

# OTI: A Model-Free and Visually Interpretable Measure of Image Attackability

Jiaming Liang<sup>1</sup>, Haowei Liu<sup>2</sup>, Chi-Man Pun<sup>1\*</sup>

<sup>1</sup>Faculty of Science and Technology, University of Macau, Macau, China

<sup>2</sup>Chongqing Key Laboratory of Image Cognition, Chongqing University of Posts and Telecommunications, China  
chinaliangjm@gmail.com, S230201069@stu.cqupt.edu.cn, cmpun@um.edu.mo

## Abstract

Despite the tremendous success of neural networks, benign images can be corrupted by adversarial perturbations to deceive these models. Intriguingly, images differ in their attackability. Specifically, given an attack configuration, some images are easily corrupted, whereas others are more resistant. Evaluating image attackability has important applications in active learning, adversarial training, and attack enhancement. This prompts a growing interest in developing attackability measures. However, existing methods are scarce and suffer from two major limitations: (1) They rely on a model proxy to provide prior knowledge (e.g., gradients or minimal perturbation) to extract model-dependent image features. Unfortunately, in practice, many task-specific models are not readily accessible. (2) Extracted features characterizing image attackability lack visual interpretability, obscuring their direct relationship with the images. To address these, we propose a novel **Object Texture Intensity (OTI)**, a model-free and visually interpretable measure of image attackability, which measures image attackability as the texture intensity of the image’s semantic object. Theoretically, we describe the principles of OTI from the perspectives of decision boundaries as well as the mid- and high-frequency characteristics of adversarial perturbations. Comprehensive experiments demonstrate that OTI is effective and computationally efficient. In addition, our OTI provides the adversarial machine learning community with a visual understanding of attackability.

**Code** — <https://github.com/chinaliangjiaming/OTI>

## 1 Introduction

Deep neural networks (DNNs) are playing a growing role in critical sectors, yet carefully crafted adversarial perturbations (Szegedy et al. 2013) by adversaries can easily fool them, which incur significant threats. Hence, robustness analysis has long been a central theme in adversarial machine learning.

Robustness analysis (Liu et al. 2025) from a *model-centric* perspective has evolved over many years, spanning formal robustness certification (Zhang et al. 2024a), robustness evaluation (Yan et al. 2024), robust architecture design (Xu et al. 2025), and adversarial defense (Liu et al.

\*Corresponding author.

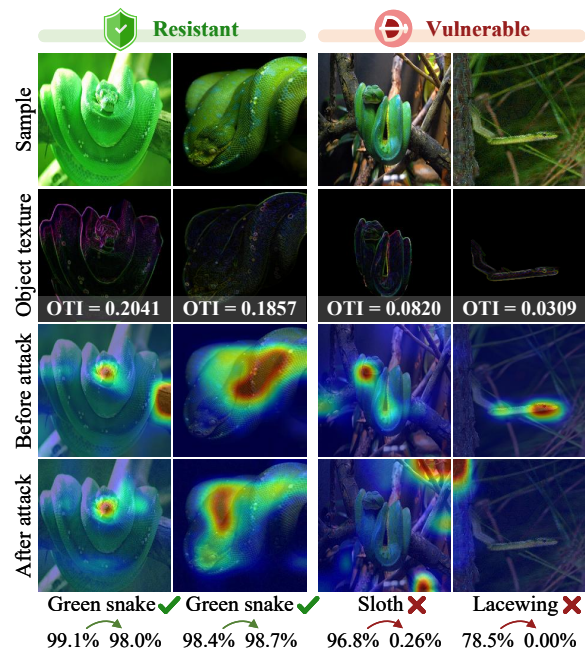


Figure 1: Measuring image attackability by object texture intensity (values in the second row). The third and fourth rows are Grad-CAM visualizations of benign images and adversarial examples under BSR attack (surrogate: R50, target: D161). The last two rows show the predicted labels of adversarial examples and the predicted probabilities of the true class for both benign images and adversarial examples.

2025). In contrast, recent work (Raina and Gales 2023) proposes a novel robustness concept called image attackability from a *sample-centric* viewpoint. Unlike the traditional focus on evaluating neural network robustness, this idea aims to assess the robustness<sup>1</sup> of *benign images* against adversarial perturbations. In defense applications, image attackability evaluation helps identify the most informative images for active learning, select the most vulnerable images for efficient generation of adversarial examples for adver-

<sup>1</sup>Higher robustness implies lower attackability and lower vulnerability. These three terms are interchangeable in this paper.

sarial training (Madry et al. 2017), and construct effective subsets for debiasing (Zhu et al. 2025). In attack applications, attackability evaluation enables more effective attacks by selecting the most vulnerable images. Furthermore, understanding image attackability can deepen our insight into the effectiveness of adversarial perturbations and thus advance the adversarial machine learning community.

