEE/CVF International Conference on Computer Vision*, pp. 8069–8079, 2019.

538 Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial
 539 examples. *CoRR*, abs/1412.6572, 2014. URL <https://api.semanticscholar.org/CorpusID:6706414>.

540 Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou,
 541 and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In
 542 *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12259–12269,
 543 2021.

544 Amira Guesmi, Muhammad Abdullah Hanif, Bassem Ouni, and Muhammad Shafique. Physical
 545 adversarial attacks for camera-based smart systems: Current trends, categorization, applications,
 546 research challenges, and future outlook. *IEEE Access*, 11:109617–109668, 2023. doi: 10.1109/
 547 ACCESS.2023.3321118.

548 Amira Guesmi, Ruitian Ding, Muhammad Abdullah Hanif, Ihsen Alouani, and Muhammad Shafique.
 549 Dap: A dynamic adversarial patch for evading person detectors. In *Proceedings of the IEEE/CVF
 550 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24595–24604, June 2024a.

551 Amira Guesmi, Muhammad Abdullah Hanif, Ihsen Alouani, Bassem Ouni, and Muhammad Shafique.
 552 Ssap: A shape-sensitive adversarial patch for comprehensive disruption of monocular depth
 553 estimation in autonomous navigation applications. In *2024 IEEE/RSJ International Conference on
 554 Intelligent Robots and Systems (IROS)*, pp. 2786–2793, 2024b. doi: 10.1109/IROS58592.2024.
 555 10802252.

556 Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in
 557 transformer. *Advances in neural information processing systems*, 34:15908–15919, 2021.

558 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 559 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 560 pp. 770–778, 2016.

561 Byeongho Heo, Sangdoo Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon
 562 Oh. Rethinking spatial dimensions of vision transformers. In *Proceedings of the IEEE/CVF
 563 international conference on computer vision*, pp. 11936–11945, 2021.

564 Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing
 565 adversarial example transferability with an intermediate level attack. In *Proceedings of the
 566 IEEE/CVF international conference on computer vision*, pp. 4733–4742, 2019.

567 Gihyun Kim, Juyeop Kim, and Jong-Seok Lee. Exploring adversarial robustness of vision transform-
 568 ers in the spectral perspective. In *Proceedings of the IEEE/CVF Winter Conference on Applications
 569 of Computer Vision*, pp. 3976–3985, 2024.

570 Goro Kobayashi, Tatsuki Kurabayashi, Sho Yokoi, and Kentaro Inui. Attention is not only a weight:
 571 Analyzing transformers with vector norms. In *Conference on Empirical Methods in Natural
 572 Language Processing*, 2020. URL [https://api.semanticscholar.org/CorpusID:
 573 222176890](https://api.semanticscholar.org/CorpusID:222176890).

574 Xiao Lin and Devi Parikh. Leveraging visual question answering for image-caption ranking. In
 575 *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October
 576 11–14, 2016, Proceedings, Part II* 14, pp. 261–277. Springer, 2016.

577 Fan Ma, Mike Zheng Shou, Linchao Zhu, Haoqi Fan, Yilei Xu, Yi Yang, and Zhicheng Yan. Unified
 578 transformer tracker for object tracking. In *Proceedings of the IEEE/CVF conference on computer
 579 vision and pattern recognition*, pp. 8781–8790, 2022.

580 Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu.
 581 Towards deep learning models resistant to adversarial attacks. In *6th International Conference on
 582 Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference
 583 Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=rJzIBfZAb>.

584 Di Ming, Peng Ren, Yunlong Wang, and Xin Feng. Boosting the transferability of adversarial attack
 585 on vision transformer with adaptive token tuning. *Advances in Neural Information Processing
 586 Systems*, 37:20887–20918, 2024.

594 Yichuan Mo, Dongxian Wu, Yifei Wang, Yiwen Guo, and Yisen Wang. When adversarial training
 595 meets vision transformers: Recipes from training to architecture. *Advances in Neural Information
 596 Processing Systems*, 35:18599–18611, 2022.

597

598 Ali Modarressi, Mohsen Fayyaz, Yadollah Yaghoobzadeh, and Mohammad Taher Pilehvar. Globenc:
 599 Quantifying global token attribution by incorporating the whole encoder layer in transformers. In
 600 Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruiz (eds.), *Proceedings
 601 of the 2022 Conference of the North American Chapter of the Association for Computational
 602 Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15,
 603 2022*, pp. 258–271. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.
 604 NAACL-MAIN.19. URL <https://doi.org/10.18653/v1/2022.nacl-main.19>.

605 Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On
 606 improving adversarial transferability of vision transformers. In *The Tenth International Conference
 607 on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
 608 URL <https://openreview.net/forum?id=D6nH3719vZy>.

609 Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy.
 610 Do vision transformers see like convolutional neural networks? *Advances in neural information
 611 processing systems*, 34:12116–12128, 2021.

612

613 Yuchen Ren, Zhengyu Zhao, Chenhao Lin, Bo Yang, Lu Zhou, Zhe Liu, and Chao
 614 Shen. Improving adversarial transferability on vision transformers via forward prop-
 615 agation refinement. In *IEEE/CVF Conference on Computer Vision and Pattern
 616 Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pp. 25071–25080.
 617 Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.02334.
 618 URL https://openaccess.thecvf.com/content/CVPR2025/html/Ren_Improving_Adversarial_Transferability_on_Vision_Transformers_via_Forward_Propagation_Refinement_CVPR_2025_paper.html.

619

620 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang,
 621 Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition
 622 challenge. *International journal of computer vision*, 115:211–252, 2015.

623

624 Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking
 625 the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer
 626 vision and pattern recognition*, pp. 2818–2826, 2016.

627

628 Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-
 629 resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference
 630 on artificial intelligence*, volume 31, 2017.

631

632 Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé
 633 Jégou. Training data-efficient image transformers & distillation through attention. In *International
 634 conference on machine learning*, pp. 10347–10357. PMLR, 2021a.

635

636 Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going
 637 deeper with image transformers. In *Proceedings of the IEEE/CVF international conference on
 638 computer vision*, pp. 32–42, 2021b.

639

640 Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry.
 641 Robustness may be at odds with accuracy. In *7th International Conference on Learning Re-
 642 presentations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL
 643 <https://openreview.net/forum?id=SyxAb30cY7>.

644

645 Xiaosen Wang and Kun He. Enhancing the transferability of adversarial attacks through variance
 646 tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 647 pp. 1924–1933, 2021.

648

649 Zhibo Wang, Hengchang Guo, Zhifei Zhang, Wenxin Liu, Zhan Qin, and Kui Ren. Feature importance-
 650 aware transferable adversarial attacks. In *Proceedings of the IEEE/CVF international conference
 651 on computer vision*, pp. 7639–7648, 2021.

648 Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang.
 649 Towards transferable adversarial attacks on vision transformers. In *Proceedings of the AAAI*
 650 *Conference on Artificial Intelligence*, volume 36, pp. 2668–2676, 2022.

651

652 Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections mat-
 653 ter: On the transferability of adversarial examples generated with resnets. In *8th International*
 654 *Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
 655 OpenReview.net, 2020. URL <https://openreview.net/forum?id=BJ1Rs34Fvr>.

656 Junyi Wu, Bin Duan, Weitai Kang, Hao Tang, and Yan Yan. Token transformation matters: Towards
 657 faithful post-hoc explanation for vision transformer. In *Proceedings of the IEEE/CVF Conference*
 658 *on Computer Vision and Pattern Recognition*, pp. 10926–10935, 2024.

