

Robust SAM: On the Adversarial Robustness of Vision Foundation Models

Jiahuan Long^{1, 2, 3}, Zhengqin Xu¹, Tingsong Jiang^{2, 3}, Wen Yao^{2, 3*}, Shuai Jia¹,
Chao Ma^{1*}, Xiaoqian Chen^{3, 4}

¹MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University,

²Defense Innovation Institute, Chinese Academy of Military Science,

³Intelligent Game and Decision Laboratory,

⁴Chinese Academy of Military Science,

jiahuanlong@sjtu.edu.cn, wendy0782@126.com, chaoma@sjtu.edu.cn

Abstract

The Segment Anything Model (SAM) is a widely used vision foundation model with diverse applications, including image segmentation, detection, and tracking. Given SAM’s wide applications, understanding its robustness against adversarial attacks is crucial for real-world deployment. However, research on SAM’s robustness is still in its early stages. Existing attacks often overlook the role of prompts in evaluating SAM’s robustness, and there has been insufficient exploration of defense methods to balance the robustness and accuracy. To address these gaps, this paper proposes an adversarial robustness framework designed to evaluate and enhance the robustness of SAM. Specifically, we introduce a cross-prompt attack method to enhance the attack transferability across different prompt types. Besides attacking, we propose a few-parameter adaptation strategy to defend SAM against various adversarial attacks. To balance robustness and accuracy, we use the singular value decomposition (SVD) to constrain the space of trainable parameters, where only singular values are adaptable. Experiments demonstrate that our cross-prompt attack method outperforms previous approaches in terms of attack success rate on both SAM and SAM 2. By adapting only 512 parameters, we achieve at least a 15% improvement in mean intersection over union (mIoU) against various adversarial attacks. Compared to previous defense methods, our approach enhances the robustness of SAM while maximally maintaining its original performance.

Introduction

SAM is a popular vision foundation model for image segmentation (Kirillov et al. 2023). It predicts the object mask of an image based on box or point prompts provided by the user. SAM has been successful in this area, leading to efforts to use it for more complex tasks beyond segmentation, such as tracking and detection (Yang et al. 2023; Ren et al. 2024; Yu et al. 2023). Recently, Meta pushed the boundaries further with the introduction of an even more capable and versatile successor to SAM, known as Segment Anything Model 2 (SAM 2), designed specifically for video segmentation (Ravi et al. 2024). However, existing studies have revealed that SAM is vulnerable to attacks by adversarial

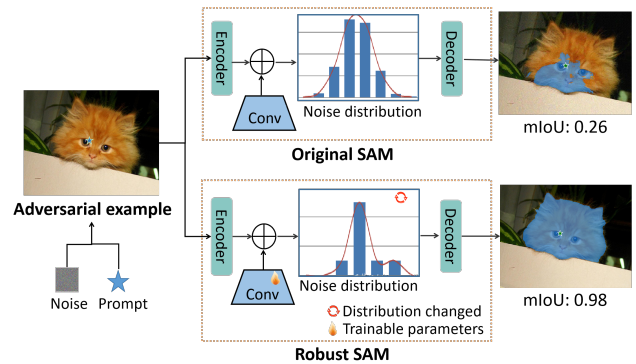


Figure 1: A comparison of the adversarial robustness of the original SAM and the proposed Robust SAM. When implementing attacks, the original SAM predicts a segmentation mask with severely compromised precision (i.e. 0.26 mIoU with the ground truth), whereas our Robust SAM maximally preserves the quality of the segmentation mask (i.e. 0.98 mIoU with the ground truth) via modified noise distribution.

examples (Zhang et al. 2023; Qiao et al. 2023; Zheng and Zhang 2023; Wang, Zhao, and Petzold 2024), which are subtle perturbations added to a benign image. These adversarial examples significantly degrade the segmentation performance of SAM. Since SAM has a wide range of applications, it is crucial to examine its adversarial robustness.

Recent studies have pointed out that the effectiveness of adversarial attacks against SAM is closely related to the type of input prompts (Dai et al. 2023; Huang et al. 2023). Indeed, our preliminary experiments indicate that existing attacks exhibit poor transferability across different types of prompts. For example, the PGD attack (Madry et al. 2017) can disrupt the segmentation mask created by a point prompt but not the one generated from a box prompt. Therefore, exploring adversarial attacks against SAM that simultaneously work on different prompt types will offer further insights into SAM’s robustness.

Beyond attack methods, it is crucial to investigate the defense mechanisms of the SAM. Previous studies have proposed different approaches to defend against adversarial examples (Dziugaite, Ghahramani, and Roy 2016; Osadchy et al. 2017; Guo et al. 2020; Ross and Doshi-Velez 2018).

*Corresponding authors.

Among these defense, full-parameter adversarial training is a commonly researched and proven effective method (Ganin et al. 2016; Bai et al. 2021) to increase a model’s robustness. However, performing full-parameter adversarial training on large pretrained model like SAM is not only computationally expensive and prone to catastrophic forgetting (Luo et al. 2023; Zhai et al. 2023). Therefore, exploring few-parameter adaptation methods to bolster the robustness of large models is both meaningful and necessary.

There are some parameter adaptation schemes that aim to improve the performance of SAM in various downstream applications like medical image segmentation (Zhang and Liu 2023; Ma et al. 2024) and camouflage detection (Chen et al. 2023). Yet, a direct adoption of these methods to enhance adversarial robustness faces non-trivial challenges. One major issue is that these adaptation schemes tend to prioritize precision on adversarial examples at the expense of accuracy on clean datasets. This trade-off is especially apparent in the case of the SA1B dataset (Kirillov et al. 2023), which serves as SAM’s foundational training dataset. Therefore, it is crucial to develop a few-parameter defense method that can effectively balance SAM’s robustness and accuracy.

