Robustness of SAM: Segment Anything Under Corruptions and Beyond

Yu Qiao
Kyung Hee University

Chaoning Zhang[*]
Kyung Hee University

Taegoo Kang
Kyung Hee University

Donghun Kim
Kyung Hee University

Shehbaz Tariq
Kyung Hee University

Chenshuang Zhang
KAIST

Choong Seon Hong
Kyung Hee University

June 14, 2023

Abstract

Segment anything model (SAM), as the name suggests, is claimed to be capable of cutting out any object. SAM is a vision foundation model which demonstrates impressive zero-shot transfer performance with the guidance of a prompt. However, there is currently a lack of comprehensive evaluation of its robustness performance under various types of corruptions. Prior works show that SAM is biased towards texture (style) rather than shape, motivated by which we start by investigating SAM's robustness against style transfer, which is synthetic corruption. With the effect of corruptions interpreted as a style change, we further evaluate its robustness on 15 common corruptions with 5 severity levels for each real-world corruption. Beyond the corruptions, we further evaluate the SAM robustness on local occlusion and adversarial perturbations. Overall, this work provides a comprehensive empirical study on the robustness of the SAM under corruptions and beyond.

1 Introduction

Foundation models Bommasani et al. [2021] have made great strides in pushing the frontiers of modern AI. In the past few years, NLP has been revolutionized by large language models (LLMs), which are trained on abundant text corpora collected from the web. In contrast to BERT Devlin et al. [2018] requiring finetuning on the downstream tasks, GPT family models Brown et al. [2020], Radford et al. [2018, 2019] demonstrate strong zero-short (or few-shot) transfer performance on unseen data distributions and new tasks. The strong capability of zero-shot transfer of such text foundation models contributes to the development of various generative AI Zhang et al. [2023a] tasks, including text generation (ChatGPT Zhang et al. [2023b] for instance), text-to-image Zhang et al. [2023c], text-to-speech Zhang et al. [2023d], and text-to-3D Li et al. [2023]. Despite some progress like CLIP Radford et al. [2021], Jia et al. [2021], Yuan et al. [2021],the progress of foundation models in the computer vision Radford et al. [2021], Jia et al. [2021], Yuan et al. [2021] lags behind. Very recently, the Meta research team released the "Segment Anything" project with the goal to build a vision foundation model for segmentation.

With masked autoencoder Zhang et al. [2022] mimicking BERT Devlin et al. [2018] to provide a unified framework for self-supervised framework in NLP and vision, the success of segment anything model (SAM) Kirillov et al. [2023] has also been recognized by some researchers as the GPT moment for vision by using prompt Radford et al. [2019]. In other words, the vision community might follow NLP to go on a path of adopting the foundation model through prompt engineering. Numerous projects have combined the SAM with other models for more complex tasks beyond mask prediction, demonstrating its popularity and compatibility. Its interactive prompt design introduces greater flexibility to segmentation tasks. However, while SAM exhibits impressive zero-shot transfer performance and high compatibility with other models, concerns arise regarding its applicability to real-world scenarios. Particularly, its robustness against various corruptions, occlusions, and perturbations (such as adversarial attacks) remains unclear.

---

[*]You are welcome to contact us through chaoningzhang1990@gmail.com

(a) *content*   (b) *style*   (c) *blended*   (d) $Mask_{content}$ (e) $Mask_{predict}$

Figure 1: The mask results of SAM in style transfer images. Figure (a) refers to the content image with the location of the point prompt marked in a green star. Figure (b) refers to the style reference image. Figure (c) is the blended images after the style transferred from Figure (b) to Figure (a). The white area in Figures (d) and (e) refer to masks predicted by SAM based on images in Figures (a) and (c) with the given prompt, respectively. The results shown in Figure (e) indicate that compared to Figure (d), SAM is more susceptible to the influence of style transfer images, even in cases where the boundaries are visually clear and distinguishable to the human eye, as shown in Figure (c).

In this work, we conduct a comprehensive study on the robustness of SAM in the presence of various types of corruptions. Prior work shows that SAM is more biased towards texture (style) rather than shape Zhang et al. [2023e], motivated from which we first evaluate the model robustness against the style change. We observe that in most cases, the robustness of SAM drops when applied to images with different styles after style transfer, even when the transferred images retain clearly visible shapes to the human eye (a toy example in Figure 1). By perceiving various real-world corruptions as a new style Benz et al. [2021a], we further evaluate the robustness of SAM under various types of common corruptions Hendrycks and Dietterich [2019]. We find that apart from zoom blur corruption, SAM maintains relatively good performance on images corrupted by various common noises, blurring, and natural disturbances, which indicates that SAM is capable of executing robustly on real-world data with most of the common corruptions. Beyond corruptions, we further evaluate the robustness of SAM under local occlusion and adversarial perturbation. Our findings indicate that SAM exhibits robustness under mild occlusion, where 10%, 20%, and 40% of pixel values are replaced with Gaussian noise. However, when the occlusion ratio reaches 60% and above, SAM's performance experiences a significant decline across the majority of example images. For adversarial perturbation, we adhere to common practices used in popular adversarial attack methods such as FGSM Goodfellow et al. [2015] attack and PGD Madry et al. [2018] attack to assess the robustness of SAM against adversarial attacks. The results demonstrate that SAM has a moderate level of resilience against FGSM attack, but not PGD attacks, even for perturbation with a very small magnitude.

2 Related works

In less than a month since the advent of SAM, there has been a surge in projects and papers investigating it from different perspectives. These investigations can be roughly categorized as follows. A mainstream line of research has focused on evaluating the capability of SAM to segment various objects in real-world situations accurately. Several studies have investigated its performance on different types of images, including medical images Ma and Wang [2023], Zhang et al. [2023f], camouflaged objects Tang et al. [2023], and transparent objects Han et al. [2023]. The findings from these works consistently indicate that SAM frequently encounters difficulties in effectively detecting objects in such challenging scenarios. Another significant research direction has focused on enhancing SAM to improve its practicality. One notable approach in this line of work is Grounded SAM IDEA-Research [2023], which incorporates text inputs to enable the detection and segmentation of various objects. This is achieved by combining Grounding DINO Liu et al. [2023] with SAM, resulting in a more versatile and capable system for segmenting objects based on textual information. Due to the absence of label predictions in the generated masks by SAM, several studies Chen et al. [2023], Park [2023] have sought to integrate SAM with other models such as BLIP Li et al. [2022] or CLIP Radford et al. [2021]. The objective of these efforts is to leverage the strengths of these additional models to enhance the performance and accuracy of SAM in object segmentation tasks. By combining SAM with BLIP or CLIP, researchers aim to address

the limitation of SAM's generated masks without labels and achieve improved segmentation results. In addition to its applications in object segmentation, SAM has also been utilized in various other areas. Several works have explored the use of SAM for image editing purposes, including image editing techniques Rombach et al. [2022], as well as inpainting tasks Yu et al. [2023]. Furthermore, SAM has been employed in object tracking within videos Yang et al. [2023], Zxyang [2023], demonstrating its potential for visual tracking applications. Additionally, researchers have leveraged SAM in the field of 3D object reconstruction from a single image Shen et al. [2023], Kang et al. [2022], highlighting its utility in generating 3D models based on limited visual input. SAM has been reported in Tariq et al. [2023] to significantly improve semantic communication by only sending and receiving the foreground objects. For a survey of SAM, please refer to Zhang et al. [2023g] where the robustness of SAM against Gaussian noise is provided. Another work Zhang et al. [2023h] has also attacked SAM with adversarial examples, showing that SAM can be vulnerable to adversarial attacks. However, it remains to perform a comprehensive evaluation on the robustness of SAM under style transfer, corruption, occlusion, perturbation and beyond.