659 Wang Xiaosen, Kangheng Tong, and Kun He. Rethinking the backward propagation for adversarial
 660 transferability. *Advances in Neural Information Processing Systems*, 36:1905–1922, 2023.

661

662 Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille.
 663 Improving transferability of adversarial examples with input diversity. In *Proceedings of the*
 664 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 2730–2739, 2019.

665 Yifeng Xiong, Jiadong Lin, Min Zhang, John E Hopcroft, and Kun He. Stochastic variance reduced
 666 ensemble adversarial attack for boosting the adversarial transferability. In *Proceedings of the*
 667 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 14983–14992, 2022.

668

669 Zhuoer Xu, Guanghui Zhu, Changhua Meng, Zhenzhe Ying, Weiqiang Wang, Ming Gu, Yihua Huang,
 670 et al. A2: Efficient automated attacker for boosting adversarial training. *Advances in Neural*
 671 *Information Processing Systems*, 35:22844–22855, 2022.

672 Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier
 673 perspective on model robustness in computer vision. *Advances in Neural Information Processing*
 674 *Systems*, 32, 2019.

675

676 Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep
 677 neural networks? *Advances in neural information processing systems*, 27, 2014.

678 Jianping Zhang, Yizhan Huang, Weibin Wu, and Michael R Lyu. Transferable adversarial attacks on
 679 vision transformers with token gradient regularization. In *Proceedings of the IEEE/CVF conference*
 680 *on computer vision and pattern recognition*, pp. 16415–16424, 2023.

681 Fangrui Zhu, Yi Zhu, Li Zhang, Chongruo Wu, Yanwei Fu, and Mu Li. A unified efficient pyramid
 682 transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference*
 683 *on Computer Vision*, pp. 2667–2677, 2021.

684

685 Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision
 686 mamba: Efficient visual representation learning with bidirectional state space model. In *Forty-first*
 687 *International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*.
 688 OpenReview.net, 2024. URL <https://openreview.net/forum?id=YbHCqn4qF4>.

689

690

691

692

693

694

695

696

697

698

699

700

701

APPENDIX

A THEORETICAL JUSTIFICATION OF FEATURE-SENSITIVE GRADIENT SCALING (FSGS)

We provide a formal argument to support the hypothesis that modulating gradients based on token-level activation norms enhances adversarial transferability. Our analysis is grounded in the relationship between semantic informativeness and gradient alignment across models.

A.1 PRELIMINARIES

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}^K$ be a surrogate classifier and $f' : \mathbb{R}^d \rightarrow \mathbb{R}^K$ a target (black-box) classifier. An adversarial perturbation $\delta \in \mathbb{R}^d$ is added to input \mathbf{x} such that $\|\delta\|_\infty \leq \epsilon$ and $f(\mathbf{x} + \delta) \neq y$.

Assume \mathbf{x} is decomposed into T tokens with embeddings $\mathbf{z}_1, \dots, \mathbf{z}_T \in \mathbb{R}^D$. Denote the loss gradients w.r.t. each token as $\mathbf{g}_i = \nabla_{\mathbf{z}_i} \mathcal{L}(f(\mathbf{x}), y)$, and similarly for f' .

A.2 SEMANTIC TOKENS AND GRADIENT ALIGNMENT

Let $S_{\text{sem}} \subseteq \{1, \dots, T\}$ be the set of semantically informative tokens (e.g., foreground object, discriminative parts). Let S_{bg} be its complement (background or irrelevant tokens).

We define the inter-model gradient alignment at token i as:

$$\text{Align}_i = \cos \theta_i = \frac{\langle \nabla_{\mathbf{z}_i} \mathcal{L}_f, \nabla_{\mathbf{z}_i} \mathcal{L}_{f'} \rangle}{\|\nabla_{\mathbf{z}_i} \mathcal{L}_f\| \cdot \|\nabla_{\mathbf{z}_i} \mathcal{L}_{f'}\|}$$

Assumption 1. *Gradients at semantically important tokens exhibit higher cross-model alignment:*

$$\mathbb{E}_{i \in S_{\text{sem}}} [\text{Align}_i] > \mathbb{E}_{i \in S_{\text{bg}}} [\text{Align}_i]$$

This is supported by empirical findings in model interpretability (Abnar & Zuidema, 2020; Lin & Parikh, 2016; Raghu et al., 2021) and our own Grad-CAM visualizations (see Section 4.4).

A.3 FEATURE-SENSITIVE GRADIENT SCALING (FSGS)

FSGS assigns a scaling factor s_i to each token based on its activation norm $\alpha_i = \|\mathbf{z}_i\|_2$:

$$s_i = \gamma_{\text{base}} + \lambda(1 - \hat{\alpha}_i), \quad \hat{\alpha}_i = \frac{\alpha_i - \min_j \alpha_j}{\max_j \alpha_j - \min_j \alpha_j + \varepsilon}$$

Tokens with high α_i (assumed to lie in S_{sem}) receive larger gradients, while low-importance tokens are suppressed.

Theorem 1 (FSGS Improves Expected Gradient Alignment). *Let $G = \sum_{i=1}^T \mathbf{g}_i$ be the unscaled gradient and $G_{\text{FSGS}} = \sum_{i=1}^T s_i \cdot \mathbf{g}_i$ the FSGS-scaled gradient. Under Assumption 1, the cosine alignment between G_{FSGS} and the target model's gradient G' satisfies:*

$$\cos \theta(G_{\text{FSGS}}, G') > \cos \theta(G, G')$$

Sketch. We decompose the total gradient into two subsets:

$$G = \sum_{i \in S_{\text{sem}}} \mathbf{g}_i + \sum_{i \in S_{\text{bg}}} \mathbf{g}_i$$

FSGS scales $i \in S_{\text{sem}}$ by higher s_i than those in S_{bg} , thus:

$$G_{\text{FSGS}} = \sum_{i \in S_{\text{sem}}} s_i \mathbf{g}_i + \sum_{i \in S_{\text{bg}}} s_i \mathbf{g}_i$$

Since $\mathbb{E}_{i \in S_{\text{sem}}} [\text{Align}_i] > \mathbb{E}_{i \in S_{\text{bg}}} [\text{Align}_i]$, amplifying contributions from S_{sem} increases the expected alignment between G_{FSGS} and G' . Therefore:

$$\cos \theta(G_{\text{FSGS}}, G') > \cos \theta(G, G') \quad (\text{by Jensen's inequality over positively weighted aligned vectors})$$

□

756 A.4 IMPLICATION
757

758 Theorem 1 provides theoretical support for the design of FSGS: boosting gradients from semantically
759 salient tokens leads to improved alignment with gradients from unseen models, thereby enhancing
760 adversarial transferability. This also explains the empirical advantage of FSGS in black-box settings,
761 particularly when transferring from ViTs to CNNs or hybrid models.

762
763 B THE OVERALL FRAMEWORK OF OPTIMIZATION ALGORITHM
764765 B.1 ALGORITHM
766

767 We now present the full optimization framework used to generate adversarial examples in our method.
768 Our algorithm builds on the momentum-based PGD attack (Dong et al., 2018a), and integrates three
769 coordinated components: (1) *Module and Layer-Wise Gradient Modulation* to adjust the contribution
770 of each transformer layer and suppress noisy deep-layer gradients, (2) *Feature-Sensitive Gradient
771 Scaling (FSGS)* to selectively enhance semantically important token gradients, and (3) *Spectral
772 Smoothness Regularization (SSR)* to constrain the perturbation’s frequency content.