This work focuses on enhancing the robustness of SAM against adversarial attacks while minimizing the impact on its baseline performance. This is achieved by training a limited set of parameters. As shown in Figure 1, selectively adapting a small portion of parameters results in a more robust SAM variant (i.e., RobustSAM), which outperforms the original SAM in segmentation under attack. Initially, our approach begins with the design of a cross-prompt adversarial attack to generate adversarial examples. These examples can effectively attack SAM under both point and box prompts, enabling a more comprehensive evaluation of SAM’s robustness. To further enhance the robustness of SAM, we propose a few-parameter adversarial defense on the adversarial examples. By adapting the convolutional layers within SAM, we effectively alter the feature distribution, thereby establishing an adversarial defense against various adversarial attacks. To minimize the number of trainable parameters, we employ Singular Value Decomposition (SVD) (Andrews and Patterson 1976) to reconstruct the parameters space of the convolution layers. By adapting only 512 singular values, we significantly improve the SAM’s robustness while maximally maintaining its original segmentation ability.

- We design a cross-prompt adversarial attack that updates noise to disrupt common key features of both point and box prompts. Compared to previous SAM attacks, it can achieve higher attack transferability across different types of prompts.
- We develop a new SVD-based defense method for vision foundation models that enhances SAM’s robustness while preserving its original performance by adapting only a small fraction of the parameters.
- We validate the proposed attack and defense methods on SAM for image segmentation, demonstrating their effectiveness. Additionally, we extend the attack to SAM 2, successfully compromising its video segmentation capabilities.

Methods	[1]	[2]	[3]	[4]	[5]	[6]	Ours
New attack method	✓	✓	✓	✓	✓		✓
New defense method						✓	✓
Attack transferability							✓
SAM1 & 2							✓
Computation efficiency						✓	✓
Robustness tradeoff							✓

Table 1: A comparison of representative works on SAM’s robustness. [1]-[6] represents the references (Zhang et al. 2023; Qiao et al. 2023; Zheng and Zhang 2023; Wang, Zhao, and Petzold 2024; Jankowski et al. 2024; Chen et al. 2024), respectively.

Related Work

Attacking SAM. As the Segment Anything Model (SAM) becomes increasingly popular in the computer vision community, its robustness against adversarial attacks begins to draw greater research attention (Zhang et al. 2023; Qiao et al. 2023; Zheng and Zhang 2023; Wang, Zhao, and Petzold 2024; Wang et al. 2023a; Jankowski et al. 2024). Zhang et al. (Zhang et al. 2023) conduct the first investigation on how to attack SAM with adversarial examples. After this work, Qiao et al. (Qiao et al. 2023) conducted an extensive study to evaluate the robustness of SAM, including various corruptions, occlusions, and perturbations. Zheng et al. (Zheng and Zhang 2023) achieve targeted adversarial attacks on SAM with black-box settings. However, these methods exhibit poor attack transferability under different prompt types. In this paper, our method can impose effective attacks against SAM under either point or box prompts.

Adapting SAM. Existing parameter adaptation methods for SAM (Wu et al. 2023; Chen et al. 2023; Zhang and Liu 2023; Peng et al. 2024; Ma et al. 2024) mainly focus on improving its performance in downstream applications, such as medical image segmentation, camouflage detection. For example, SAMed (Zhang and Liu 2023) customize SAM to the medical image segmentation task using a low-rank adaptation (LoRA) (Hu et al. 2021). Similarly, SAM-Adapter (Chen et al. 2023) introduce new adapter layers into SAM’s transformer blocks for camouflage detection. However, applying these adaptation methods to improve SAM’s robustness can result in overfitting on adversarial datasets, which in turn may compromise its performance on clean datasets.

To address the above research gaps, we propose a comprehensive robustness framework for SAM that integrates a cross-prompt adversarial attack and an SVD-based defense method. Unlike existing approaches, our attack not only targets image and video segmentation on SAM and SAM 2 but also demonstrates higher attack transferability across different types of prompts. Additionally, our few-parameter defense method effectively preserves SAM’s original performance on clean datasets while enhancing its robustness against adversarial datasets. A comparison of the features and contributions of our method against prior works is provided in Table 1.