Unfortunately, existing image attackability measures are scarce and face two major limitations: (1) *Model-dependent*. IAARS (Raina and Gales 2023) relies on a model proxy to train an image attackability discriminator and leverages its generalization capability to assess the attackability of unseen images on unseen models. ZGP (Karunanayake et al. 2025) measures image attackability as the proportion of elements in the gradient of the proxy’s loss with respect to the input whose absolute values are below a specified threshold. Both methods require a trained model proxy, which is often infeasible in many practical scenarios, especially where public models are not readily available (e.g., in the field of medical imaging) and the cost of training models is prohibitively high. Ideally, image attackability measures should work based only on the intrinsic attributes of the image itself and the task. (2) *Visually uninterpretable*. Features characterizing image attackability extracted by existing methods lack visual interpretability for humans. IAARS encodes these features into the neural-network discriminator, and ZGP utilizes gradients. It is difficult to establish a visual connection between the extracted features and the image itself. Therefore, we aim to develop a *model-free* and visually interpretable image attackability measure.

Empirical results on visual tasks suggest that the smaller the semantic object (Fan et al. 2021) and the weaker the texture intensity, the more attackable the image. We explain these phenomena from the perspective of decision boundary (Yan et al. 2024) and the mid- and high-frequency nature (Wang et al. 2020) of adversarial perturbations. From the perspective of decision boundary, smaller semantic objects and weaker textures make images difficult to classify. As a result, their degradation vectors are greater and these images are closer from the decision boundary, leading to greater attackability. From the perspective of frequency domain, since adversarial perturbations are generally believed to rely heavily on mid- and high-frequency components, stronger semantic textures safeguard images against the adverse impact of the mid- and high-frequency components of adversarial perturbations and thus contribute to improved robustness.

Accordingly, we propose a novel **Object Texture Intensity (OTI)** measure, which quantifies image attackability by the sum of absolute texture values within the semantic object region, as shown in Figure 1. A lower OTI suggests a smaller semantic object and weaker texture intensity, indicating greater attackability. Extensive experiments demonstrate that OTI can effectively distinguish the image attackability across various advanced attacks, diverse domains, different tasks, varying dataset sizes, in both defense and non-defense scenarios, and across different configurations. The main contributions are summarized as follows:

- To the best of our knowledge, this is the first work to reveal the relationship between the texture intensity of semantic objects and image attackability.
- We describe the relationship from the perspectives of decision boundary theory as well as the mid- and high-frequency nature of perturbations.
- Based on this relationship, we propose a novel image attackability measure, Object Texture Intensity (OTI). This is the first model-free and visually interpretable measure.
- Our OTI provides a visual reference for human understanding of image attackability. Extensive experiments demonstrate that although OTI is model-free and simple, it is highly effective across different tasks, various attacks, and diverse configurations.

## 2 Related Work

In this section, we provide a comprehensive review of adversarial attacks and robustness analysis.

### 2.1 Adversarial Attacks

Adversarial attacks (Szegedy et al. 2013) refer to the process of applying imperceptible perturbations to benign samples to induce decision deviations in the model. In the current landscape, according to the source of perturbations, adversarial attacks are generally classified into three major categories: gradient-based, query-based, and generation-based attacks.

*Gradient-based attacks* exploit the input gradients of the loss to carry out gradient ascent, thus disrupting the model’s decision. In the white-box settings, access to the full gradient information of the model enables nearly perfect attacks (Kurakin, Goodfellow, and Bengio 2018; Dong et al. 2018). In the black-box settings, adversaries typically exploit the adversarial transferability, whereby adversarial examples effective on one model are likely to succeed on other models. To improve transferability, transformation-based strategies (Lin et al. 2024; Liang and Pun 2025), model-related strategies (Ma et al. 2024), ensemble-based strategies (Chen et al. 2024; Tang et al. 2024), gradient-optimization strategies (Wang et al. 2024a), and objective-optimization strategies (Zhou et al. 2018; Li et al. 2023) are employed.

*Query-based attacks* rely solely on the outputs of the target model and iteratively update the perturbation until the attack succeeds or the query budget is exhausted. According to the type of output returned by the target model, query-based attacks can be categorized into decision-based and score-based attacks. Decision-based attacks (Wang et al. 2025b) only have access to the predicted class label, whereas score-based attacks (Reza et al. 2025) can access the predicted probability distribution.

*Generation-based attacks* employ neural-network generators to craft adversarial examples. These generators are typically trained to either increase the output loss of a surrogate model (Li et al. 2024) or to shift its intermediate feature representations (Li, Ma, and Jiang 2025).

This paper will evaluate the proposed measure under the aforementioned various advanced adversarial attacks.




Figure 2: Illustrations of how semantic object area ratio and texture intensity affect the robustness of benign samples. DeCoWA is used as the attack with R50 as the surrogate model. The target model is D161 in (a) and is SwinT in (b).

## 2.2 Model-Centric Robustness Analysis

Conventional studies on adversarial robustness primarily focus on a model-centric perspective.

*Formal robustness certification* (Zhang et al. 2024a) uses formal methods to mathematically prove whether a neural network satisfies certain robustness properties.

*Robustness evaluation* relies on empirical methods to estimate a model’s robustness. Representative approaches include evaluating the model’s accuracy under selected adversarial attacks (Xie et al. 2025) and estimating the geometric distance between a sample and the model’s decision boundary (Yan et al. 2024).

*Robust architecture design* (Xu et al. 2025) explores how architectural choices affect model robustness, seeking structural variations that improve robustness even under standard training.