3 Experiment Setups

3.1 Evaluation Metric

Mask prediction.

SAM is a model designed to perform promptable segmentation by generating masks using both images and prompts as inputs, where the prompts are the necessary component for predictions. We denote a certain input image as x, the necessary prompts as $prompt$, and the model parameters as ω. It should be noted that the generated masks focus on object segmentation without providing semantic labels for individual masks. Therefore, for a given image, the prediction process as be formulated as follows:

$$mask_{i,j} = SAM(x, promt; \omega),$$ (1)

where $mask_{i,j}$ is the predicted mask with the shape of the original image size. The subscripts i and j indicate the coordinates of each pixel in the predicted mask. The pixel $x_{i,j}$ in the original image x is marked as part of the mask area if the predicted value $mask_{i,j}$ for $x_{i,j}$ is positive (greater than zero). Otherwise, it is marked as background. We denote the final predicted masks as $mask_{pred}$, which is a binary matrix with the same shape as the original image.

Evaluation metric. To quantitatively evaluate the effects of various corruptions and other factors on SAM, we employ the Intersection over Union (IoU) metric commonly used in segmentation tasks. Specifically, we calculate the IoU between the predicted masks of a certain clean image, denoted as $mask_{clean}$, and the predicted masks of a certain corrupted image, denoted as $mask_{predict}$. This allows us to assess the changes in masks caused by corruption and other factors. It should be noted that since the SAM lacks class labels and can produce multiple binary masks, we focus on regions with the highest similarity. The mIoU is then obtained by averaging the IoU scores from N pairs of data samples, as shown in Equation 2.

$$mIoU = \frac{1}{N} \sum_{i=1}^{N} IoU(mask_{clean}, mask_{predict}),$$ (2)

where $mask_{(\cdot)}$ is a binary matrix indicating whether a pixel is predicted to be masked, and $mask_{clean}$ and $mask_{predict}$ are the masked region of content images and the images after corruptions or other factors, respectively. The maximum of mIoU equals 1.0 when no corruptions or other factors are applied. Note that we use Equation 2 to evaluate the performance of SAM under style transfer, corruption, occlusion, and perturbation.

3.2 Implementation Details

Style transfer. Synthetic corruption, also interpreted as style transfer, is closely related to texture synthesis and transfer Huang and Belongie [2017], Elad and Milanfar [2017], which achieve style transfer through non-photorealistic rendering techniques Kyprianidis et al. [2012] while preserving the overall global shape of the image Geirhos et al. [2019]. We perform style transfer by taking the style of an image from an arbitrary style image and transferring it onto a content image to synthesize an output image. The resulting output image combines the style characteristics of the former with the content of the latter. Note that we employ AdaIn Huang and Belongie [2017] as the style transfer method to evaluate the robustness of SAM against style transfer images. We perform style transfer on each content image using five different styles, with 45 randomly selected images for each style. This results in a total of 225 style transfer images for each content image.

(a) *cartoon*   (b) *naturalism*   (c) *painting*   (d) *photorealism*   (e) *suprematism*

Figure 2: One example from each style dataset.

Common corruptions. The image encoder of SAM is constructed using a pre-trained Vision Transformer (ViT) Dosovitskiy et al. [2020] model. Here, we employ three different SAM backbones (ViT-B, ViT-L, ViT-H) to evaluate the robustness of SAM under common corruptions. To generate the dataset for common corruption robustness, we follow the methodology described in Hendrycks and Dietterich [2019]. In detail, for robustness under common corruptions, we apply various types of noise corruption, including Gaussian and shot noise, as well as corruptions such as glass and motion blur. We also consider weather conditions like snow and fog, along with digital corruptions such as pixelation and JPEG compression. Each type of corruption is evaluated at 5 different severity levels (the higher the level, the greater the impact on the image), accounting for the expected diversity in both the variety and intensity of corruption found in real-world data. This results in a total of 75 corrupted images for each original image in the dataset.

Local occlusion. Various methods can be considered to define occlusion. Here, we adopt a straightforward approach to evaluate the robustness of SAM under local occlusion. We start with an image, denoted as $\mathbf{x} = \{x_i\}_{i=1}^{N}$, where N represents the total number of patches within the image. We then randomly select M patches (M < N) and replace the pixel values of these patches with Gaussian noise. Subsequently, we evaluate how well SAM maintains mask boundaries when provided with an occluded image at a specific occlusion ratio, $\frac{M}{N}$. This empirical study yields valuable insights into mask quality in occlusion-heavy scenarios, enabling us to evaluate the real-world reliability and utility of SAM. Note that we evaluate the robustness of SAM at 5 different occlusion levels: 10%, 20%, 40%, 60%, and 80%, and we employ three different SAM backbones (ViT-B, ViT-L, ViT-H) for evaluation.

Adversarial perturbation. We deploy PGD Madry et al. [2018] attack and FGSM Goodfellow et al. [2015] attack. Following prior research Benz et al. [2021b] on attacking vision models in the white-box setting, we consider a range of epsilon values, denoted as $\delta = \{d/255 | d \in \{0.1, 0.3, 0.5, 0.8, 1, 3, 5, 8\}\}$, for images normalized to the range of $[0, 1]$. For the PGD attack, if not specified, we adopt the PGD-20 attack, which indicates that the number of iterations is set to 20.

Table 1: Style Transfer mIoU Results.