773 Let $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ be a clean input, $y \in \{1, \dots, K\}$ its ground-truth label, and f the surrogate
774 model. We seek a perturbation δ satisfying $\|\delta\|_\infty \leq \epsilon$, such that the adversarial input $\mathbf{x}^{\text{adv}} = \mathbf{x} + \delta$
775 misleads f and transfers effectively to other black-box models.

776
777 OPTIMIZATION PROCEDURE

778 The perturbation is optimized over T steps using projected gradient descent with momentum. At each
779 step $t \in \{1, \dots, T\}$, the perturbed input is smoothed using a differentiable Gaussian blur operator:

$$781 \quad \mathbf{x}^{(t)} = \mathcal{G}_\sigma(\mathbf{x} + \delta^{(t-1)})$$

782 where $\mathcal{G}_\sigma(\cdot)$ denotes Gaussian blurring with standard deviation σ , enforcing low-frequency spectral
783 structure (SSR).

784 Next, the gradient of the loss is computed:

$$786 \quad \mathbf{g}^{(t)} = \nabla_{\mathbf{x}} \mathcal{L}(f(\mathbf{x}^{(t)}), y)$$

788 This gradient is intercepted via backward hooks at key ViT modules (Attention, QKV, MLP). For
789 each transformer block l , the following sequence is applied to each module:

- 791 1. **Module-wise Weakening:** The gradient $\mathbf{g}^{(l)}$ for each module is first scaled using a module-
792 specific weakening factor $\omega^{(l)} \in (0, 1]$ (e.g., $\omega_{\text{attn}}^{(l)}$, $\omega_{\text{qkv}}^{(l)}$, $\omega_{\text{mlp}}^{(l)}$). This captures prior knowledge
793 about the sensitivity of each module.
- 794 2. **Layer-wise Modulation:** The weakened attention gradient is then further modulated by
795 a layer-specific coefficient $\tau_l \in [0, 1]$, which reduces the influence of deeper transformer
796 layers:

$$797 \quad \mathbf{g}^{(l)} \leftarrow \tau_l \cdot (\omega^{(l)} \cdot \mathbf{g}^{(l)})$$

- 798 3. **Feature-Sensitive Gradient Scaling (FSGS):**

799 A layer-aware gradient modulation mechanism that scales gradients based on token-wise
800 importance. FSGS promotes perturbation alignment with semantically salient features while
801 suppressing low-level, architecture-specific signals that degrade cross-model transferability.
802 Let $\mathbf{Z} \in \mathbb{R}^{T \times D}$ be the token embedding matrix at a given transformer block. We define
803 the raw importance score of token i as: $\alpha_i = \|\mathbf{z}_i\|_2$. The importance scores are min-max
804 normalized across tokens: $\hat{\alpha}_i = \frac{\alpha_i - \min_j \alpha_j}{\max_j \alpha_j - \min_j \alpha_j + \varepsilon}$, where ε is a small constant to avoid
805 division by zero.

806 Let $l \in \{1, \dots, L\}$ denote the index of the current transformer block, and let $\mathcal{E} \subset \{1, \dots, L\}$
807 be the set of early layers (e.g., $\mathcal{E} = \{1, \dots, k\}$). Define an indicator function:

$$809 \quad \beta^{(l)} = \begin{cases} 1 & \text{if } l \in \mathcal{E} \quad (\text{early layer}) \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

810 The final scaling factor for token i at layer l is then computed as: $s_i^{(l)} = \gamma_{\text{base}} + \lambda \cdot$
 811 $[(1 - \beta^{(l)}) \cdot \hat{\alpha}_i + \beta^{(l)} \cdot (1 - \hat{\alpha}_i)]$. And the FSGS-modulated gradient is: $\mathbf{g}_i^{(l),\text{FSGS}} = s_i^{(l)} \cdot$
 812 $\mathbf{g}_i^{(l)}$
 813

814 All module gradients are aggregated to form the total input gradient $\mathbf{g}^{(t)}$, and momentum is applied:
 815

$$816 \mathbf{m}^{(t)} = \mu \cdot \mathbf{m}^{(t-1)} + \frac{\mathbf{g}^{(t)}}{\|\mathbf{g}^{(t)}\|_1}$$

817 The perturbation is updated and projected onto the ℓ_∞ -norm ball:
 818

$$819 \delta^{(t)} = \text{Clip}_\epsilon \left(\delta^{(t-1)} + \eta \cdot \text{sign}(\mathbf{m}^{(t)}) \right)$$

820 **Description.** Algorithm 1 summarizes our full optimization loop. The key innovation lies in the
 821 sequential application of module-wise weakening, layer-wise modulation, and semantic-aware scaling
 822 via FSGS. All components are implemented via backward hooks, ensuring compatibility with any
 823 transformer-based model.
 824

825 **Hyper-parameters.** Table 6 summarizes the model-specific hyperparameters used in TESSER for
 826 each architecture. These include module-wise weakening factors ($\omega^{(\cdot)}$), FSGS scaling parameters
 827 (λ), spectral smoothness regularization strength (σ), attention truncation depth (l_{cut}), base gradient
 828 scaling (γ_{base}), as well as optimization parameters: momentum decay (μ) and step size (η). Values
 829 are carefully selected to balance gradient modulation and attack stability per architecture.
 830

831 B.2 COMPUTATIONAL COST

832 To evaluate the efficiency of our proposed method, we report the average time (in seconds) required to
 833 generate a single adversarial example using FSGS, FSGS+SSR, and the ATT across different models.
 834 As shown in Table 7, our methods incur minimal overhead compared to ATT, with only a slight
 835 increase when applying SSR. In particular, even in deeper architectures like CaiT-S/24, FSGS+SSR
 836 remains significantly more efficient than ATT.
 837

838 We provide the detailed environment configuration used for all evaluations. All experiments were
 839 conducted on NVIDIA Tesla T4 GPUs hosted on Google Colab. We present the key software
 840 dependencies and their corresponding versions:
 841

- 842 • Python: 3.11.12
- 843 • PyTorch: 2.6.0
- 844 • Torchvision: 0.21.0
- 845 • NumPy: 2.0.2
- 846 • Pillow: 11.2.1
- 847 • Timm: 1.0.15
- 848 • SciPy: 1.15.3

849 C ADDITIONAL EXPERIMENTS

850 C.1 QUANTITATIVE ANALYSIS FOR SSR

851 To evaluate the effect of spectral regularization strength, we conduct an ablation study by varying
 852 the Gaussian blur standard deviation σ used in Spectral Smoothness Regularization (SSR). Table 8
 853 reports the average attack success rates (ASR) on ViTs, CNNs, and defended CNNs for $\sigma \in$
 854 $\{0.5, 0.6, 0.7, 0.8\}$ across all surrogate models.
 855

856 We observe a consistent trend: increasing σ improves transferability to CNNs and defended CNNs,
 857 while slightly reducing ASR on ViTs. This trade-off reflects the role of SSR in suppressing high-
 858 frequency architecture-specific noise: improving cross-architecture generalization but marginally
 859

864
865
866
867
868

Algorithm 1: TESSER: Transfer-Enhancing Adversarial Optimization from Vision Transformers via Spectral and Semantic Regularization