*Adversarial defenses* focus on arming models with carefully designed mechanisms to withstand adversarial threats. Representative strategies encompass adversarial training (Liu et al. 2025), input purification (Song et al. 2024), randomized smoothing (Chaouai and Tamaazousti 2024), and adversarial detection (Li, Angelov, and Suri 2024).

## 2.3 Sample-Centric Robustness Analysis

Unlike model-centric robustness analysis, recent studies reconsider from a sample-centric perspective. They aim to find measures for estimating image attackability. It is promising in active learning, adversarial training, and sample selection-based efficient attacks. It mainly includes two types: (1) evaluating the attackability of benign samples and (2) predicting the transferability of adversarial examples. ET (Levy et al. 2024) utilizes multiple surrogate models to approximate the expected transferability of adversarial examples. IAARS (Raina and Gales 2023) measures the attackability of benign samples as the smallest perturbation required to change the prediction of the model proxy. They use a given dataset and model proxy to train a neural-network discriminator. Then, they utilize the discriminator to evaluate the attackability of unseen images on unseen models. ZGP (Karunanayake et al. 2025) measures the attackability as the element ratio in the image gradients of the proxy’s loss, whose absolute value is below a certain threshold.

However, ET, IAARS and ZGP are model-dependent. Moreover, the attackability features they extract cannot be

directly linked to the image itself in a visually interpretable manner. To address these, this paper aims to develop a model-free and visually interpretable attackability measure.

## 3 Methodology

### 3.1 Problem Definition

Given an image set  $D = \{(x, y) \in \mathcal{X} \times \mathcal{Y}\}$ , where  $\mathcal{X}$  and  $\mathcal{Y}$  represent the image space and label space, respectively, for a target model  $M$  and an attack  $A$ , we aim to find a model-free and visually interpretable measure  $\phi^*(x)$  that satisfies

$$\phi^* = \arg \max_{\phi} \frac{|\{(x, y) \in D_{\alpha}(\phi) | M(A(x)) \neq y\}|}{\alpha |D|}, \quad (1)$$

in contrast to the conventional model-dependent one  $\phi^*(x, M')$ . Here,  $|\cdot|$  indicates the cardinality of a set.  $D_{\alpha}(\phi)$  is the top- $\alpha$  subset of attackable images from the set  $D$ , sorted according to the measure  $\phi$ .  $0 < \alpha < 1$  is the sampling rate (SR). For convenience, we define the objective function in Equation 1 as  $r(D_{\alpha}(\phi), M, A)$ .

However, optimization problem 1 is intractable, and we thus relax it to a more solvable problem as below

$$\begin{aligned} \hat{\phi} &= \phi, \\ \text{s.t. } r(D_{\alpha}(\phi), M, A) - r(D_{\alpha}(\phi_{\text{random}}), M, A) &> \eta, \end{aligned} \quad (2)$$

where  $\phi_{\text{random}}$  is the random evaluation strategy, and  $\eta$  is the acceptable threshold for the attack success rate (ASR) gain.

### 3.2 Motivation

Comprehensive empirical studies across various attacks, tasks, datasets, and other configurations reveal that two types of images tend to be more attackable:

- 1) Images with smaller semantic object area;
- 2) Images with weaker textures and lower contrast.

Illustrative examples based on the DeCoWA attack (Lin et al. 2024) are presented in Figure 2. For the harp class, images that resist adversarial perturbations typically have larger harp objects than those that are successfully attacked. In the case of sturgeon class, low-contrast and weak-texture images taken underwater are generally more attackable compared to clearly captured images taken above water. Motivated by these two sample-intrinsic and human-perceivable trends, we propose the following measures.

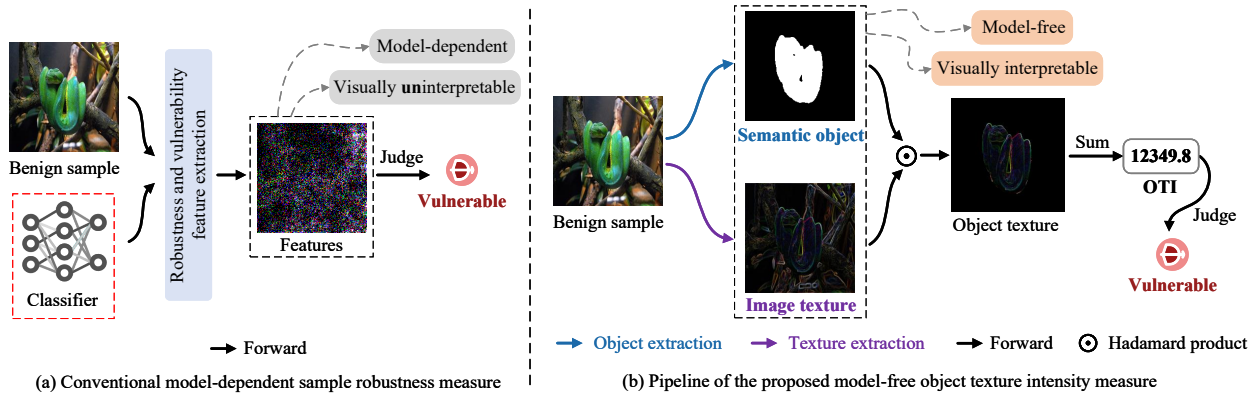


Figure 3: Comparison of overall pipelines between conventional methods and our proposed method.