Style transfer	Cartoon	Naturalism	Painting	Photorealism	Suprematism
Chair	0.5948	0.4319	0.5872	0.5824	0.6330
Bear	0.8867	0.6309	0.7715	0.8743	0.8766
Bottle	0.6764	0.6325	0.8040	0.6713	0.8535
Bicycle	0.7340	0.8051	0.6689	0.8258	0.6839
Car	0.6954	0.7497	0.6993	0.7164	0.7588
Airplane	0.5443	0.4222	0.5258	0.5591	0.6747
Average	0.6886	0.6121	0.6761	0.7049	0.7468

4 Robustness of SAM Against Corruptions

4.1 Style Transfer

It has been found that SAM is more biased towards texture (style) rather than shape Zhang et al. [2023e]. Motivated by their finding, we first evaluate the model's robustness against the style change. In contrast, to Zhang et al. [2023e] that synthesizes an image with texture contrast on different regions, we change the texture of the whole image into a new one. Specifically, we leverage style transfer to change the texture of the image while keeping the shape of the image content. Style transfer is accomplished by combining the content of one image with the style of another, resulting in a new image that preserves the underlying structure of the content image while incorporating the visual characteristics and artistic style of the style image. For the content images, we utilize the dataset provided in Geirhos et al. [2019],

which serves as the source of content images for our style transfer experiments. For the style images, we follow the work in Westlake et al. [2016] and select five different ones to use in our experiments.

Specifically, Figure 2 showcases an example from the aforementioned set of 5 different style images used in our experiments. In addition, Figure 3 displays the content images, along with their corresponding mask results. These masks can be considered as ground truth, serving as a basis for calculating the mIoU scores. Consequently, Figure 4 illustrates the results of synthesized images generated from the same content image using different style transfers. These images demonstrate the application of various styles to the content image. It should be noted that the objects in these synthesized images in Figure 4 are distinguishable and clear in human vision. Since the performance of SAM on such images is currently unclear, we show the mask result in Figure 5 (additional results regarding style transfer and their corresponding mask results can be found in Appendix A.1). The masks highlight the regions affected by the style transfer process, providing a visual representation of the changes to the performance of SAM. An overall illustration can be observed that even visually perceivable objects exhibit varying segmentation performance under different style transformations. Further, we investigate the mIoU under different style transformations, and the results are shown in Table 1. On average, SAM achieves a moderate mIoU performance across all object categories and style transfers. The highest mIoU scores are observed in the Photorealism and Suprematism styles, while the lowest scores are seen in the Naturalism style.

These findings emphasize the impact of style transfer on SAM's segmentation performance, with certain styles yielding better results than others. Although the average mIoU indicates a reasonable overall performance, there is room for improvement in specific style transfer cases and adaptability to diverse visual styles.

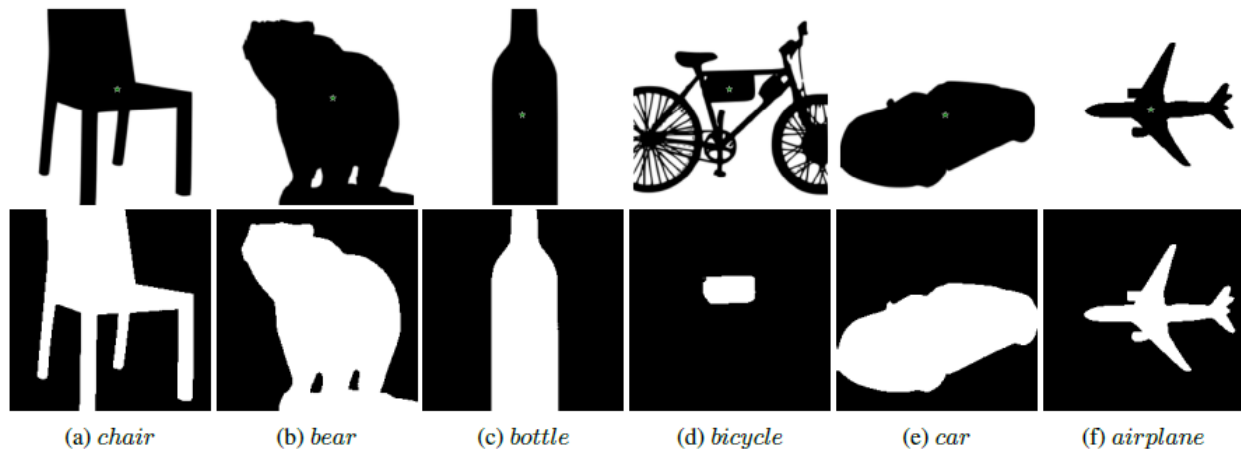| (a) chair | (b) bear | (c) bottle | (d) bicycle | (e) car | (f) airplane |

Figure 3: The content images used for style transfer. Figures (a) to (f) refer to the content image with the location of the point prompt marked in a green star and their corresponding masked results. The masked results without style transfer are used as the ground truth in the style transfer evaluation.

4.2 Common Corruptions

Various corruptions can be perceived as different new styles Benz et al. [2021a]. Style change caused by the technique of style transfer in the above can be seen as a synthetic corruption that rarely occurs in practice. Therefore, we further evaluate the generalization capability of SAM under common corruptions Hendrycks and Dietterich [2019].

To generate the dataset for assessing common corruption robustness, we start by randomly selecting 100 images from the SA-1B dataset introduced in the SAM paper Kirillov et al. [2023]. We evaluate the model's robustness under 15 different corruptions, each with 5 severity levels. These corruptions are broadly categorized into 4 categories: digital, noise, weather, and blur, as described in Hendrycks and Dietterich [2019]. In detail, the digital category of common corruptions includes adjustments to brightness, contrast, pixelation, and jpeg compression. Figure 6 provides visual examples of both the original and corrupted images, demonstrating the impact of the corruptions at severity level 5. Further, Figure 7 showcases the resulting mask generated by the ViT-H model for each corresponding image. These images are then subjected to the corruptions explained above. After perturbing the images, we feed them into the model to generate segmentation masks. To report the result for a batch of a certain image, we take an average of IoUs for all batches. Finally, the mIoU scores across all the selected images are reported in Table 2 and Table 3, respectively.