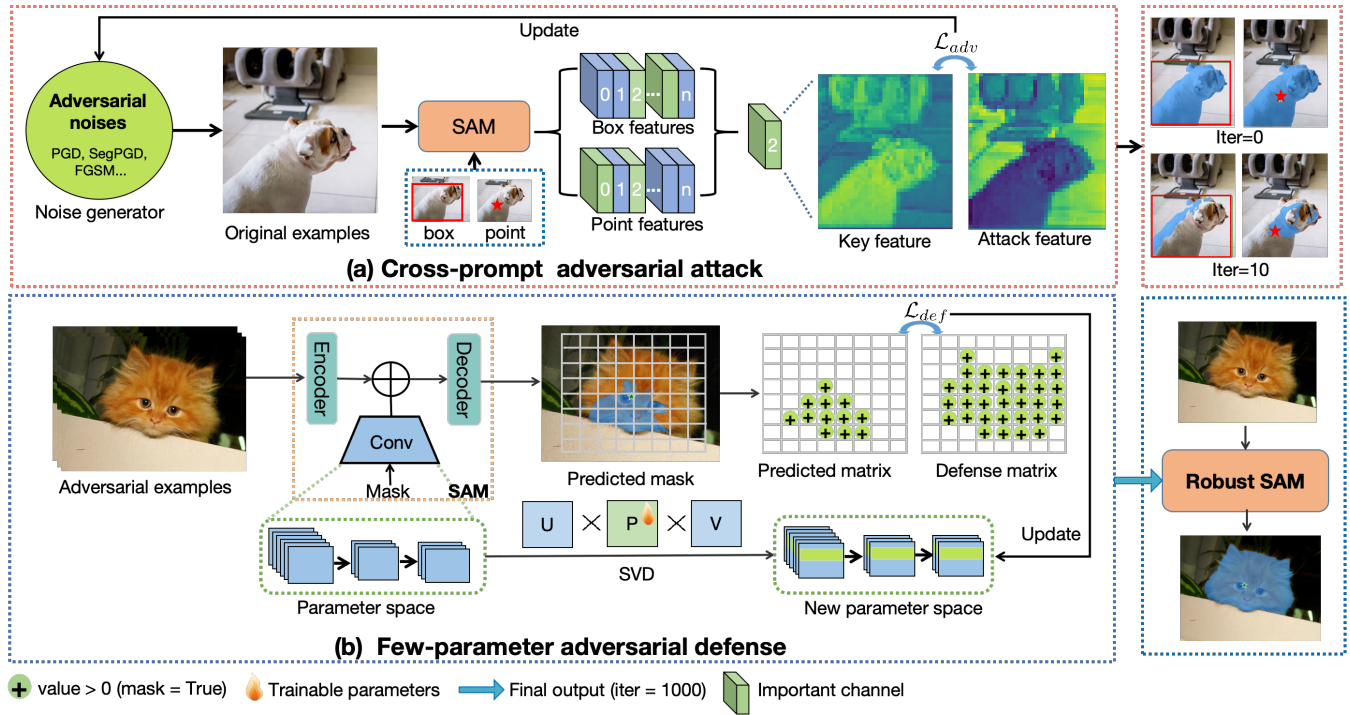


Figure 2: The overall framework of our proposed adversarial attack and defense pipelines. (a) illustrates one iteration of the cross-prompt adversarial attack. Original examples are perturbed by a noise generator, and these perturbed examples are then input into SAM, disrupting key features for both box and point prompts. The effectiveness and progression of the attack are demonstrated over multiple iterations. (b) depicts the adversarial defense process. Within the SAM model, the parameters of the convolutional layer (Conv) are decomposed via the singular value decomposition (SVD). Only the matrix P in the new parameter space is updated, while other parameters in the model remain frozen. After the parameter adaptation, the fine-tuned SAM is validated using adversarial examples to evaluate its robustness.

Methodology

Cross-Prompt Adversarial Attack

Attack Define. Previous attacks on image segmentation models have primarily aimed to alter the category labels of predicted segmentation masks (Rossolini et al. 2023; Nesti et al. 2022). However, SAM’s predicted masks do not include categorical labels; instead, they indicate only the shape of plausible objects. As such, our attack is formulated with the objective of disrupting the shape of the segmentation masks generated by SAM. Moreover, as SAM works under the guidance of two types of user-provided prompt (i.e., point and box prompts), the proposed attack shall generalize to both types of prompts.

Consider a fixed set of images $\mathcal{X} = \{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^N\}$, where $N \in \mathbb{N}$ is the total number of images. For any two masks A and B corresponding to the same image \mathbf{X}^i , their mean Intersection over Union (mIoU) is defined as:

$$\text{mIoU}(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (1)$$

The indicator function $I(i)$ for image \mathbf{X}^i is defined as:

$$I(i) = \begin{cases} 1 & q^{\text{point}}(i) < 0.5 \ \& \ q^{\text{box}}(i) < 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where

$$q^{\text{point}}(i) = \text{mIoU}(S^{\text{point}}(\mathbf{X}_{adv}^i), S^{\text{point}}(\mathbf{X}^i)), \quad (3)$$

$$q^{\text{box}}(i) = \text{mIoU}(S^{\text{box}}(\mathbf{X}_{adv}^i), S^{\text{box}}(\mathbf{X}^i)). \quad (4)$$

Here, \mathbf{X}_{adv}^i is the adversarial example for benign image \mathbf{X}^i ; $S^{\text{point}}(\mathbf{X}_{adv}^i)$ and $S^{\text{box}}(\mathbf{X}_{adv}^i)$ denote SAM’s predicted masks for \mathbf{X}_{adv}^i under point and box prompts, respectively; likewise, $S^{\text{point}}(\mathbf{X}^i)$ and $S^{\text{box}}(\mathbf{X}^i)$ denote SAM’s predicted masks for \mathbf{X}^i under point and box prompts, respectively. Essentially, the indicator function defines an attack to be successful only if the mIoU for the predicted masks of the original image and its corresponding adversarial examples concurrently fall below 0.5 under the point and box prompts. Based on the indicator function, we define the attack success rate (ASR) as $ASR = \frac{1}{N} \sum_{i=1}^N I(i)$.

Attack Motivation. Existing works (Wang et al. 2023b, 2021) indicate that disrupting the important features can dramatically reduce the model accuracy. Building on this insight, **our cross-prompt attack is designed to affect the common key features of the model across both point and box prompts, thereby enhancing the attack transferability.** To achieve this, we propose a novel key feature selection algorithm, detailed as follows.