### 3.3 Object Texture Intensity

We quantify the first trend using the **Object Area Ratio (OAR)**, which indicates the proportion of the image area occupied by task-relevant objects. OAR is formally defined as

$$\text{OAR}(x^{(C \times H \times W)}) = \frac{1}{C \times H \times W} \|\text{object}(x)\|_1. \quad (3)$$

Here, the function  $\text{object}(\cdot)$  is a binarization function that converts the input  $x$  into a single-channel object segmentation map. It is worth noting that  $\text{object}(\cdot)$  can be realized by semantic segmentation or salient object detection (Liu et al. 2021) networks, gradient activation maps (Selvaraju et al. 2017), other coarse annotations, or precise manual labels, as long as it produces an approximately correct binarized object segmentation map.

For the second trend, we quantify it using **Image Texture Intensity (ITI)**, which indicates the sum of absolute texture values in an image. ITI is formulated as

$$\text{ITI}(x^{(C \times H \times W)}) = \frac{1}{C \times H \times W} \|f * x\|_1, \quad (4)$$

where  $f$  denotes a texture extraction operator and  $*$  denotes the convolution operation. While the choice of  $f$  is flexible, we find that the Sobel operator (Gonzales and Wintz 1987) suffices to effectively differentiate image attackability. To expose the intrinsic relationship between texture intensity and attackability, as well as improve computational efficiency, we adopt the Sobel operator as  $f$  in this work.

OAR and ITI measure distinct aspects of an image. By integrating both, we can better leverage their complementary properties to construct a stronger measure, **Object Texture Intensity (OTI)**, which quantifies the texture intensity of the semantic object in an image. OTI is formulated as

$$\text{OTI}(x) = \frac{1}{C \times H \times W} \|\text{object}(x) \odot (f * x)\|_1, \quad (5)$$

where  $\odot$  denotes the Hadamard product. Figure 3(b) illustrates a complete pipeline of OTI.




Figure 4: Illustration explaining the effectiveness of OTI based on model decision boundaries.

### 3.4 Explanation of the Measure

In this section, we explain OTI’s effectiveness from the perspectives of decision boundaries as well as the mid- and high-frequency characteristics of adversarial perturbations.

*Explanation based on decision boundaries.* In the model-centric robustness analysis, it is widely believed that a primary source of model vulnerability is the small input margin (Elsayed et al. 2018; Yan et al. 2024), where the decision boundary is too close to the input sample, making the model susceptible to even minor perturbations. Shifting to a sample-centric perspective, as illustrated in Figure 4, attackable samples can be pushed across the decision boundary with only a few steps of perturbation, whereas robust samples require significantly more steps. Thus, the image attackability can essentially be characterized by the image’s distance to the decision boundary. Going a step further, a robust image can be converted into a vulnerable counterpart by a degradation vector. Applying certain transformations to a robust image can bring it closer to the decision boundary, thereby increasing its attackability. For example, downscaling the semantic object reduces the amount of discriminative information and lowers the signal-to-noise ratio (Pawlowski et al. 2019), bringing the sample closer to

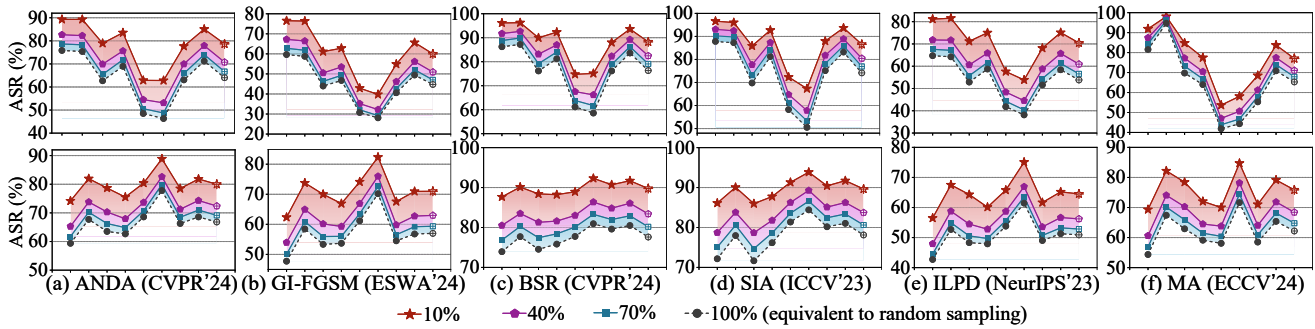


Figure 5: ASRs under varying sampling rates  $\alpha$  in cross-model untargeted attacks. The surrogate models are ResNet-50 (top row) and ViT-B/16 (bottom row). The x-axis represents the target models, with the last one indicating the average.

the decision boundary. Likewise, image blurring leads to the loss of high-frequency components (Zhang et al. 2024b), which also moves the sample closer to the decision boundary. Therefore, OTI essentially estimates a relative degradation vector based on the object size and image texture, which in turn approximates the distance to the decision boundary. A lower OTI indicates a larger degradation vector, meaning the sample lies closer to the decision boundary and is thus more attackable.