We find that the output of SAM on both the original images and the images corrupted with commonly occurring noise, blur, and natural disruptions exhibit a remarkably high level of similarity. Notably, we observe an IoU of approximately
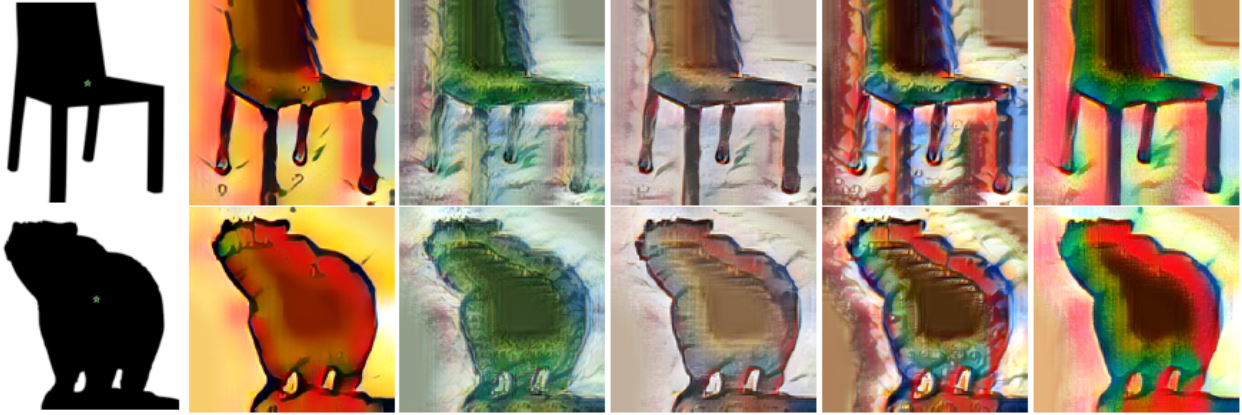
Figure 4: The synthetic images after style transfer. Figure (a) refers to the original content image with the location of the point prompt marked in a green star. Figures (b) to (f) refer to the synthetic images of the content image with styles such as cartoon, naturalism, painting, photorealism, and suprematism, respectively.
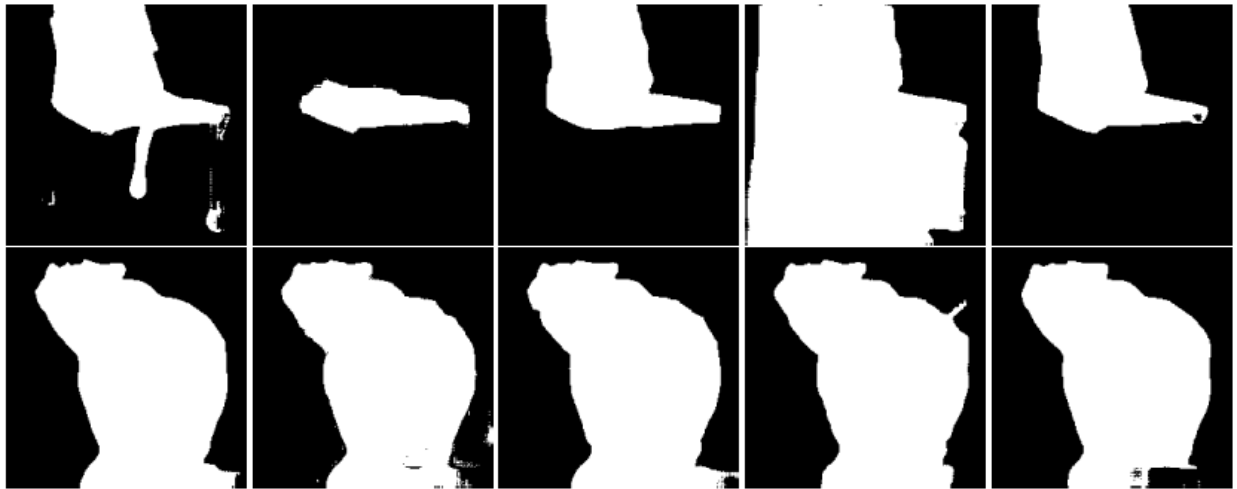

Figure 5: The mask of images after style transfer.

0.8 at the higher severity level (level 5) for the majority of added corruptions except for zoom blur. These preliminary results suggest that SAM possesses robust capabilities and is likely to perform well on real-world data that commonly exhibit such corruptions. Consequently, these findings indicate the potential to reduce the efforts required for curating datasets significantly, thanks to the SAM's ability to handle common corruptions effectively. Moreover, it is generally observed that as the severity levels of corruption increase, the robustness of the model tends to decrease, which is an expected behavior. However, it is observed that the model performs poorly, even at the lowest severity level (level 1), for certain specific corruptions such as zoom blur. This observation may be attributed to the significant distortion introduced by this corruption, which may be more severe compared to other types of corruptions. Obviously, significant disruption becomes readily apparent at the maximum severity level for the zoom blur corruption, as depicted in Figure 6 (g). Our results are consistent with the findings in concurrent works Huang et al. [2023], Wang et al. [2023] that have reported the robustness of SAM in a similar setup.

5 Beyond Corruptions

5.1 Local Occlusion

The robustness of models against occlusions is a pivotal and intricate aspect of computer vision research, with significant consequences in practical applications Naseer et al. [2021], Guo et al. [2023], Gunasekaran and Jaiman [2023], Kunhardt

Table 2: Common Corruption mIoU Results - Severity level 5.

Model	Gaussian Noise	Shot Noise	Impulse Noise	Jpeg Compression	Snow	Frost	Fog
ViT-H	0.7311	0.7442	0.7351	0.6789	0.6950	0.7544	0.8334
ViT-L	0.7239	0.7341	0.7256	0.6527	0.6923	0.7500	0.8330
ViT-B	0.6709	0.6856	0.6756	0.5590	0.6430	0.7068	0.8070
Model	Defocus blur	Motion Blur	Zoom blur	Glass blur	Brightness	Contrast	Pixelate
ViT-H	0.7454	0.7073	0.3739	0.6900	0.8171	0.6757	0.8835
ViT-L	0.7447	0.7077	0.3698	0.6902	0.8177	0.6779	0.8623
ViT-B	0.7166	0.6731	0.3556	0.6566	0.7843	0.6351	0.8285

Table 3: Common Corruption mIoU Results - ViT-H, All Severity Levels.

Level	Gausian Noise	Shot Noise	Impulse Noise	Glass blur	Zoom blur	Frost	Fog
1	0.8871	0.8902	0.8693	0.8741	0.5097	0.8536	0.8890
2	0.8577	0.8589	0.8386	0.8414	0.4562	0.8146	0.8710
3	0.8207	0.8268	0.8187	0.7600	0.4227	0.7845	0.8585
4	0.7824	0.7793	0.7752	0.7389	0.3910	0.7841	0.8539
5	0.7311	0.7442	0.7351	0.6900	0.3739	0.7544	0.8334
Level	Defocus blur	Motion Blur	Snow	Jpeg Compression	Brightness	Contrast	Pixelate
1	0.8966	0.8899	0.8334	0.8855	0.9430	0.8881	0.9469
2	0.8658	0.8442	0.7753	0.8528	0.9054	0.8613	0.9460
3	0.8185	0.7905	0.7521	0.8269	0.8760	0.8244	0.9150
4	0.7806	0.7382	0.6982	0.7588	0.8450	0.7544	0.9024
5	0.7454	0.7073	0.6950	0.6789	0.8171	0.6757	0.8835

(a) original (b) gaussian noise (c) shot noise (d) impulse noise (e) defocus blur