Figure 3: Demonstrations of output differences for point and box prompts. It illustrates the negative impact on SAM’s output when the feature map of each individual channel is set to zero (i.e., $f_i(\mathbf{X}) = 0$). **Red boxes** highlight the key features that have the greatest impact on mIoU for the point and box prompts.

Key Feature Selection. SAM employs a Vision Transformer as its encoder to extract global image features and utilizes a decoder to convert these features into precise segmentation masks. Given an image \mathbf{X} , the SAM encoder outputs a set of features $\mathcal{F} = \{f_i(\mathbf{X})\}$, where $i \in \{1, 2, \dots, N\}$, with N representing the total number of channels. To evaluate the importance of each feature, we generate modified feature sets $\{\tilde{\mathcal{F}}_i\}$, where each set $\tilde{\mathcal{F}}_i$ is constructed by setting feature $f_i(\mathbf{X})$ to 0 while keeping all other features unchanged. For each channel i , we calculate the output differences for the point prompt and box prompt:

$$\Delta y_{\text{point},i} = \text{Decoder}_{\text{point}}(\mathcal{F}) - \text{Decoder}_{\text{point}}(\tilde{\mathcal{F}}_i), \quad (5)$$

$$\Delta y_{\text{box},i} = \text{Decoder}_{\text{box}}(\mathcal{F}) - \text{Decoder}_{\text{box}}(\tilde{\mathcal{F}}_i), \quad (6)$$

The channels of common key features are calculated by:

$$C_{\text{point}} = \text{TopK}(\Delta y_{\text{point},i}), \quad C_{\text{box}} = \text{TopK}(\Delta y_{\text{box},i}), \quad (7)$$

$$C_{\text{common}} = C_{\text{point}} \cap C_{\text{box}}, \quad (8)$$

where the function TopK selects top K feature channels with the largest output differences. Figure 3 provides a specific example, illustrating the output differences for both point and box prompts when each feature channel is set to zero.

Attack Pipeline. As shown in Figure 2(a), SAM extracts two types of features—point features and box features—from an input image. Although these two types of features are fundamentally similar, they contribute differently to SAM’s output depending on the prompt type. Next, we use our key feature selection algorithm to extract the common key features from both types. Subsequently, we define \mathbf{A} as an attack feature that uses negative values to disrupt the key features, leading to less effective feature representations. Finally, we optimize the adversarial noise using \mathcal{L}_{adv} , which quantifies the discrepancy between the key features and the attack features. \mathcal{L}_{adv} is given by:

$$\mathcal{L}_{adv} = \text{MSE}(\mathbf{A}_i, f_i(\mathbf{X} + \delta)), \quad i \in C_{\text{common}} \quad (9)$$

where $f_i(\mathbf{X} + \delta)$ represents the feature map of the i -th channel produced by SAM’s encoder given the adversarial image $\mathbf{X} + \delta$. The adversarial noise δ is updated in the next

iteration as $\delta^* = \text{Clip}_{\mathbf{X},\epsilon}(\delta + \alpha \cdot \text{sign}(\nabla_{\delta} \mathcal{L}_{adv}))$ with a step size α . The function $\text{Clip}_{\mathbf{X},\epsilon}(\cdot)$ constrains the noise within $[-\epsilon, \epsilon]$ and ensures the adversarial image $\mathbf{X} + \delta$ remains within the valid range of $[0, 255]$.

Few-Parameter Adversarial Defense

Defense Define. The objective of our defense method is to enhance SAM’s robustness, enabling it to recover the correct masks for adversarial examples. We define a set of adversarial examples $\mathcal{X} = \{\mathbf{X}_{adv}^1, \mathbf{X}_{adv}^2, \dots, \mathbf{X}_{adv}^N\}$, where $N \in \mathbb{N}$ is the total number of adversarial examples in the set. Our goal is to train a more robust version of SAM (denoted by RobustSAM) that maximizes $\text{mIoU}(\text{RS}^p(\mathcal{X}), \mathcal{GT})$ under different types of prompts p , where $p \in [\text{point}, \text{box}]$. $\text{RS}^p(\cdot)$ is the output masks of the RobustSAM under prompt p , and \mathcal{GT} denotes the set of ground truth masks corresponding to the adversarial examples in \mathcal{X} .

Defense Motivation. Existing works (Xie et al. 2019; Ilyas et al. 2019) indicate that adversarial noise, while almost imperceptible when added into the pixel space, induces apparent perturbations in the feature space. As illustrated in Figure 4, the feature map of a clean image (a) differs significantly from that of an image corrupted by adversarial noise (b), where the texture features of the “dog” and “cat” are obscured by the added noise. As such, **the motivation behind our defense method is to mitigate the noise impact by altering the feature distribution of the adversarial example, with minimal parameter adaptation.** Figure 4 (c) demonstrates that the proposed RobustSAM maintains the segmentation quality of adversarial examples by adjusting the feature distribution.

Trainable Parameters Selection. To achieve our goal of improving SAM’s robustness with minimal parameter adaptations, selecting the appropriate trainable parameters is crucial. The convolutional layers, as the final stage of SAM’s feature extraction network (i.e., encoder), play an important role to adjusting feature distribution. These layers not only reduce the dimensionality of the feature space but also contain relatively few parameters ($\approx 600\text{K}$), accounting for just 1/100th of SAM’s total parameters.