*Explanation based on mid- and high-frequency characteristics of adversarial perturbations.* The mid- and high-frequency components are generally considered to contain higher energy in adversarial perturbations (Wang et al. 2020), especially the recent study (Chen et al. 2025) that highlights a positive correlation between high-frequency components and attack performance. OTI quantifies the intensity of semantically relevant mid- and high-frequency information. A lower OTI suggests that the image contains fewer mid- and high-frequency components, making it less capable of resisting the high-energy mid- and high-frequency components of adversarial perturbations, and thus more attackable.

In summary, from the perspective of decision boundaries, OTI characterizes the magnitude of the degradation vector. From the frequency-domain perspective, OTI reflects the amount of semantically relevant mid- and high-frequency components. An image’s attackability to adversarial perturbations can thus be inferred from its semantic object size and the intensity of its texture. In this way, we establish a visual connection between an image and its attackability.

## 4 Experiments and Results

### 4.1 Setup

**Datasets.** To evaluate our OTI on sets of different scales, we adopt the large-scale ImageNet Validation dataset (Deng et al. 2009) and a small-scale subset of 1,000 images for natural image classification. Additionally, to assess OTI across different tasks and domains, we use Kvasir-SEG (Jha et al. 2019) for gastrointestinal polyp segmentation.

**Evaluation benchmarks.** Diverse advanced attacks from various strategies are adopted as benchmarks, including gradient-optimization methods ANDA (Fang et al.

2024) and GI-FGSM (Wang et al. 2024a), transformation-based methods SIA (Wang, Zhang, and Zhang 2023) and BSR (Wang et al. 2024b), objective-optimization method ILPD (Li et al. 2023), model-related method MA (Ma et al. 2024), ensemble-based methods Ens (Liu et al. 2016), AdaEA (Chen et al. 2023), SMR (Tang et al. 2024), and CWA (Chen et al. 2024), query-based methods ADBA (Wang et al. 2025a), TtBA (Wang et al. 2025b), and methods specifically designed for targeted attacks SU (Wei et al. 2023), Logit Margin (Weng et al. 2023), CFM (Byun et al. 2023) and FFT (Zeng, Chen, and Peng 2024).

**Models.** Single surrogate models used are ViT-B/16 and R50. Ensemble surrogate models used include  $E_1$ :{R18, R50, R101},  $E_2$ :{MV2, IncV3, BeiT-B/16}, and  $E_3$ :{ViT-S/16, ViT-S/32, ViT-B/32}. Undefended target models include R152, D161, IncRes-V2, ConvNeXt-B, Swin-B, BeiT-B/16, XCiT-S, and Poolformer. The defended adversarially trained models include ConvNeXt-B + ConvStem (Singh, Croce, and Hein 2023), ViT-B + ConvStem (Singh, Croce, and Hein 2023), ConvNeXt-B (Liu et al. 2025), SwinT-B (Liu et al. 2025), RaWideR101-2 (Peng et al. 2023), and ConvNeXtV2-L + Swin-L (Bai et al. 2024). The target models used for segmentation include: U-Net, UNet-CCT, UNet-URPC, MWCNN, ATTU-Net, and ResU-Net. The experimental results are organized by default in the same order as presented here.

**Metrics.** For classification tasks, we use attack success rate (ASR) as the metric, defined as the proportion of successfully attacked samples over the total number of samples. For segmentation tasks, we adopt the precision, F1-score, and IoU as metrics. Given a sampling rate  $\alpha$  for vulnerable images, a higher ASR or lower precision, F1-score, and IoU suggest that OTI achieves more accurate identification and thus demonstrates a stronger capability in evaluating image attackability. Experiments are conducted on NVIDIA A100 Tensor Core GPU, and the results reported are based on a single run.

**Details of Implementation.** For the ImageNet validation dataset, where the semantic and salient objects largely overlap, we use the saliency detector VST (Liu et al. 2021) as the object segmenter. For Kvasir-SEG, we simply use the manually annotated groundtruth as the segmentation map.




Figure 6: ASRs under varying sampling rates  $\alpha$  in ensemble-based attacks. The x-axis represents the target models, with the last one indicating the average.

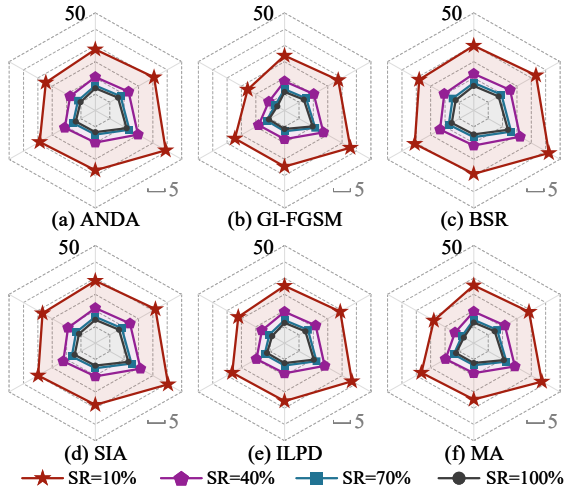


Figure 7: ASRs under varying sampling rates  $\alpha$  in attacking adversarially trained models. The first surrogate model is at the top of the hexagon, with others arrange clockwise.