(f) motion blur (g) zoom blur (h) glass blur (i) snow (j) frost

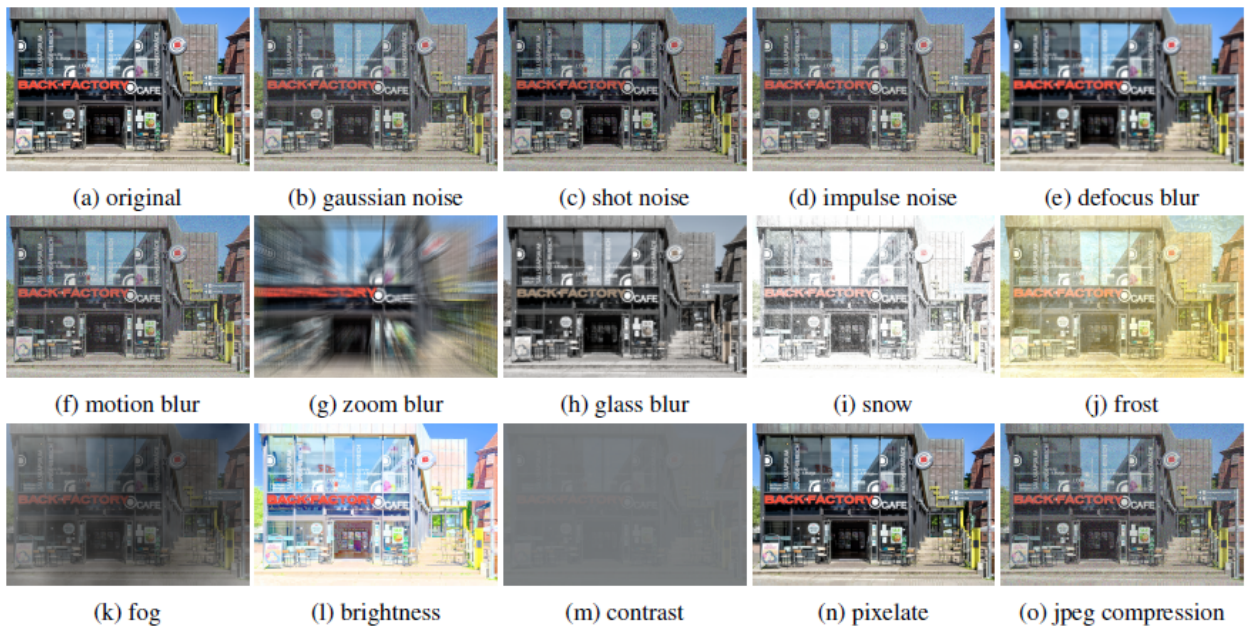(k) fog (l) brightness (m) contrast (n) pixelate (o) jpeg compression

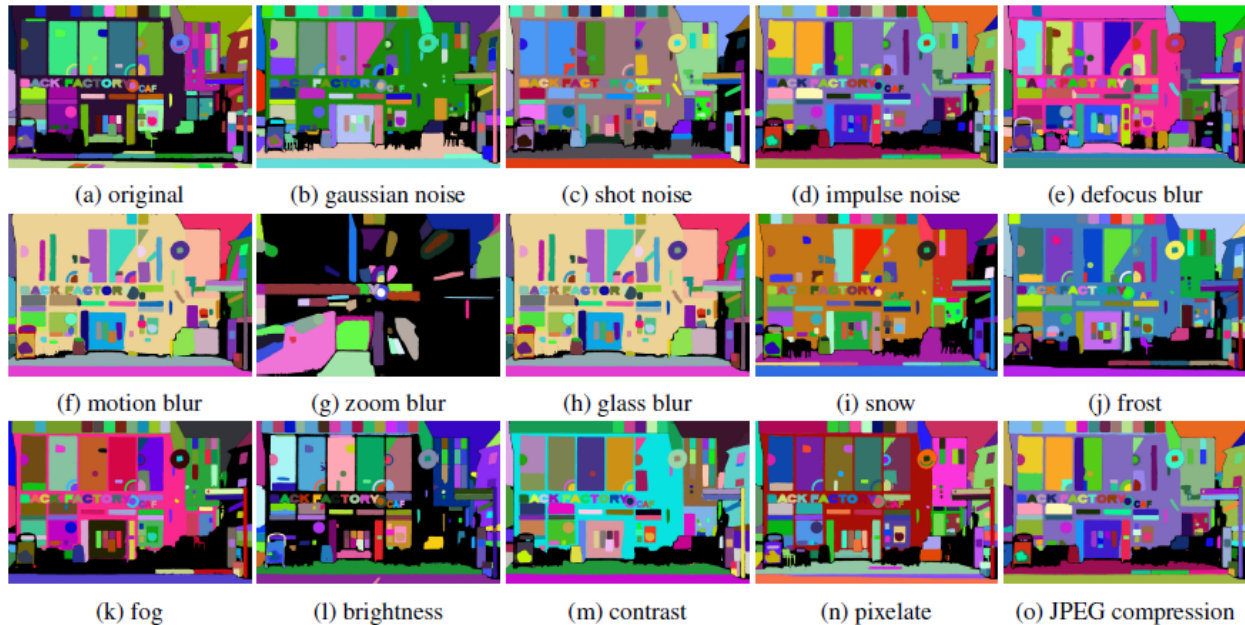Figure 6: Original and corrupted images with severity level 5.

Figure 7: The masks of SAM for original and corrupted images with severity level 5.

et al. [2021], Athalye et al. [2018]. Considering the frequent occurrence of occlusions in natural images, it is essential to assess the performance of models under obscured conditions Zeng et al. [2021], Ke et al. [2021a,b]. In this regard, we aim to investigate the robustness of the SAM model in the presence of occlusions.

We use the same dataset for evaluation as introduced in common corruption Section 4.2. In particular, for occlusion robustness, we apply occlusion perturbations to these images, as described in Section 3.2, to generate the occluded dataset for further analysis and evaluation. Figure 8 showcases two example images and their corresponding mask results. The results for local occlusion robustness evaluation are presented in Table 4.

An overall trend can be observed that as the occlusion ratio increases, the mIoU values for all models gradually decrease. This indicates that the model's performance in terms of segmentation accuracy is negatively impacted by the presence of occlusion. Among the models, ViT-H achieves the highest mIoU values across all occlusion ratios in most cases, while ViT-B exhibits the lowest mIoU values. However, it should be noted that when the occlusion ratio is around 40%, the models can achieve competitive performance. Overall, these results highlight the significance of occlusion in challenging the robustness of the models. Higher occlusion ratios lead to a greater degradation in segmentation performance, emphasizing the need for improved robustness in computer vision models when dealing with occluded scenes.

Table 4: Occlusion mIoU results.