869 **Input:** Input image \mathbf{x} , label y , model f ,
 870 Steps T , step size η , perturbation bound ϵ ,
 871 Gaussian blur \mathcal{G}_σ , momentum μ ,
 872 Base scale γ_{base} ,
 873 Module-specific FSGS strengths $\lambda_{\text{qkv}}, \lambda_{\text{attn}}, \lambda_{\text{mlp}}$,
 874 Early-layer set \mathcal{E} , attention cutoff layer l_{cut} ,
 875 Module weakening factors $\omega^{(l)}$,
 876 SSR loss function \mathcal{L}_{SSR}
 877 **Output:** Adversarial example \mathbf{x}^{adv}
 878 **Initialize:** $\delta^{(0)} = 0, \mathbf{m}^{(0)} = 0$
 879 **for** $t = 1$ **to** T **do**
 880 1. **Apply SSR:**
 881 $\mathbf{x}^{(t)} = \mathcal{G}_\sigma(\mathbf{x} + \delta^{(t-1)})$
 882 2. **Forward pass and compute classification loss:**
 883 $\mathcal{L}_{\text{cls}}^{(t)} = \mathcal{L}(f(\mathbf{x}^{(t)}), y)$
 884 3. **Backward pass with hooks at QKV, Attention, and MLP modules:**
 885 **foreach** block $l \in \{1, \dots, L\}$ **do**
 886 **foreach** module $m \in \{\text{qkv}, \text{attn}, \text{mlp}\}$ **do**
 887 3.1 **Extract token features and gradients:**
 888 $\mathbf{Z}^{(l,m)} = [\mathbf{z}_1^{(l,m)}, \dots, \mathbf{z}_T^{(l,m)}]$
 889 $\mathbf{G}^{(l,m)} = [\mathbf{g}_1^{(l,m)}, \dots, \mathbf{g}_T^{(l,m)}]$
 890 3.2 **Compute token importance:**
 891 $\alpha_i = \|\mathbf{z}_i^{(l,m)}\|_2, \hat{\alpha}_i = \frac{\alpha_i - \min_j \alpha_j}{\max_j \alpha_j - \min_j \alpha_j + \epsilon}$
 892 3.3 **Apply module-wise weakening:**
 893 $\mathbf{G}^{(l,m)} \leftarrow \omega^{(l)} \cdot \mathbf{G}^{(l,m)}$
 894 3.4 **Selective Attention Truncation (only if $m = \text{attn}$):**
 895 **if** $l \geq l_{\text{cut}}$ **then**
 896 $\mathbf{G}^{(l,\text{attn})} \leftarrow 0$
 897 **end**
 898 3.5 **Compute FSGS scaling:**
 899 **foreach** token $i \in \{1, \dots, T\}$ **do**
 900 **if** $l \in \mathcal{E}$ **then**
 901 $s_i = \gamma_{\text{base}} + \lambda_m \cdot (1 - \hat{\alpha}_i)$
 902 **end**
 903 **else**
 904 $s_i = \gamma_{\text{base}} + \lambda_m \cdot \hat{\alpha}_i$
 905 **end**
 906 $\mathbf{g}_i^{(l,m)} \leftarrow s_i \cdot \mathbf{g}_i^{(l,m)}$
 907 **end**
 908 **end**
 909 4. **Aggregate gradients across all modules:**
 910 $\mathbf{g}^{(t)} = \sum_{l,m} \text{Aggregate}(\mathbf{G}^{(l,m)})$
 911 5. **Momentum update:**
 912 $\mathbf{m}^{(t)} = \mu \cdot \mathbf{m}^{(t-1)} + \frac{\mathbf{g}^{(t)}}{\|\mathbf{g}^{(t)}\|_1}$
 913 6. **Perturbation update with projection:**
 914 $\delta^{(t)} = \text{Clip}_\epsilon(\delta^{(t-1)} + \eta \cdot \text{sign}(\mathbf{m}^{(t)}))$
 915 **end**
 916 **return** $\mathbf{x}^{\text{adv}} = \mathbf{x} + \delta^{(T)}$
 917

918
 919 Table 6: Model-specific hyperparameter settings used for TESSER. $\omega^{(\cdot)}$ denotes the weakening factor
 920 for each module, λ is the FSGS scaling parameter, σ controls the strength of spectral regularization,
 921 l_{cut} is the attention truncation depth, γ_{base} is the minimum gradient scaling factor, μ is the momentum
 922 decay used in PGD, and η is the step size for perturbation update.

Hyperparameter	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S
$\omega^{(\text{attn})}$	0.45	0.25	0.3	0.4
$\omega^{(\text{qkv})}$	0.5	0.5	1.0	0.8
$\omega^{(\text{mlp})}$	0.7	0.7	0.6	0.5
λ_{attn}	0.4	0.45	0.5	0.45
λ_{qkv}	0.5	0.5	0.5	0.5
λ_{mlp}	0.55	0.55	0.65	0.6
σ (SSR)	0.5	0.7	0.7	0.7
l_{cut}	10	9	4	8
γ_{base}	0.5	0.5	0.5	0.5
μ	1.0	1.0	1.0	1.0
η	1.6/255	1.6/255	1.6/255	1.6/255

936
 937 Table 7: Computational cost (in seconds) for generating a single adversarial example across different
 938 models and methods. FSGS refers to our feature-sensitive gradient scaling, SSR refers to spectral
 939 smoothness regularization, and ATT denotes state of the art.

Model	FSGS	FSGS + SSR	ATT (Ming et al., 2024)
ViT-B/16	0.5	0.52	0.93
PiT-B	0.54	0.6	1.05
CaiT-S/24	1.24	1.27	1.88
Visformer-S	0.35	0.38	1.14

947 weakening model-specific alignment. For instance, in ViT-B/16, increasing σ from 0.5 to 0.8
 948 decreases ViT ASR from 83.21% to 81.21%, but improves CNN ASR from 58.77% to 61.82% and
 949 defended CNN ASR from 46.63% to 52.33%. A similar pattern is observed in PiT-B and CaiT-S/24.

950 Notably, the improvement on defended CNNs is particularly pronounced. For Visformer-S, the ASR
 951 on defended models improves from 46.96% at $\sigma = 0.5$ to 61.5% at $\sigma = 0.8$, a gain of over 14%.
 952 These results confirm that SSR strengthens black-box transferability and robustness by encouraging
 953 low-frequency perturbations that are less dependent on the surrogate model’s internal architecture.

954 In practice, setting σ between 0.6 and 0.8 offers a favorable trade-off, preserving sufficient ViT ASR
 955 while achieving substantial improvements on CNNs and defended models. This ablation supports
 956 the effectiveness of spectral regularization and its role in enhancing generalization under diverse
 957 adversarial settings.

959 C.2 COMPARISON OF ATTACK EFFICIENCY WHEN USING INPUT DIVERSITY TECHNIQUE

961 To further assess the effectiveness and generality of our proposed TESSER framework, we evaluate
 962 its performance when combined with an input diversity enhancement strategy, specifically PatchOut
 963 (PO) (Wei et al., 2022). This technique introduces random masking during inference to improve the
 964 robustness and transferability of adversarial perturbations.

966 Table 9 presents the average Attack Success Rate (ASR) of different attack methods augmented
 967 with PO, tested across ViTs, CNNs, and defended CNNs. The experiments span four representative
 968 surrogate models: ViT-B/16, CaiT-S/24, PiT-B, and Visformer-S.