## 4.2 Evaluation of Robustness Discrimination

**Under single-surrogate untargeted attacks.** In this experiment, we evaluate OTI under single-surrogate untargeted attacks. Specifically, we rank the 50K benign images from the ImageNet Validation dataset by OTI, and select the top- $\alpha$  most attackable images. These images are then used to craft adversarial examples on the surrogate model, which are transferred to attack black-box target models. We report ASRs across varying  $\alpha$  values. For comparison, we also report ASRs of the random sampling strategy (equivalent to OTI-based sampling of  $\alpha = 100\%$ ). A larger ASR gap between OTI-based and random sampling at different sampling rates indicates better performance of OTI. We select 6 advanced benchmarks. The surrogate model and target models are as described in the Setup. The perturbation budget is  $L_\infty = 10/255$ . The results are illustrated in Figure 5. Experimental results show that images selected at lower sampling rates achieve higher ASRs. At  $\alpha = 10\%$ , the average ASR improves by 13.39% compared to random sampling, demonstrating OTI’s effectiveness.

**Under single-surrogate targeted attacks.** This experiment evaluates OTI to identify image attackability to targeted attacks, following a process similar to that used for




Figure 8: ASRs under varying sampling rates  $\alpha$  in targeted attacks. The x-axis represents the target models, with the last one indicating the average.

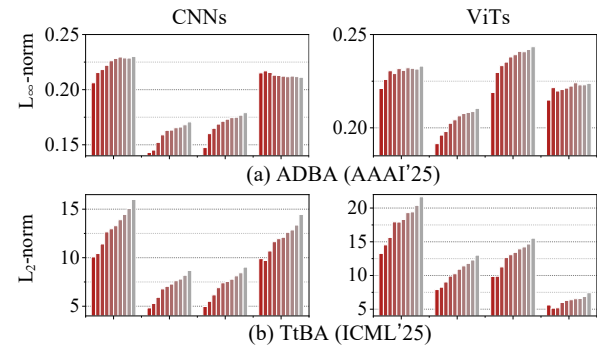


Figure 9: The query-based perturbation norm required for samples at different sampling rates. Each group corresponds to a target model. With each group, the 10 bars represent sampling rates ranging from 10% to 100% in 10% intervals.

untargeted attacks. The surrogate model is R50 and the perturbation budget is  $L_\infty = 16/255$ . The results shown in Figure 8 indicate that, despite the increased difficulty of targeted attacks, subsets selected based on OTI achieve higher ASRs. On average, at  $\alpha = 10\%$ , OTI-based sampling improves the ASR by 6.79% compared to random sampling.

**Under ensemble-based untargeted attacks.** This experiment investigates OTI to evaluate attackability to ensemble-based attacks. Four advanced ensemble-based benchmarks and three ensemble surrogate models are as described in Setup. The perturbation budget is  $L_\infty = 10/255$ . Experimental results presented in Figure 6 indicate that OTI is also effective in evaluating image attackability to ensemble-based attacks, achieving an average improvement of 13.55% over random sampling at  $\alpha = 10\%$ .




Figure 10: Visualization of attacks on Kvasir-SEG. The first row shows high-OTI samples, and the second row shows low-OTI samples. White text indicates OTI values, and yellow text indicates F1-scores. Each group shows: benign sample, groundtruth, benign prediction and attacked prediction. The surrogate model is URPC and target model is U-Net.




Figure 11: Ablation study of OTI. The surrogate models are ResNet-50 (top row) and ViT-B/16 (bottom row). The x-axis indicates the sampling rate, and the y-axis shows the average ASR across 8 targets.

**On attacking adversarially trained models.** In the above experiments, target models are all undefended. To test OTI’s performance on defended models, we use R50 as the surrogate model and evaluate on six advanced adversarially trained models as described in Setup. Results in Figure 7 confirm that OTI reliably assesses image attackability even with defenses. On average, OTI-based sampling achieves a 12.24% higher ASR than random sampling at  $\alpha = 10\%$ .

**Query optimization in query-based attacks.** Since OTI effectively measures image attackability, more attackable images require less perturbation. Intuitively, OTI-selected vulnerable images need smaller perturbations. We evaluate on two SOTA query-based methods with  $L_\infty$  and  $L_2$  norms respectively. The dataset used is the subset of ImageNet Validation. Results in Figure 9 confirm our thoughts.

**On attacking gastrointestinal polyp segmentation.** Extensive experiments have shown OTI’s effectiveness in classification tasks. We further investigate whether OTI remains effective in segmentation tasks with non-natural images, and whether it performs well with a smaller dataset scale compared to previous datasets of 50K and 1K images. To this end, we test on the 200-image Kvasir-SEG dataset using the MI-FGSM attack with a perturbation budget of  $L_\infty = 4/255$ . The surrogate model is U-Net. The target models are shown in Setup. Visualization results are shown in Figure 10, with detailed data in Figure 12. Both confirm that OTI remains effective in non-natural image segmentation tasks.

### 4.3 Ablation Study

This experiment aims to ablate OTI to explore the relationship among OAR, ITI, and OTI. We use untargeted single-




Figure 12: Metrics of MI-FGSM based Kvasir-SEG adversarial examples at different sampling rates. The x-axis represents sampling rates from 10% to 100% in 10% intervals.

surrogate attacks as the benchmark. The results are shown in Figure 9. It can be observed that in all cases, OTI outperforms OAR and ITI, indicating that OAR and ITI are complementary robustness features. Additionally, OAR generally outperforms ITI because OAR focuses on the quantity of semantically relevant elements, while ITI considers global and possibly semantic irrelevant texture information. OAR, ITI, and OTI all significantly outperform random strategy, demonstrating their effectiveness.