Model	0.1	0.2	0.4	0.6	0.8
ViT-H	0.8934	0.9234	0.6075	0.5655	0.2290
ViT-L	0.9151	0.8950	0.5434	0.6012	0.1039
ViT-B	0.8775	0.8996	0.5229	0.4696	0.2266

5.2 Adversarial Perturbation

It is widely known that deep recognition models can be fooled by adversarial attacks to change the predicted label with subtle perturbations Goodfellow et al. [2015], Zhang et al. [2020]. To assess model robustness, two commonly used methods are Fast Gradient Sign Method (FGSM) Goodfellow et al. [2015] and Projected Gradient Descent (PGD) Madry et al. [2018]. Note that we also use the same dataset as introduced in common corruption Section 4.2 to investigate the robustness under adversarial perturbations. In detail, the FGSM attack is a single-step method that perturbs the input image based on the gradient of the model. By calculating the sign of the gradient with respect to the input, the
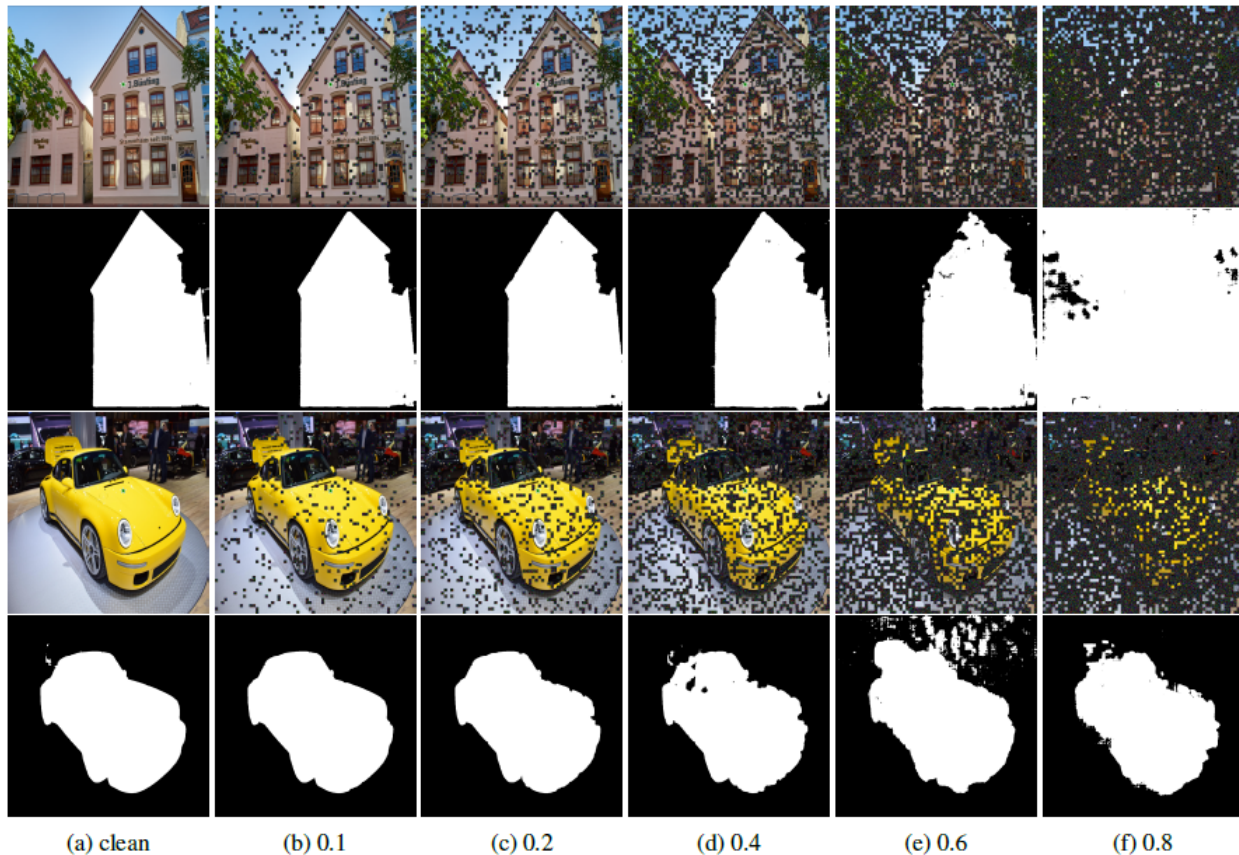
Figure 8: The example images under different levels of occlusion ratio. Figure (a) refers to the clean image without occlusion and the location of the point prompt marked in a green star. Figures (b) to (e) refer to the random occluded images with varying occlusion ratios. The pixel values of the occluded region are set to Gaussian noise with a mean of 0 and variance of 1.

FGSM attack introduces perturbations to the image. PGD can be seen as an iterative extension of the FGSM attack and is sometimes referred to as I-FGSM Zhang et al. [2023h]. It performs multiple iterations of the FGSM attack, progressively refining the perturbations. At each iteration, the perturbation is updated based on the model's gradient with respect to the perturbed image while also constraining the perturbation within a specified range. Note that in this work, we consistently refer to this method as PGD for the sake of consistency. When specifying the number of iterations, the PGD attack is denoted as PGD-N, where N represents the iteration number. Here, we define $\mathcal{F}((\boldsymbol{x}, y); \omega)$ as the target model to be attacked where y represents the corresponding labels for the input image $\boldsymbol{x}$.

$$\delta = \mathbb{E}_{(\boldsymbol{x};y) \in \mathcal{D}}[max\, \mathcal{L}(\mathcal{F}(\omega; \widetilde{\boldsymbol{x}}); y))], \tag{3}$$

where $\mathcal{D}$ is the dataset consisting of the image pair $(x; y)$, $\mathcal{L}(\cdot)$ is the loss function (such as cross-entropy loss), and $\widetilde{\boldsymbol{x}}$ defined as $x + \delta$ is the AE of x, which is generated by PGD as follows:

$$\boldsymbol{x}^{t+1} = \Pi_{\boldsymbol{x}+\delta}\left(\boldsymbol{x}^t + \alpha\,\mathrm{sign}(\nabla_{\boldsymbol{x}}\mathcal{L}(\mathcal{F}((\omega; \boldsymbol{x}); y)))\right), \tag{4}$$

where α represents the step size, $\boldsymbol{x}^t$ denotes the AE generated at step t, $\Pi_{\boldsymbol{x}+\delta}$ represents the projection function that projects the AE onto the ϵ-ball centered at $\boldsymbol{x}^0$, and $\mathrm{sign}(\cdot)$ denotes the sign function. Additionally, to ensure that the perturbation δ is imperceptible (or quasi-imperceptible) to the human eye, it is commonly constrained by an upper bound ϵ on the ℓ_∞-norm, i.e., $||\delta||_\infty \le \epsilon$.