Additionally, we can further compress these trainable parameters using Singular Value Decomposition (SVD). In our approach, the parameters of the convolutional layers are decomposed into three matrices: \mathbf{U} , \mathbf{P} , and \mathbf{V} . \mathbf{U} and \mathbf{V} are orthogonal matrices that represent the bases for the input and transformed spaces, respectively, with each column capturing specific patterns or features in the input data. \mathbf{P} is a diagonal matrix containing singular values that scale these bases. The decomposition process can be formulated as:

$$\mathbf{W} = \sum_{i=1}^d \mathbf{u}_i p_i (\mathbf{v}_i)^T = \mathbf{U} \mathbf{P} \mathbf{V}^T, \quad (10)$$

where $\mathbf{W} \in \mathbb{R}^{d \times k}$ denotes the parameter space of a convolutional layer with d input channels and k output channels. $\mathbf{U} \in \mathbb{R}^{d \times d}$ and $\mathbf{V} \in \mathbb{R}^{k \times d}$ are the basis matrices while $\mathbf{P} = \text{diag}(p_1, p_2, \dots, p_i, \dots, p_d)$ is the diagonal matrix adjusting these bases. To sum up, adjusting the matrix




Figure 4: Effects of attack and defense on the intermediate features in SAM. Top to bottom: feature maps of (a) a clean image, (b) an image corrupted by adversarial noise, and (c) an image corrupted by adversarial noise but defended by our proposed defense method.

\mathbf{P} is equivalent to adjusting the entire parameter space of the convolutional layers in specific directions. Our defense improves the adversarial robustness of SAM by training only the diagonal matrix \mathbf{P} .

Defense Pipeline. As illustrated in Figure 2(b), RobustSAM encompasses two key processes: parameter adaptation and parameter decomposition. During parameter adaptation process, SAM takes as input an adversarial example and outputs a prediction mask for it. To encourage correct mask prediction by SAM, we create a defense matrix with positive values that correspond to the pixel locations of the ground truth masks. Based on the mean squared error (MSE), we then calculate a loss \mathcal{L}_{def} between the predicted matrix and the defense matrix. During the parameter decomposition process, we focus on the convolutional layers in SAM’s feature extraction network. Rather than adapting all parameters, we apply singular value decomposition (SVD) to decompose the convolutional layer’s parameter space into three matrices \mathbf{U} , \mathbf{P} , and \mathbf{V} . Among them, we specifically set diagonal matrix \mathbf{P} as the trainable parameters while other parameters in the model remain frozen. The optimization objective for the proposed adversarial defense is given by:

$$\mathcal{L}_{def} = \text{MSE}(\mathbf{D}, \text{RS}^p(\mathbf{X}_{adv})), \quad (11)$$

$$\mathbf{P}^* = \arg \min_{\mathbf{P}} \mathcal{L}_{def}, \quad (12)$$

where $\text{RS}^p(*)$ represents the output masks of RobustSAM under different types of prompts p , \mathbf{D} is the defense matrix with positive values according to the ground truth masks, and \mathbf{P} is a diagonal matrix in the new parameters space of RobustSAM.

Our defense improves the adversarial robustness of SAM by training only diagonal matrix \mathbf{P} . We stress that this design enjoys the following three merits: (1) *Varied feature distribution.* The matrix \mathbf{P} acts as a coefficient matrix that directly adjusts the feature space, allowing us to modify the feature distribution and reduce the impact of adversarial noise. (2) *Balanced robustness and performance.* By training only \mathbf{P} and keeping the foundational matrices \mathbf{U} and \mathbf{V} unchanged, we maintain the core feature space while adapting the model’s response to adversarial examples, prevent-

ing overfitting to noise. (3) *Reduced training cost.* Training only \mathbf{P} significantly lowers the training cost, as we improve SAM’s robustness with only 512 parameters, compared to the 93735K parameters of the model or 600K of the convolutional layers.

Experiments

Experimental Setup

Datasets and Metrics. To evaluate the robustness of SAM under different types of prompts (i.e., point and box prompts), we randomly sample 2000 images from the SA1B (Kirillov et al. 2023), VOC (Everingham et al. 2010), COCO (Lin et al. 2014), and DAVIS (Pont-Tuset et al. 2017) datasets. Each image is paired with a ground truth segmentation mask of a single object, along with corresponding prompts. The VOC dataset is split into 70% for training and 30 for evaluation. SA1B and VOC are used to assess SAM’s robustness under independent and identically distributed (i.i.d.) conditions, while COCO, DAVIS is used for out-of-distribution (o.o.d.) conditions. For attack, we use the attack success rate (ASR) and the mean Intersection over Union (mIoU) as the evaluation metric. For defense, mIoU is used as the evaluation metric.

Baselines. We choose two victim models for our experiments: SAM (Kirillov et al. 2023) and SAM 2 (Ravi et al. 2024). Existing attacks against SAM can be categorized based on the type of prompt inputs: point prompt attack (PPA) (Zhang et al. 2023; Wang et al. 2023a) and box prompt attack (BPA) (Qiao et al. 2023). These methods are variants of gradient-based adversarial attacks, such as PGD (Madry et al. 2017) and FGSM (Goodfellow, Shlens, and Szegedy 2014). For a fair comparison, PPA, BPA and our cross-prompt attack are configured with the same attack strength and number of iterations. For defense comparison, we introduce various parameter adaptation schemes to improve the robustness, including SAM-LoRA (Zhang and Liu 2023), SAM-Adapter (Wu et al. 2023), MedSAM (Ma et al. 2024) and SAM-COBOT (Peng et al. 2024). The main difference among these approaches is the selection of parameters within SAM for adaptation. For more details on these attack and defense methods, please refer to the Appendix.