## 5 Conclusions and Limitations

This paper proposes OTI, a model-free and visually interpretable measure for evaluating the benign image attackability to adversarial perturbations. Our OTI aligns with human intuition and performs well across various tasks, domains, and data sizes. However, it is limited to image data, making the development of similar measures for audio, text, and other modalities a valuable future direction.

## Acknowledgments

This work was supported in part by the Science and Technology Development Fund, Macau SAR, under Grant 0193/2023/RIA3 and 0079/2025/AFJ, and the University of Macau under Grant MYRG-GRG2024-00065-FST-UMDF.

## References

- Bai, Y.; Zhou, M.; Patel, V. M.; and Sojoudi, S. 2024. Mixednuts: Training-free accuracy-robustness balance via nonlinearly mixed classifiers. *arXiv preprint arXiv:2402.02263*.
- Byun, J.; Kwon, M.-J.; Cho, S.; Kim, Y.; and Kim, C. 2023. Introducing competition to boost the transferability of targeted adversarial examples through clean feature mixup. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24648–24657.
- Chaouai, Z.; and Tamaazousti, M. 2024. Universal Robustness via Median Randomized Smoothing for Real-World Super-Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9059–9068.
- Chen, B.; Yin, J.; Chen, S.; Chen, B.; and Liu, X. 2023. An adaptive model ensemble adversarial attack for boosting adversarial transferability. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4489–4498.
- Chen, H.; Zhang, Y.; Dong, Y.; Yang, X.; Su, H.; and Zhu, J. 2024. Rethinking Model Ensemble in Transfer-based Adversarial Attacks. In *The Twelfth International Conference on Learning Representations*.
- Chen, L.; Yang, H.; Wang, H.; Cao, Y.; Li, S.; and Luo, Y. 2025. Intriguing Frequency Interpretation of Adversarial Robustness for CNNs and ViTs. *arXiv preprint arXiv:2506.12875*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9185–9193.
- Elsayed, G.; Krishnan, D.; Mobahi, H.; Regan, K.; and Bengio, S. 2018. Large margin deep networks for classification. *Advances in neural information processing systems*, 31.
- Fan, D.-P.; Ji, G.-P.; Cheng, M.-M.; and Shao, L. 2021. Concealed object detection. *IEEE transactions on pattern analysis and machine intelligence*, 44(10): 6024–6042.
- Fang, Z.; Wang, R.; Huang, T.; and Jing, L. 2024. Strong transferable adversarial attacks via ensembled asymptotically normal distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24841–24850.
- Gonzales, R. C.; and Wintz, P. 1987. *Digital image processing*. Addison-Wesley Longmagonzales1987digitaln Publishing Co., Inc.
- Jha, D.; Smedsrud, P. H.; Riegler, M. A.; Halvorsen, P.; De Lange, T.; Johansen, D.; and Johansen, H. D. 2019. Kvasir-seg: A segmented polyp dataset. In *International conference on multimedia modeling*, 451–462. Springer.
- Karunanayake, N.; Silva, B.; Ginige, Y.; Seneviratne, S.; and Chawla, S. 2025. Quantifying and Exploiting Adversarial Vulnerability: Gradient-Based Input Pre-Filtering for Enhanced Performance in Black-Box Attacks. *ACM Transactions on Privacy and Security*.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2018. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, 99–112. Chapman and Hall/CRC.
- Levy, M.; Amit, G.; Elovici, Y.; and Mirsky, Y. 2024. Ranking the Transferability of Adversarial Examples. *ACM Transactions on Intelligent Systems and Technology*, 15(5): 1–21.
- Li, Q.; Guo, Y.; Zuo, W.; and Chen, H. 2023. Improving adversarial transferability via intermediate-level perturbation decay. *Advances in Neural Information Processing Systems*, 36: 32900–32912.
- Li, T.; Ma, X.; and Jiang, Y.-G. 2025. AIM: Additional Image Guided Generation of Transferable Adversarial Attacks. *arXiv preprint arXiv:2501.01106*.
- Li, Y.; Angelov, P.; and Suri, N. 2024. Self-supervised representation learning for adversarial attack detection. In *European Conference on Computer Vision*, 236–252. Springer.
- Li, Z.; Wang, W.; Li, J.; Chen, K.; and Zhang, S. 2024. Ucg: A universal cross-domain generator for transferable adversarial examples. *IEEE Transactions on Information Forensics and Security*, 19: 3023–3037.
- Liang, J.; and Pun, C.-M. 2025. IC Attack: In-place and Cross-pixel Augmentations for Highly Transferable Transformation-based Attacks. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 5345–5354.
- Lin, Q.; Luo, C.; Niu, Z.; He, X.; Xie, W.; Hou, Y.; Shen, L.; and Song, S. 2024. Boosting adversarial transferability across model genus by deformation-constrained warping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3459–3467.
- Liu, C.; Dong, Y.; Xiang, W.; Yang, X.; Su, H.; Zhu, J.; Chen, Y.; He, Y.; Xue, H.; and Zheng, S. 2025. A comprehensive study on robustness of image classification models: Benchmarking and rethinking. *International Journal of Computer Vision*, 133(2): 567–589.
- Liu, N.; Zhang, N.; Wan, K.; Shao, L.; and Han, J. 2021. Visual saliency transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4722–4732.
- Liu, Y.; Chen, X.; Liu, C.; and Song, D. 2016. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*.
- Ma, A.; Farahmand, A.-m.; Pan, Y.; Torr, P.; and Gu, J. 2024. Improving adversarial transferability via model alignment. In *European Conference on Computer Vision*, 74–92. Springer.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.