For the aforementioned adversarial perturbation tasks, their objective is to manipulate the images through subtle perturbations at either the image or pixel level, leading to incorrect label predictions by the model. However, attacking the model as in common tasks becomes nearly impossible due to the lack of semantic labels for the masks generated

by SAM. We follow this work Zhang et al. [2023h] and adopt a direct approach to cause the model to be unable to detect objects, thereby removing the generated masks under adversarial perturbations. Specifically, if the predicted value $mask_{i,j}$ for pixel $x_{i,j}$ is positive, it signifies that the pixel is masked. Therefore, the objective of successfully removing the mask is achieved when all the predicted values $mask_{i,j}$ within the designated region become negative. Here, we redefine the loss function, and the Equation 3 are as reformulated as follows:

$$\delta = \mathbb{E}_{||\delta||_\infty \le \epsilon}[max \, ||clamp(SAM(x + \delta, \boldsymbol{prompt}; \omega), min = \theta) - \theta||^2], \tag{5}$$

where θ is a negative threshold, $clamp(\cdot)$ is the clamp operation to restrict the predicted values' range. Setting the clamp operation in this way is because it makes no sense to increase the prediction values that are already below θ to be close to θ. In other words, we can evaluate the robustness of SAM by performing these adversarial attacks.

Adversarial robustness evaluation. To evaluate the robustness of SAM against FGSM and PGD attacks, we evaluate the robustness of SAM against FGSM and PGD attacks. The visualization of the results includes clean images, adversarial images, and predicted masks, as shown in Figure 9. Additional results can be found in Appendix A.2. Note that for these two attacks, we set the θ to be -10, which is also used in their work Zhang et al. [2023h]. Specifically, the clean image is given a prompt indicated by a green star in Figure 9(a). The generated adversarial images through FGSM and PGD attacks, as depicted in Figure 9(b) and Figure 9(c), respectively. Table 5 presents the mIoU results for different perturbation magnitudes (δ) used in the attacks. For the FGSM attack, we observe that as the perturbation magnitude increased from 0.1/255 to 8.0/255, the mIoU decreased from 0.8227 to 0.4406. This indicates a decrease in the model's performance as it becomes more susceptible to adversarial perturbations. In the case of the PGD-20 attack, the mIoU values are significantly lower compared to FGSM. Starting from 0.8254 at $\delta = 0.1/255$, the mIoU decreased sharply to 0.0019 at $\delta = 8.0/255$. This demonstrates a high vulnerability of the SAM model to more sophisticated attacks, where the model's segmentation performance deteriorates rapidly.

Overall, these results in both Figure 9 and Table 5 highlight the limited robustness of SAM against both FGSM and PGD attacks, with the PGD-20 attack showing a higher impact than FGSM on the model's robustness.

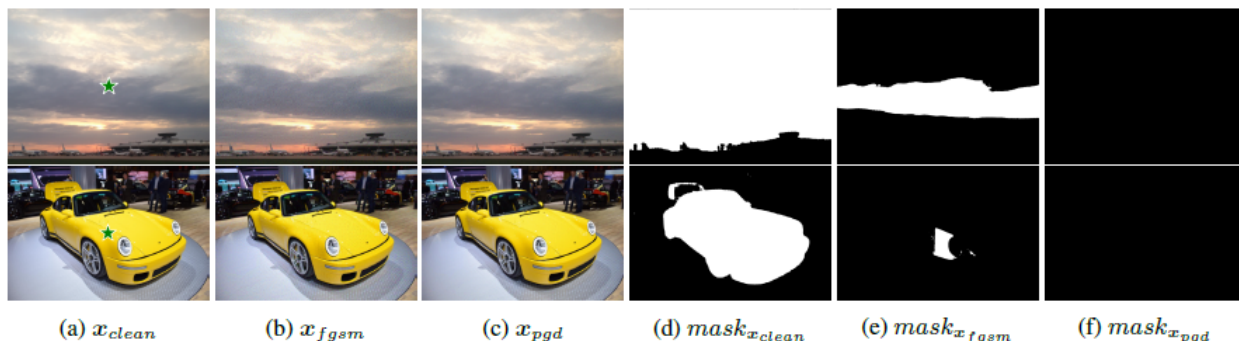| (a) x_{clean} | (b) x_{fgsm} | (c) x_{pgd} | (d) $mask_{x_{clean}}$ | (e) $mask_{x_{fgsm}}$ | (f) $mask_{x_{pgd}}$ |

Figure 9: Robustness evaluation against FGSM and PGD-20 attacks with $\delta = 8.0/255$. Figure (a) shows the original image with a green star indicating the location of the point prompt. Figures (b) and (c) display the adversarial images generated using FGSM and PGD attacks, respectively. The white areas in Figures (d), (e), and (f) represent masks predicted by SAM based on the provided prompt and the images from Figures (a), (b), and (c), respectively. The results in Figure (e) and (f) demonstrate that SAM is susceptible to adversarial attacks, as indicated by the reduced white area compared to Figure (d).

Table 5: Adversarial attack mIoU results.

δ	0.1/255	0.3/255	0.5/255	0.8/255	1.0/255	3.0/255	5.0/255	8.0/255
FGSM	0.8227	0.6091	0.5445	0.4968	0.4808	0.4513	0.4386	0.4406
PGD-20	0.8254	0.5745	0.3875	0.2190	0.1723	0.0068	0.0012	0.0019

6 Conclusions

In this work, we are among the early pioneers to evaluate the robustness of the segment anything model (SAM), for which we provide a comprehensive evaluation. Specifically, we first explore the robustness of SAM to synthetic

corruptions using the style transfer method. We find that in the majority of cases, SAM is affected by style transfer, even when the corruption is visually distinguishable to the human eye. Second, we investigate the robustness of SAM to real-world corruption, such as common corruption. Our findings indicate that SAM is robust to 14 common corruptions except for zoom blur. Third, we evaluate the robustness of SAM under corruption; we find that even when the occlusion ratio reaches 40%, SAM still achieves competitive performance in this type of corruption. These findings suggest that SAM is not necessarily capable of segmenting anything while also motivating the community to optimize SAM for specific scenarios further. Finally, we further investigate the adversarial robustness of SAM and find that SAM exhibits limited robustness against FGSM and PGD attacks.

References

Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 2020.

Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.

Chaoning Zhang, Chenshuang Zhang, Sheng Zheng, Yu Qiao, Chenghao Li, Mengchun Zhang, Sumit Kumar Dam, Chu Myaet Thwal, Ye Lin Tun, Le Luang Huy, et al. A complete survey on generative ai (aigc): Is chatgpt from gpt-4 to gpt-5 all you need? *arXiv preprint arXiv:2303.11717*, 2023a.

Chaoning Zhang, Chenshuang Zhang, Chenghao Li, Yu Qiao, Sheng Zheng, Sumit Kumar Dam, Mengchun Zhang, Jung Uk Kim, Seong Tae Kim, Jinwoo Choi, et al. One small step for generative ai, one giant leap for agi: A complete survey on chatgpt in aigc era. *arXiv preprint arXiv:2304.06488*, 2023b.

Chenshuang Zhang, Chaoning Zhang, Mengchun Zhang, and In So Kweon. Text-to-image diffusion models in generative ai: A survey. *arXiv preprint arXiv:2303.07909*, 2023c.