Implementation Details. For the attack setting, we set the total number of iteration steps to 20, and perturbation intensity ϵ to 16/255 for PPA, BPA, and our cross-prompt attack. The attack feature \mathbf{A} are set to the negative values of the key features, and K in TOPK function is set to 5. For all few-parameter adaptation methods, we randomly sample 70% of the adversarial examples in the VOC dataset for adapting SAM. Our training employs the Adam optimizer (Kingma and Ba 2014). The initial learning rate is set to 1.0×10^{-3} , and the weight decay is 5×10^{-5} with one image per mini-batch. The number of training epochs is set to 500.

Evaluation of Cross-Prompt Adversarial Attack

Figure 5 qualitatively evaluates the performance of the proposed cross-prompt attack method. PPA (a) can effectively disrupt the mask predicted by SAM under point prompts

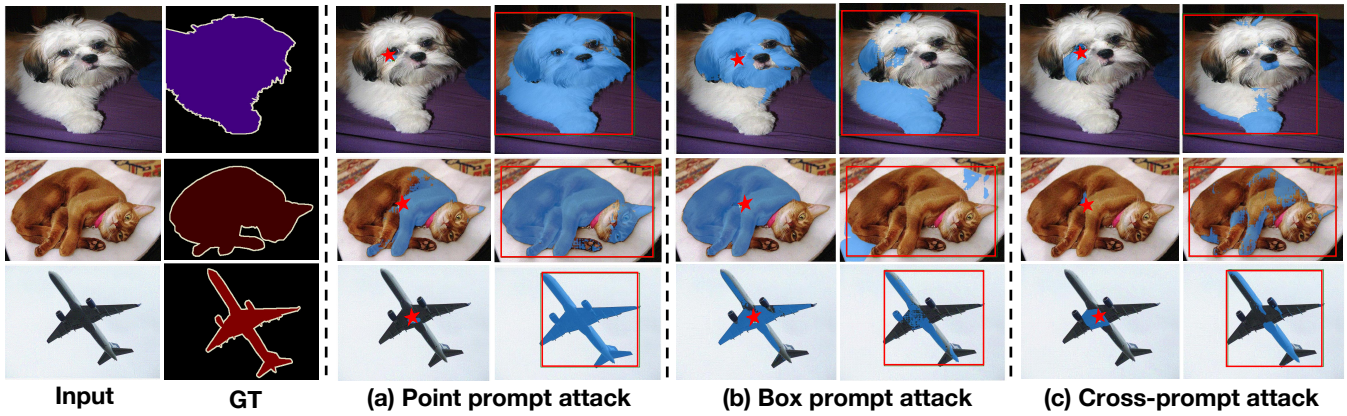


Figure 5: Illustrations of various adversarial attacks against SAM by employing both point and box prompts. An attack is considered successful if it removes more than 50% of the ground truth masks for either point or box prompts. (a) illustrates that the point prompt attacks effectively compromise the ground truth masks under point prompts but are ineffective when encountering box prompts. (b) depicts the box prompt attacks that successfully disrupt the ground truth masks with point prompts but fail to attack under point prompts. (c) showcases that our cross-prompt attack successfully degrades the performance for both point and box prompts simultaneously.

		COCO			VOC			
		ϵ	4	8	16	4	8	16
SAM	PPA		1.4	2.8	11.4	2.6	5.4	8.3
	BPA		1.9	12.8	35.7	2.4	19.1	44.4
	CPA (Ours)		5.7	15.7	60.2	6.2	30.3	62.9
SAM2	PPA		0	0.3	5.9	0.2	0.5	3.1
	BPA		0	6.6	18.5	0	9.8	20.0
	CPA (Ours)		1.3	8.5	32.8	2.2	18.7	34.4

Table 2: Comparisons of ASR with point prompt attack (PPA), box prompt attack (BPA), and our cross-prompt attack (CPA) on multiple datasets. ϵ denotes the perturbation intensity.

(i.e., the ground truth masks of “dog”, “cat” are removed under point prompts). However, PPA fails to achieve the same level of disruption under box prompts. BPA (b) effectively removes more than 50% of the ground truth mask, but fails to do so under a point prompt. In contrast, our cross-prompt attack (c) successfully disrupts the ground truth masks under both point and box prompts, outperforming PPA and BPA. This indicates a significant improvement in adversarial attacks against SAM.

Figure 6 illustrates our cross-prompt attack targeting video segmentation on SAM 2 (Ravi et al. 2024). It demonstrates that our attack can disrupt the temporal consistency of video segmentation. For example, it misleads SAM 2 into incorrectly tracking the race car’s mask as its wheels and causes the man’s mask to be tracked as the skateboard. Therefore, our attack effectively compromises the expected continuity and accuracy of video segmentation.

Table 2 quantitatively evaluate the performance of the proposed cross-prompt attack method on SAM (Kirillov et al. 2023) and SAM 2 (Ravi et al. 2024). Compared to




Figure 6: Illustration of our cross-prompt attack against the video segmentation on SAM 2 (refer to Supplementary Materials for a video demo).

point prompt attack (PPA) and box prompt attack (BPA), the cross-prompt attack exhibits a higher attack success rate across all datasets and attack intensities (i.e., $\epsilon = 4/255, 8/255, 16/255$).