969 Across all surrogate models and evaluation categories, TESSER+PO consistently achieves the highest
 970 ASR. For example, using PiT-B as the surrogate, TESSER+PO achieves an ASR of **94.83%** on ViTs,
 971 **87.7%** on CNNs, and **61.43%** on defended CNNs, representing improvements of more than **+10%**
 972 over the strongest baseline ATT+PO. Similar trends are observed with the other surrogate models,

972 Table 8: Average attack success rate (ASR) (%) against ViTs, CNNs, and defended CNNs across
 973 varying Gaussian blur strength σ . Increasing σ generally improves transferability to CNNs and
 974 defended models by enforcing low-frequency perturbations, while slightly reducing white-box ASR
 975 on ViTs.

976

Model	σ	ViTs	CNNs	Def-CNNs	Model	σ	ViTs	CNNs	Def-CNNs
ViT-B/16	0.5	83.21	58.77	46.63	CaiT-S/24	0.5	94.82	68.12	45.6
	0.6	83.12	61.85	49.86		0.6	94.57	71.9	51.76
	0.7	81.95	61.62	51.13		0.7	94.2	73.87	54.86
	0.8	81.21	61.82	52.33		0.8	93.82	73.55	57.06
PiT-B	0.5	90.3	80.85	49	Visformer-S	0.5	75.93	76.77	46.96
	0.6	91.63	83.4	54.9		0.6	78.47	80.45	53.9
	0.7	91.36	83.27	57.06		0.7	78.67	81.67	58.46
	0.8	90.02	83.45	58.6		0.8	83.33	81.22	61.5

985

986 Table 9: The average attack success rate (%) against ViTs, CNNs, and defended CNNs by various
 987 transfer-based attacks with input diversity enhancement strategy. The best results are highlighted in
 988 bold. “PO” denotes PatchOut (Wei et al., 2022).

989

Model	Attack	ViTs	CNNs	Def-CNNs	Model	Attack	ViTs	CNNs	Def-CNNs
ViT-B/16	MIM+PO	61.3	31.3	21.7	CaiT-S/24	MIM+PO	70.3	44.0	29.3
	VMI+PO	69.1	42.8	30.9		VMI+PO	76.8	57.8	38.4
	SGM+PO	64.8	29.2	18.9		SGM+PO	85.1	49.2	29.3
	PNA+PO	70.8	42.6	29.9		PNA+PO	81.6	56.6	39.3
	TGR+PO	76.0	46.7	33.3		TGR+PO	88.8	60.5	40.5
	ATT+PO	77.1	51.7	37.1		ATT+PO	91.1	71.9	54.3
	Ours+PO	85.18\uparrow	64.17\uparrow	52.16\uparrow		Ours+PO	91.15\uparrow	72.9\uparrow	56.46\uparrow
	MIM+PO	47.3	32.5	17.5		MIM+PO	54.9	45.7	23.4
PiT-B	VMI+PO	59.5	46.2	35.8		VMI+PO	64.8	56.6	32.6
	SGM+PO	70.0	45.6	21.3		SGM+PO	51.6	44.3	15.0
	PNA+PO	73.1	57.8	32.7		PNA+PO	68.8	61.8	32.3
	TGR+PO	82.3	68.9	41.3		TGR+PO	70.4	64.3	33.5
	ATT+PO	84.2	75.2	48.4		ATT+PO	70.5	79.3	44.5
	Ours+PO	94.83\uparrow	87.7\uparrow	61.43\uparrow		Ours+PO	84.42\uparrow	79.4\uparrow	58.06\uparrow

1001

1002 including CaiT-S/24 and ViT-B/16, where TESSER+PO continues to outperform baselines by wide
 1003 margins.

1004

1005 These results demonstrate two key insights: (1) TESSER is orthogonal to input diversity methods, as
 1006 its performance improves further when used with PO, and (2) our gradient modulation and spectral
 1007 regularization strategies remain effective under randomized input transformations, indicating strong
 1008 generalization.

1009

1010 In particular, on defended CNNs, traditionally difficult targets due to adversarial training, TESSER
 1011 + PO outperforms all baselines by significant margins (e.g., + 7% over ATT + PO with ViT-B/16).
 1012 This highlights that FSGS and SSR lead to perturbations that survive stochastic augmentations while
 1013 preserving transferability and robustness.

1014

C.3 ADVERSARIAL ATTACK EFFICIENCY AND CONFIDENCE DYNAMICS

1015

1016 To better assess the quality of adversarial examples beyond final attack success rate (ASR), we
 1017 evaluate the efficiency and effectiveness of the generated perturbations in terms of iteration-wise
 1018 model response. Specifically, we compare TESSER and ATT based on:

1019

- **Attack efficiency:** How quickly the model’s prediction flips and stabilizes to an adversarial
 1020 label across iterations.
- **Attack effectiveness:** The final confidence of the model in the adversarial label after
 1021 optimization completes.

1023

1024 Table 10 summarizes the average iteration at which the target model stabilizes on the adversarial
 1025 label (i.e., no further label flipping) and the average confidence on the adversarial class after 10 attack
 1026 steps.

1026 Table 10: Comparison of attack efficiency and effectiveness between TESSER and ATT. We report
 1027 the average iteration where the model prediction stabilizes on the adversarial label (lower is better)
 1028 and the final model confidence (%) in the adversarial class (higher is better).

Method	Stabilization Iteration (↓)	Final Confidence (%) (↑)
ATT	6.8	87.93
TESSER (Ours)	5.1	91.37

1034 Table 11: TESSER attack success rate (%) with and without module-wise gradient weakening ω
 1035 against eight ViT models and the average attack success rate (%) of all black-box models. The best
 1036 results are highlighted in **bold**.

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	w/o ω	99.6*	40.9	71	45.7	69.3	64.8	40.2	72.5	63
	w ω	100*	61.7	94	68.3	92.5	85.6	72.2	91.4	83.2↑
PiT-B	w/o ω	58.1	99.9*	69.2	74.7	69	74	66.9	70.3	72.76
	w ω	74.9	100.0*	91.6	93.2	92.1	95	92.4	91.7	91.4↑

1042
 1043
 1044 TESSER reaches a stable adversarial label approximately 1.7 iterations earlier than ATT, confirming its
 1045 improved gradient alignment and optimization direction. Additionally, the final adversarial confidence
 1046 achieved by TESSER is consistently higher, indicating stronger and more decisive misclassification.
 1047 This validates that our semantic gradient modulation not only accelerates convergence but also
 1048 increases attack effectiveness by pushing perturbations toward model-relevant, transferable features.
 1049

1050 D ADDITIONAL ABLATION STUDIES

1051 D.1 QUALITATIVE COMPARISON

1052 To further analyze the effectiveness and interpretability of our proposed method, we present qualitative
 1053 comparisons between TESSER (FSGS+SSR) and the state-of-the-art ATT (Ming et al., 2024) across
 1054 a diverse set of samples from ImageNet. Figure 3 shows clean and adversarial images, Grad-CAM
 1055 heatmaps, and FFT visualizations for perturbations.

1056 **Semantic Alignment.** In nearly all examples, the adversarial images generated by TESSER show
 1057 stronger alignment with semantically meaningful regions (e.g., bird bodies, faces, objects of interest)
 1058 compared to ATT. This is reflected in the Grad-CAM visualizations guided by the adversarial label.
 1059 Despite being misclassified, the Grad-CAM of TESSER adversarial examples remains spatially
 1060 focused on relevant visual features, validating the effectiveness of FSGS in preserving semantically
 1061 informative gradients during attack optimization.