- Pawlowski, N.; Bhooshan, S.; Ballas, N.; Ciompi, F.; Glocker, B.; and Drozdal, M. 2019. Needles in haystacks: On classifying tiny objects in large images. *arXiv preprint arXiv:1908.06037*.
- Peng, S.; Xu, W.; Cornelius, C.; Hull, M.; Li, K.; Duggal, R.; Phute, M.; Martin, J.; and Chau, D. H. 2023. Robust principles: Architectural design principles for adversarially robust cnns. *arXiv preprint arXiv:2308.16258*.
- Raina, V.; and Gales, M. 2023. Identifying Adversarially Attackable and Robust Samples. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*.
- Reza, M. F.; Jin, R.; Wu, T.; and Dai, H. 2025. GSBA  $K$ : top- $K$  Geometric Score-based Black-box Attack. In *The Thirteenth International Conference on Learning Representations*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Singh, N. D.; Croce, F.; and Hein, M. 2023. Revisiting adversarial training for imagenet: Architectures, training and generalization across threat models. *Advances in Neural Information Processing Systems*, 36: 13931–13955.
- Song, K.; Lai, H.; Pan, Y.; and Yin, J. 2024. Mimicdiffusion: Purifying adversarial perturbation via mimicking clean diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24665–24674.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tang, B.; Wang, Z.; Bin, Y.; Dou, Q.; Yang, Y.; and Shen, H. T. 2024. Ensemble diversity facilitates adversarial transferability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24377–24386.
- Wang, F.; Zuo, X.; Huang, H.; and Chen, G. 2025a. Adba: Approximation decision boundary approach for black-box adversarial attacks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7628–7636.
- Wang, F.; Zuo, X.; Huang, H.; and Chen, G. 2025b. TtBA: Two-third Bridge Approach for Decision-Based Adversarial Attack. In *Forty-second International Conference on Machine Learning*.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8684–8694.
- Wang, J.; Chen, Z.; Jiang, K.; Yang, D.; Hong, L.; Guo, P.; Guo, H.; and Zhang, W. 2024a. Boosting the transferability of adversarial attacks with global momentum initialization. *Expert Systems with Applications*, 255: 124757.
- Wang, K.; He, X.; Wang, W.; and Wang, X. 2024b. Boosting adversarial transferability by block shuffle and rotation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24336–24346.
- Wang, X.; Zhang, Z.; and Zhang, J. 2023. Structure invariant transformation for better adversarial transferability. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4607–4619.
- Wei, Z.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2023. Enhancing the self-universality for transferable targeted attacks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12281–12290.
- Weng, J.; Luo, Z.; Li, S.; Sebe, N.; and Zhong, Z. 2023. Logit margin matters: Improving transferable targeted adversarial attack by logit calibration. *IEEE Transactions on Information Forensics and Security*, 18: 3561–3574.
- Xie, Y.; Zheng, W.; Huang, H.; Ye, G.; and Ma, X. 2025. Towards million-scale adversarial robustness evaluation with stronger individual attacks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 30702–30711.
- Xu, K.; Zhang, C.; Chen, Z.; Wang, Z.; Xiao, C.; and Liang, C. 2025. Rethinking the Adversarial Robustness of Multi-Exit Neural Networks in an Attack-Defense Game. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 10265–10274.
- Yan, J.; Yin, H.; Zhao, Z.; Ge, W.; and Zhang, J. 2024. Enhance adversarial robustness via geodesic distance. *IEEE Transactions on Artificial Intelligence*, 5(8): 4202–4216.
- Zeng, H.; Chen, B.; and Peng, A. 2024. Enhancing targeted transferability via feature space fine-tuning. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4475–4479. IEEE.
- Zhang, H.; Qin, Z.; Gao, S.; and Clark, A. 2024a. Seev: Synthesis with efficient exact verification for relu neural barrier functions. *Advances in Neural Information Processing Systems*, 37: 101367–101392.
- Zhang, L.; Luo, Y.; Shen, H.; Wang, T.; and Larson, K. 2024b. A fourier perspective of feature extraction and adversarial robustness. In *International Joint Conferences on Artificial Intelligence Organization*, 1715–1723.
- Zhou, W.; Hou, X.; Chen, Y.; Tang, M.; Huang, X.; Gan, X.; and Yang, Y. 2018. Transferable adversarial perturbations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 452–467.
- Zhu, H.; Liang, S.; Wang, W.; Zhang, Z.; Li, F.; and Wang, S.-L. 2025. Evading Data Provenance in Deep Neural Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1249–1260.