Evaluation of Few-Parameter Adversarial Defense

Table 3 compares the effectiveness of various parameter adaptation methods against adversarial attacks, including BPA, PPA, and our cross-prompt attack (CPA). As shown, RobustSAM achieves the highest mIoU, manifesting its superiority over existing methods (i.e., MedSAM, SAM-Adapter, SAM-LoRA, SAM-COBOT) across VOC, SA1B and DAVIS datasets. We observe that under BPA and PPA attacks on the SA1B dataset, the mIoU of RobustSAM is slightly lower than that of SAM-LoRA. This may be attributed to LoRA’s low-rank adaptations (Hu et al. 2021),

	VOC			SA1B			DAVIS		
	BPA	PPA	CPA(ours)	BPA	PPA	CPA(ours)	BPA	PPA	CPA(ours)
Baseline [†]	16.9	29.5	6.4	17.2	23.7	8.1	18.3	30.1	7.5
MedSAM (Ma et al. 2024)	22.1	24.3	8.7	19.5	24.0	11.8	23.5	26.7	10.2
SAM-Adapter (Chen et al. 2023)	27.1	34.4	25.3	27.5	38.1	16.9	29.4	36.5	28.7
SAM-LoRA (Zhang and Liu 2023)	32.8	44.0	27.0	44.3	57.5	31.9	33.6	45.3	29.3
SAM-COBOT (Peng et al. 2024)	31.1	46.3	30.4	40.9	58.7	34.8	32.1	48.9	33.7
RobustSAM (Ours)	39.5	50.8	35.0	41.5	52.2	36.0	40.2	52.7	37.9

Table 3: Comparison of mIoU with different parameter adaptation methods on adversarial datasets. “Baseline[†]” denotes the original performance of the undefended SAM.

	VOC		SA1B	
	Point	Box	Point	Box
Baseline [†]	53.5	87.1	32.0	94.0
MedSAM	23.4	38.2	3.9	8.6
SAM-Adapter	62.7	57.9	13.1	49.4
SAM-LoRA	66.5	63.4	15.6	58.0
SAM-COBOT	68.6	61.1	19.7	59.4
RobustSAM (Ours)	71.3	80.9	28.7	74.0

Table 4: Comparison of mIoU with different parameter adaptation methods on clean datasets. “Baseline[†]” denotes the original performance of SAM. “Point” and “Box” represent the model outputs under point prompts and box prompts, respectively.

which might be more effective at capturing the feature characteristics specific to the SA1B dataset. Nonetheless, RobustSAM surpasses SAM-LoRA by a significant margin on the VOC and DAVIS datasets. Additionally, we found that MedSAM underperforms compared to other methods, likely due to its adaptation of both the encoder and decoder parameters in SAM, leading to severe overfitting. In conclusion, experiments demonstrates that enhancing the robustness of SAM does not depend on the number of trainable parameters, but rather on which parameters are adapted. This finding aligns with our previous defense theory, where we can enhanced robustness by modifying the feature distribution using a small fraction of the parameters.

Table 4 presents the fundamental performance comparison on the clean datasets after applying various parameter adaptation methods. The results demonstrate that RobustSAM can maximally preserve its fundamental segmentation ability on clean datasets. For instance, on the VOC dataset under Box prompts, SAM-LoRA reduces the baseline mIoU from 87.1 to 63.4 on the clean dataset, whereas our method results in a smaller decrease, from 87.1 to 80.9. This indicates that our defense successfully improves robustness on adversarial datasets while preserving strong segmentation performance on clean datasets.

Table 5 presents the computational comparison of different parameter adaptation method. The results show that the

Method	Params (K)	Time (Fps)
Baseline [†]	93735	3.3
MedSAM (Ma et al. 2024)	93729	3.3
SAM-Adapter (Wu et al. 2023)	3550	6.2
SAM-LoRA (Zhang and Liu 2023)	147	5.2
SAM-COBOT (Peng et al. 2024)	1.3	10.8
RobustSAM (Ours)	0.5	11.5

Table 5: Computation efficiency analysis for different parameter adaptation methods. “Baseline[†]”: adapting SAM’s all parameters.

complexity of our method is significantly lower than that of other parameter adaptation methods. For example, the parameters of RobustSAM are 600 times less than that of SAM-Adapter. This significant reduction in overhead stems from that our RobustSAM emphasizes adapting the core parameters, rather than adjusting a large number of parameters.

Conclusion

In this paper, we propose a comprehensive adversarial robustness framework for SAM including both adversarial attack and defense methods. The proposed cross-prompt attack effectively reveals the vulnerabilities of SAM and SAM 2 across different prompt types. To counter these vulnerabilities, we propose RobustSAM, a few-parameter defense approach that involves training a small subset of parameters in SAM’s feature extraction network. By applying SVD, we further control the number of trainable parameters, ensuring a balance between enhanced robustness and maintenance of the baseline performance of SAM. Our experimental results demonstrate that RobustSAM significantly boosts SAM’s robustness against adversarial attacks by adapting only 512 parameters. We expect that this work bring new insights into existing research on the robustness of SAM and serve as a baseline that facilitates further research in this direction.

Acknowledgements

This work was supported in part by NSFC (62322113, 62376156, 62401361).