1062 **Spectral Coherence.** The FFT visualizations reveal that TESSER perturbations exhibit smoother
 1063 and more coherent frequency profiles, with lower high-frequency energy content than those generated
 1064 by ATT. This is consistently supported by the computed High-Frequency Energy Ratio (HFER),
 1065 which is reduced by 6–16% across examples when using FSGS+SSR. Lower HFER confirms that
 1066 SSR suppresses architecture-specific, high-frequency noise that often undermines transferability.

1067 These additional ablations reinforce our core claim: FSGS guides perturbations toward transferable,
 1068 semantically meaningful features, while SSR regularizes their spectral profile to avoid overfitting
 1069 to model-specific noise. Together, these properties lead to adversarial examples that are more
 1070 interpretable and more effective in black-box transfer scenarios.

1071 D.2 IMPACT OF MODULE-WISE GRADIENT WEAKENING

1072 We compare ASR with and without ω (i.e., setting all $\omega = 1$ disables gradient weakening). On
 1073 ViT-B/16, using ω improves ASR from 63.0% to 83.2% (Table 11), confirming the effectiveness of
 1074 selective gradient suppression.

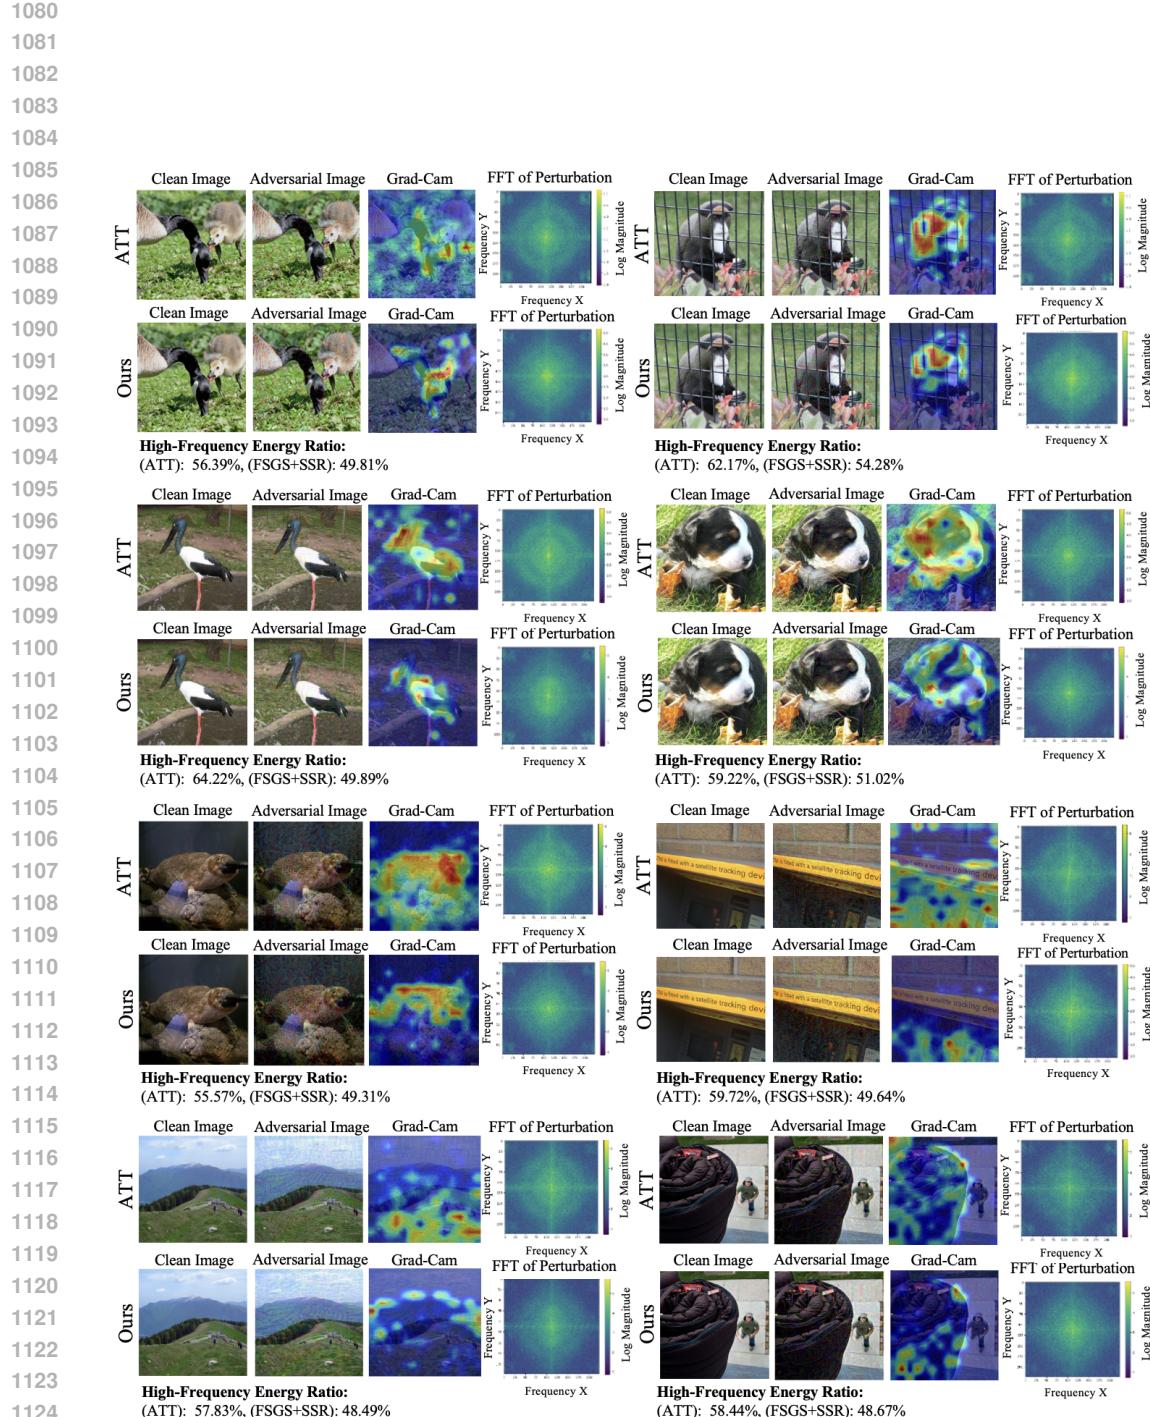


Figure 3: Qualitative comparison between ATT and our TESSER method (FSGS+SSR). Each block shows clean image, adversarial image, Grad-CAM heatmap, and FFT of the perturbation. TESSER yields semantically aligned and spectrally smooth perturbations, with consistently lower high-frequency energy ratios.

Table 12: TESSER attack success rate (%) with and without selective attention truncation l_{cut} against eight ViT models and the average attack success rate (%) of all black-box models. The best results are highlighted in **bold**.

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	w/o l_{cut}	100*	60.4	87.3	67.7	88.3	85.4	65.4	89.9	80.55
	w l_{cut}	100*	61.7	94	68.3	92.5	85.6	72.2	91.4	83.2↑
PiT-B	w/o l_{cut}	73.1	99.7*	82.1	86	84	86.4	83.2	84.1	84.82
	w l_{cut}	74.9	100.0*	91.6	93.2	92.1	95	92.4	91.7	91.4↑

Table 13: TESSER attack success rate (%) with and without rescaling factor λ against eight ViT models and the average attack success rate (%) of all black-box models. The best results are highlighted in **bold**.

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	w/o λ	84.1*	32.1	54.5	42.3	55.1	56	43.1	56.9	53.01
	w λ	100*	61.7	94	68.3	92.5	85.6	72.2	91.4	83.2↑
PiT-B	w/o λ	51.2	92.9*	61.2	67.9	63.6	67.9	66.5	62.6	66.72
	w λ	74.9	100.0*	91.6	93.2	92.1	95	92.4	91.7	91.4↑

D.3 IMPACT OF SELECTIVE ATTENTION TRUNCATION

On PiT-B, disabling attention truncation (i.e., no l_{cut}) leads to an average 7% drop in ASR (Table 12), validating the importance of focusing on early-layer token gradients.

D.4 IMPACT OF RESCALING FACTOR

Setting $\lambda = 0$ disables adaptive FSGS scaling (only γ_{base} is used as a fixed multiplier). On ViT-B/16, enabling λ improves ASR by an average of 30% (Table 13), highlighting the value of adaptive gradient modulation in improving attack effectiveness.

In addition, we empirically validate the effectiveness of our scaling strategy by comparing it to a random scaling baseline. As shown in Table 14, our method significantly outperforms random scaling across all target models, achieving consistently higher ASR and demonstrating stronger transferability.

E EVALUATING THE TRANSFERABILITY OF DIFFERENT ATTACK METHODS FOR TARGETED ATTACKS

While our main experiments focus on untargeted attacks, both FSGS and SSR are model-agnostic and loss-independent components applied during backpropagation. Therefore, they are fully compatible with targeted attack formulations—only the loss needs to be adapted. To validate this, we conducted targeted attack experiments using the target label set as (true label + 1). As shown in Table 15, our method (TESSER) achieves a significantly higher targeted ASR of 43.08%, outperforming PGD (16.06%), MIM (22.01%), and ATT (33.33%), demonstrating the effectiveness and transferability of our approach in targeted settings as well. The table will be included in the revised version.

Table 14: TESSER Attack Success Rate (%) with our Scaling strategy vs. Random Scaling. **Bold** = better of the two scalings for the same surrogate–target pair. * denotes white-box (surrogate equals target).

Surrogate	Scaling	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg
ViT-B/16	Random	86.4*	30.6	52.3	37.8	53.0	55.4	37.8	55.7	51.12
	ours	100*	61.7	94.0	68.3	92.5	85.6	72.2	91.4	83.2↑
PiT-B	Random	28.4	100*	34.6	44.8	33.9	44.1	38.3	38.2	45.28
	ours	74.9	100.0*	91.6	93.2	92.1	95.0	92.4	91.7	91.4↑

1188 Table 15: The attack success rate (%) of various transfer-based targeted attacks against eight ViT
 1189 models and the average attack success rate (%) of all black-box models. The best results are
 1190 highlighted in **bold**.

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	PGD	96.1*	2.1	6.8	2.7	5.6	6	1.7	7.5	16.06
	MIM	99.4*	6.6	16.1	6.3	13.9	11.8	4.4	17.6	22.01
	ATT	99.5*	8.7	35.6	9.6	33.6	30.9	6.7	42.1	33.33
	Ours	99.6*	17.9	47.3	21	48.6	42.1	16.1	52.1	43.08↑
PiT-B	PGD	0.8	96.6*	1.1	2.3	0.9	2.1	1.3	0.7	13.22
	MIM	5.3	99.9*	5.1	8	4.9	6.5	4	5.5	17.4
	ATT	12.1	100*	14.6	20.2	13.2	18.8	12.2	16.8	25.98
	Ours	20.4	100.0*	26	33	28.4	30.2	26.3	27.2	47.86↑

1199
 1200 Table 16: The attack success rate (%) of Autoattack (AA) vs. TESSER against eight ViT models and
 1201 the average attack success rate (%) of all black-box models. The best results are highlighted in **bold**.
 1202

Model	Attack	ViT-B/16	PiT-B	CaiT-S/24	Visformer-S	DeiT-B	TNT-S	LeViT-256	ConViT-B	Avg _{bb}
ViT-B/16	AA	99.9*	10.5	42.9	13.2	33.6	38.4	17	40.4	36.98
	Ours	100*	61.7	94	68.3	92.5	85.6	72.2	91.4	83.2↑
PiT-B	AA	10	98.5*	13.5	24	11.3	21.2	22.4	13.9	26.85
	Ours	74.9	100.0*	91.6	93.2	92.1	95	92.4	91.7	91.4↑

1209 F EVALUATING TESSER PERFORMANCE VS. AUTOATTACK

1211 While AutoAttack (AA) is a strong white-box evaluation benchmark, it is not optimized for transfer-
 1212 based black-box settings. To enable a fair comparison, we evaluate both TESSER and AutoAttack
 1213 under the same transfer setup with a fixed perturbation budget of epsilon = 16/255. As reported in
 1214 Table 16, TESSER achieves over 50% higher average ASR compared to AutoAttack across multiple
 1215 target models. For example, when attacking PiT-B from a ViT-B/16 surrogate, TESSER achieves
 1216 an ASR of 61.7%, compared to only 10.5% for AutoAttack. This gap is expected, as TESSER is
 1217 explicitly designed to optimize black-box transferability, whereas AutoAttack is tailored for white-box
 1218 robustness evaluation.

1220 G EVALUATION TESSER TRANSFERABILITY TO VISUAL STATE SPACE 1221 MODELS

1223 To further increase the architectural dissimilarity, we evaluate TESSER transferability to Vision
 1224 Mamba (Zhu et al., 2024), a state-space-based architecture with bidirectional SSMs and position,
 1225 aware embeddings, representing a class of models distinct from transformers. As shown in Table 17,
 1226 TESSER consistently achieves the highest ASR 87.1% and 76.7% on both Vim-Tiny and Vim-Small,
 1227 respectively compared to 80.9% and 69% for sota ATT attack, demonstrating robust transfer even
 1228 under significant architectural and representational divergence.

1230 Table 17: Comparative experiments of different attack methods on VIM. “clean” indicates that clean
 1231 images are classified and all results indicate the percentage of classification errors (*i.e.*, ASR).
 1232

Model	Attack	VIM-tiny	VIM-small
ViT-B/16	clean	3.1	0.9
	MIM	45.6	42
	ATT	80.9	69
	TESSER	87.1↑	76.7↑
PiT-B	MIM	32.3	34.9
	ATT	53.4	55.1
	TESSER	77.4↑	80.7↑

1242 H SSR EFFECTIVENESS IN THE FREQUENCY DOMAIN

1244 H.1 CONTROLLED BAND-PASS STRESS TEST (SPECTRAL ANALYSIS OF PERTURBATIONS)

1246 To quantitatively evaluate the spectral behavior introduced by the proposed **Spectral Smoothness**
 1247 **Regularization (SSR)** and to compare it with existing baselines (**ATT** and **TGR**), we design a
 1248 controlled frequency-domain stress test that isolates the contribution of different spectral bands in
 1249 adversarial perturbations.

1251 **Experimental Setup.** For each attack method, we compute perturbations $\delta = x_{\text{adv}} - x_{\text{clean}}$ and
 1252 apply frequency-domain filtering prior to reconstruction. Specifically, we generate ideal circular
 1253 *low-pass (LP)* and *high-pass (HP)* masks in the Fourier domain with normalized cutoff radii $r \in$
 1254 $\{0.05, 0.1, 0.2, 0.3, 0.5\} \times r_{\text{max}}$. Each perturbation is filtered as

$$1255 \quad \tilde{\delta} = \mathcal{F}^{-1}(\mathcal{M} \odot \mathcal{F}(\delta)),$$

1257 where \mathcal{M} is the LP or HP mask and \mathcal{F} denotes the 2D Fourier Transform. The filtered perturbation
 1258 is then projected back into the valid L_{∞} -bounded region, producing new adversarial examples
 1259 $x' = \text{clip}(x_{\text{clean}} + \tilde{\delta})$. We then re-evaluate the *Attack Success Rate (ASR)* of each filtered adversarial
 1260 image on the target model. This controlled procedure directly measures how strongly each attack
 1261 relies on specific frequency bands.