References

- Andrews, H.; and Patterson, C. 1976. Singular value decomposition (SVD) image coding. *IEEE transactions on Communications*.
- Bai, T.; Luo, J.; Zhao, J.; Wen, B.; and Wang, Q. 2021. Recent advances in adversarial training for adversarial robustness. *arXiv preprint arXiv:2102.01356*.
- Chen, T.; Zhu, L.; Ding, C.; Cao, R.; Zhang, S.; Wang, Y.; Li, Z.; Sun, L.; Mao, P.; and Zang, Y. 2023. SAM Fails to Segment Anything?—SAM-Adapter: Adapting SAM in Underperformed Scenes: Camouflage, Shadow, and More. *arXiv preprint arXiv:2304.09148*.
- Chen, W.-T.; Vong, Y.-J.; Kuo, S.-Y.; Ma, S.; and Wang, J. 2024. RobustSAM: Segment Anything Robustly on Degraded Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dai, H.; Ma, C.; Liu, Z.; Li, Y.; Shu, P.; Wei, X.; Zhao, L.; Wu, Z.; Zhu, D.; Liu, W.; et al. 2023. Samaug: Point prompt augmentation for segment anything model. *arXiv preprint arXiv:2307.01187*.
- Dziugaite, G. K.; Ghahramani, Z.; and Roy, D. M. 2016. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (VOC) challenge. *International journal of computer vision*.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Guo, M.; Yang, Y.; Xu, R.; Liu, Z.; and Lin, D. 2020. When nas meets robustness: In search of robust architectures against adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, Y.; Cao, Y.; Li, T.; Juefei-Xu, F.; Lin, D.; Tsang, I. W.; Liu, Y.; and Guo, Q. 2023. On the robustness of segment anything. *arXiv preprint arXiv:2305.16220*.
- Ilyas, A.; Santurkar, S.; Tsipras, D.; Engstrom, L.; Tran, B.; and Madry, A. 2019. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*.
- Jankowski, K.; Sobieski, B.; Kwiatkowski, M.; Szulc, J.; Janik, M.; Baniecki, H.; and Biecek, P. 2024. Red-Teaming Segment Anything Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*.
- Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; and Zhang, Y. 2023. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint arXiv:2308.08747*.
- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment anything in medical images. *Nature Communications*.
- 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*.
- Nesti, F.; Rossolini, G.; Nair, S.; Biondi, A.; and Buttazzo, G. 2022. Evaluating the robustness of semantic segmentation for autonomous driving against real-world adversarial patch attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Osadchy, M.; Hernandez-Castro, J.; Gibson, S.; Dunkelman, O.; and Pérez-Cabo, D. 2017. No bot expects the Deep-CAPTCHA! Introducing immutable adversarial examples, with applications to CAPTCHA generation. *IEEE Transactions on Information Forensics and Security*.
- Peng, Z.; Xu, Z.; Zeng, Z.; Xie, L.; Tian, Q.; and Shen, W. 2024. Parameter efficient fine-tuning via cross block orchestration for segment anything model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3743–3752.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*.
- Qiao, Y.; Zhang, C.; Kang, T.; Kim, D.; Tariq, S.; Zhang, C.; and Hong, C. S. 2023. Robustness of sam: Segment anything under corruptions and beyond. *arXiv preprint arXiv:2306.07713*.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; et al. 2024. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*.
- Ross, A.; and Doshi-Velez, F. 2018. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Proceedings of the AAAI conference on artificial intelligence*.
- Rossolini, G.; Nesti, F.; D’Amico, G.; Nair, S.; Biondi, A.; and Buttazzo, G. 2023. On the real-world adversarial robustness of real-time semantic segmentation models for autonomous driving. *IEEE Transactions on Neural Networks and Learning Systems*.

Wang, A.; Islam, M.; Xu, M.; Zhang, Y.; and Ren, H. 2023a. Sam meets robotic surgery: an empirical study on generalization, robustness and adaptation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer.

Wang, D.; Yao, W.; Jiang, T.; and Chen, X. 2023b. Improving transferability of universal adversarial perturbation with feature disruption. *IEEE Transactions on Image Processing*.

Wang, Y.; Zhao, Y.; and Petzold, L. 2024. An empirical study on the robustness of the segment anything model (sam). *Pattern Recognition*, 110685.

Wang, Z.; Guo, H.; Zhang, Z.; Liu, W.; Qin, Z.; and Ren, K. 2021. Feature importance-aware transferable adversarial attacks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7639–7648.

Wu, J.; Fu, R.; Fang, H.; Liu, Y.; Wang, Z.; Xu, Y.; Jin, Y.; and Arbel, T. 2023. Medical sam adapter: Adapting segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.12620*.

Xie, C.; Wu, Y.; Maaten, L. v. d.; Yuille, A. L.; and He, K. 2019. Feature denoising for improving adversarial robustness. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.

Yang, J.; Gao, M.; Li, Z.; Gao, S.; Wang, F.; and Zheng, F. 2023. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*.

Yu, T.; Feng, R.; Feng, R.; Liu, J.; Jin, X.; Zeng, W.; and Chen, Z. 2023. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*.

Zhai, Y.; Tong, S.; Li, X.; Cai, M.; Qu, Q.; Lee, Y. J.; and Ma, Y. 2023. Investigating the catastrophic forgetting in multimodal large language models. *arXiv preprint arXiv:2309.10313*.

Zhang, C.; Zhang, C.; Kang, T.; Kim, D.; Bae, S.-H.; and Kweon, I. S. 2023. Attack-sam: Towards evaluating adversarial robustness of segment anything model. *arXiv preprint arXiv:2305.00866*.

Zhang, K.; and Liu, D. 2023. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*.

Zheng, S.; and Zhang, C. 2023. Black-box Targeted Adversarial Attack on Segment Anything (SAM). *arXiv preprint arXiv:2310.10010*.