1262 The measured ASR (%) across cutoff frequencies for both HP and LP masks is summarized in
 1263 Table 18.

1265 Method	1266 HP (High-Pass Filter)					1267 LP (Low-Pass Filter)				
	0.05	0.10	0.20	0.30	0.50	0.05	0.10	0.20	0.30	0.50
1268 ATT	100.0	98.8	64.9	15.0	0.5	0.0	2.0	19.4	75.5	99.7
1269 TESSER (SSR)	99.2	97.5	88.7	59.7	3.5	4.2	23.2	77.5	95.4	99.6
1270 TGR	99.8	99.6	96.7	83.2	15.9	3.9	17.8	60.6	88.4	99.5

1271 Table 18: ASR (%) under controlled frequency-domain filtering. HP columns show the effect of
 1272 retaining only high-frequency components, while LP columns show the effect of retaining only
 1273 low-frequency components. Cutoff values correspond to normalized frequency radii (r/r_{max}).

1274 A high ASR under HP filtering indicates that the attack relies on fine-grained, noise-like high-
 1275 frequency content, while a high ASR under LP filtering suggests that the attack’s discriminative
 1276 power comes from smoother, structured low-frequency variations. Several key observations can be
 1277 made:

- 1280 • **ATT:** exhibits a strong high-frequency bias. Its ASR collapses rapidly under low-pass
 1281 filtering (from 99% to 0%), indicating that its perturbations primarily occupy the high-
 1282 frequency spectrum.
- 1283 • **TESSER (with SSR):** maintains significantly higher ASR under LP filtering (77.5% at
 1284 $r = 0.2$ vs. 19.4% for ATT), confirming that SSR distributes energy across lower and mid-
 1285 frequency bands and enforces *spectral smoothness*. This evidences that TESSER generates
 1286 more structured, semantically aligned perturbations rather than high-frequency noise.
- 1287 • **TGR:** shows intermediate behavior, partially resilient to LP filtering but still decaying
 1288 faster than TESSER, highlighting that TESSER’s SSR achieves stronger low-frequency
 1289 regularization.

1291 This experiment demonstrates that **Spectral Smoothness Regularization** effectively redistributes per-
 1292 turbation energy from high-frequency, noise-like components to smoother low-frequency structures.
 1293 Such spectral redistribution validates the intended effect of SSR and correlates with the observed
 1294 improvements in transferability and robustness across models and quantization levels. Overall, SSR
 1295 produces *spectrally smoother, semantically coherent perturbations*, whereas ATT remains heavily
 1296 reliant on fragile high-frequency artifacts.

1296
1297

H.2 VERIFYING THE EFFECTIVENESS OF SSR IN THE FREQUENCY DOMAIN.

1298
1299
1300
1301
1302

We thank the reviewer for encouraging a more controlled spectral analysis. To quantify the frequency redistribution introduced by **Spectral Smoothness Regularization (SSR)**, we computed the relative energy of perturbations from **TESSER**, **ATT**, and **TGR** after applying ideal low-pass (LPF) and high-pass (HPF) filters with a cutoff of $0.25 \times \text{Nyquist}$. The measured spectral energy distribution is reported below.

1303
1304
1305
1306
1307

Method	LPF Energy (%) ↓	HPF Energy (%) ↓
ATT	47.6	52.4
TGR	40.0	60.0
TESSER (w/ SSR)	56.3	43.7

1308
1309

Table 19: Spectral energy distribution of perturbations under LPF/HPF decomposition (cutoff = $0.25 \times \text{Nyquist}$).

1310
1311
1312
1313
1314
1315

Compared to ATT and TGR, **TESSER** shows a **+8%–+16% increase in low-frequency energy and a -9%–-16% reduction in high-frequency energy**, confirming that SSR effectively shifts perturbation power toward smoother and more transferable frequency bands. When SSR is disabled, TESSER’s cross-architecture ASR on CNNs drops from **74.4%** to **70.3%**, demonstrating that this spectral adjustment materially enhances transferability.

1316
1317
1318
1319

I ADDITIONAL EMPIRICAL VALIDATION OF SEMANTIC AND STRUCTURAL ALIGNMENT

1320
1321

I.1 EMPIRICAL VALIDATION OF THE SEMANTIC MEANING OF TOKEN NORMS

1322
1323
1324
1325
1326
1327
1328

To validate the core assumption underlying FSGS, we performed a layer-wise analysis measuring the Spearman correlation between token feature norms and token gradient norms across all ViT-B/16 blocks. As shown in Table 20, early layers exhibit weak or even negative correlation, confirming that high-norm shallow tokens do not encode semantic content (consistent with the reviewer’s dual-task critique). In contrast, correlation increases steadily through mid-level layers and peaks in deep layers (up to 0.49), indicating that high-norm tokens reliably capture semantic importance precisely where the network forms object-level representations.

1329
1330
1331

This provides direct empirical evidence supporting our design choice: *FSGS does not operate on shallow layers, but applies semantic-aware scaling only in mid-to-deep blocks, where token magnitude meaningfully reflects semantic relevance.*

1332
1333
1334

Table 20: Spearman correlation between token feature norms and token gradient norms across ViT-B/16 layers.

1335
1336
1337

Layer	0	1	2	3	4	5	6	7	8	9	10	11
Spearman ρ	-0.03	0.03	0.05	0.10	0.25	0.21	0.29	0.37	0.36	0.49	0.34	0.00

1338
1339
1340

I.2 ATTENTION ROLLOUT ANALYSIS (ViT-SPECIFIC SEMANTIC ALIGNMENT)

1341
1342
1343
1344
1345

To directly address the reviewer’s concern regarding the suitability of Grad-CAM for ViTs, we evaluate semantic alignment using *Attention Rollout* (Abnar & Zuidema, 2020), a ViT-native interpretability method. For each clean image and its adversarial counterpart, we compute the CLS→patch rollout map and measure (i) Spearman correlation and (ii) intersection-over-union (IoU) between the two saliency maps.

1346
1347
1348
1349

Rollout maps are structurally stable due to their cumulative nature across layers and heads, which explains the overall high correlations across all methods. Despite this stability, meaningful relative differences arise across attacks. As shown in Table 21, TESSER achieves the strongest semantic alignment (Spearman = 0.9894, IoU = 0.8258), outperforming ATT (0.9581 / 0.8045) and substantially exceeding TGR (0.9322 / 0.6596). These findings demonstrate that *TESSER preserves global*

1350 *semantic attention flow more faithfully than prior methods*, complementing our Grad-CAM results
 1351 and providing ViT-specific evidence for the semantic coherence of TESSER’s perturbations.
 1352

1353 Table 21: Attention Rollout-based semantic alignment between clean and adversarial images. Higher
 1354 is better.

1355

	Method	Rollout Spearman	Rollout IoU
1356	ATT	0.9581	0.8045
1357	TGR	0.9322	0.6596
1358	TESSER	0.9894	0.8258

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403