Chenshuang Zhang, Chaoning Zhang, Sheng Zheng, Mengchun Zhang, Maryam Qamar, Sung-Ho Bae, and In So Kweon. A survey on audio diffusion models: Text to speech synthesis and enhancement in generative ai. *arXiv preprint arXiv:2303.13336*, 2023d.

Chenghao Li, Chaoning Zhang, Atish Waghwase, Lik-Hang Lee, Francois Rameau, Yang Yang, Sung-Ho Bae, and Choong Seon Hong. Generative ai meets 3d: A survey on text-to-3d in aigc era. *arXiv preprint arXiv:2305.06131*, 2023.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021.

Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.

Chaoning Zhang, Chenshuang Zhang, Junha Song, John Seon Keun Yi, Kang Zhang, and In So Kweon. A survey on masked autoencoder for self-supervised learning in vision and beyond. *arXiv preprint arXiv:2208.00173*, 2022.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.

Chaoning Zhang, Yu Qiao, Shehbaz Tariq, Sheng Zheng, Chenshuang Zhang, Chenghao Li, Hyundong Shin, and Choong Seon Hong. Understanding segment anything model: Sam is biased towards texture rather than shape. 2023e.

Philipp Benz, Chaoning Zhang, Adil Karjauv, and In So Kweon. Revisiting batch normalization for improving corruption robustness. *WACV*, 2021a.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *ICLR*, 2019.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.

Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023.

Yizhe Zhang, Tao Zhou, Peixian Liang, and Danny Z Chen. Input augmentation with sam: Boosting medical image segmentation with segmentation foundation model. *arXiv preprint arXiv:2304.11332*, 2023f.

Lv Tang, Haoke Xiao, and Bo Li. Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709*, 2023.

Dongsheng Han, Chaoning Zhang, Yu Qiao, Maryam Qamar, Yuna Jung, SeungKyu Lee, Sung-Ho Bae, and Choong Seon Hong. Segment anything model (sam) meets glass: Mirror and transparent objects cannot be easily detected. *arXiv preprint*, 2023.

IDEA-Research. Grounded segment anything, 2023. URL `https://github.com/IDEA-Research/Grounded-Segment-Anything`. GitHub repository.

Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023.

Jiaqi Chen, Zeyu Yang, and Li Zhang. Semantic-segment-anything, 2023. URL `https://github.com/fudan-zvg/Semantic-Segment-Anything`. GitHub repository.

Curt Park. segment anything with clip, 2023. URL `https://github.com/Curt-Park/segment-anything-with-clip`. GitHub repository.

Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pages 12888–12900. PMLR, 2022.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.

Tao Yu, Runseng Feng, Ruoyu Feng, Jinming Liu, Xin Jin, Wenjun Zeng, and Zhibo Chen. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*, 2023.

Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*, 2023.

Zxyang. Segment and track anything, 2023. URL `https://github.com/z-x-yang/Segment-and-Track-Anything`. GitHub repository.

Qiuhong Shen, Xingyi Yang, and Xinchao Wang. Anything-3d: Towards single-view anything reconstruction in the wild. *arXiv preprint arXiv:2304.10261*, 2023.

Minki Kang, Dongchan Min, and Sung Ju Hwang. Any-speaker adaptive text-to-speech synthesis with diffusion models. *arXiv preprint arXiv:2211.09383*, 2022.

Shehbaz Tariq, Brian Estadimas Arfeto, Chaoning Zhang, and Hyundong Shin. Segment anything meets semantic communication. *arXiv preprint arXiv:2306.02094*, 2023.

Chaoning Zhang, Sheng Zheng, Chenghao Li, Yu Qiao, Taegoo Kang, Xinru Shan, Chenshuang Zhang, Caiyan Qin, Francois Rameau, Sung-Ho Bae, et al. A survey on segment anything model (sam): Vision foundation model meets prompt engineering. 2023g.

Chenshuang Zhang, Chaoning Zhang, Taegoo Kang, Donghun Kim, Sung-Ho Bae, and In So Kweon. Attack-sam: Towards evaluating adversarial robustness of segment anything model. *arXiv preprint*, 2023h.

Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017.

Michael Elad and Peyman Milanfar. Style transfer via texture synthesis. *IEEE Transactions on Image Processing*, 26 (5):2338–2351, 2017.

J Kyprianidis, J Collomosse, T Wang, and T Isenberg. A taxonomy of artistic stylization techniques for images and video. *State of the'art*, 2012.

Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Philipp Benz, Soomin Ham, Chaoning Zhang, Adil Karjauv, and In So Kweon. Adversarial robustness comparison of vision transformer and mlp-mixer to cnns. *arXiv preprint arXiv:2110.02797*, 2021b.

Nicholas Westlake, Hongping Cai, and Peter Hall. Detecting people in artwork with cnns. In *European Conference on Computer Vision*, pages 825–841. Springer, 2016.

Yihao Huang, Yue Cao, Tianlin Li, Felix Juefei-Xu, Di Lin, Ivor W Tsang, Yang Liu, and Qing Guo. On the robustness of segment anything. *arXiv preprint arXiv:2305.16220*, 2023.

Yuqing Wang, Yun Zhao, and Linda Petzold. An empirical study on the robustness of the segment anything model (sam). *arXiv preprint arXiv:2305.06422*, 2023.

Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *arXiv preprint arXiv:2105.10497*, 2021.

Xingwu Guo, Ziwei Zhou, Yueling Zhang, Guy Katz, and Min Zhang. Occrob: Efficient smt-based occlusion robustness verification of deep neural networks. In *Tools and Algorithms for the Construction and Analysis of Systems: 29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Paris, France, April 22–27, 2023, Proceedings, Part I*, pages 208–226. Springer, 2023.

Karthick Prasad Gunasekaran and Nikita Jaiman. Now you see me: Robust approach to partial occlusions. *arXiv preprint arXiv:2304.11779*, 2023.

Owen Kunhardt, Arturo Deza, and Tomaso Poggio. The effects of image distribution and task on adversarial robustness. *arXiv preprint arXiv:2102.10534*, 2021.

Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *ICML*, 2018.

Dan Zeng, Raymond Veldhuis, Luuk Spreeuwers, and Richard Arendsen. Occlusion-invariant face recognition using simultaneous segmentation. *IET biometrics*, 10(6):679–691, 2021.

Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Deep occlusion-aware instance segmentation with overlapping bilayers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4019–4028, 2021a.

Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Occlusion-aware video object inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14468–14478, 2021b.

Chaoning Zhang, Philipp Benz, Tooba Imtiaz, and In-So Kweon. Understanding adversarial examples from the mutual influence of images and perturbations. In *CVPR*, 2020.

A appendix

A.1 Style transfer

In this section, we will show more style transfer results in Figure 10 and their corresponding masks in Figure 11. A similar observation highlights that SAM is greatly affected by image style and still has some way to go before reaching the level of human vision.
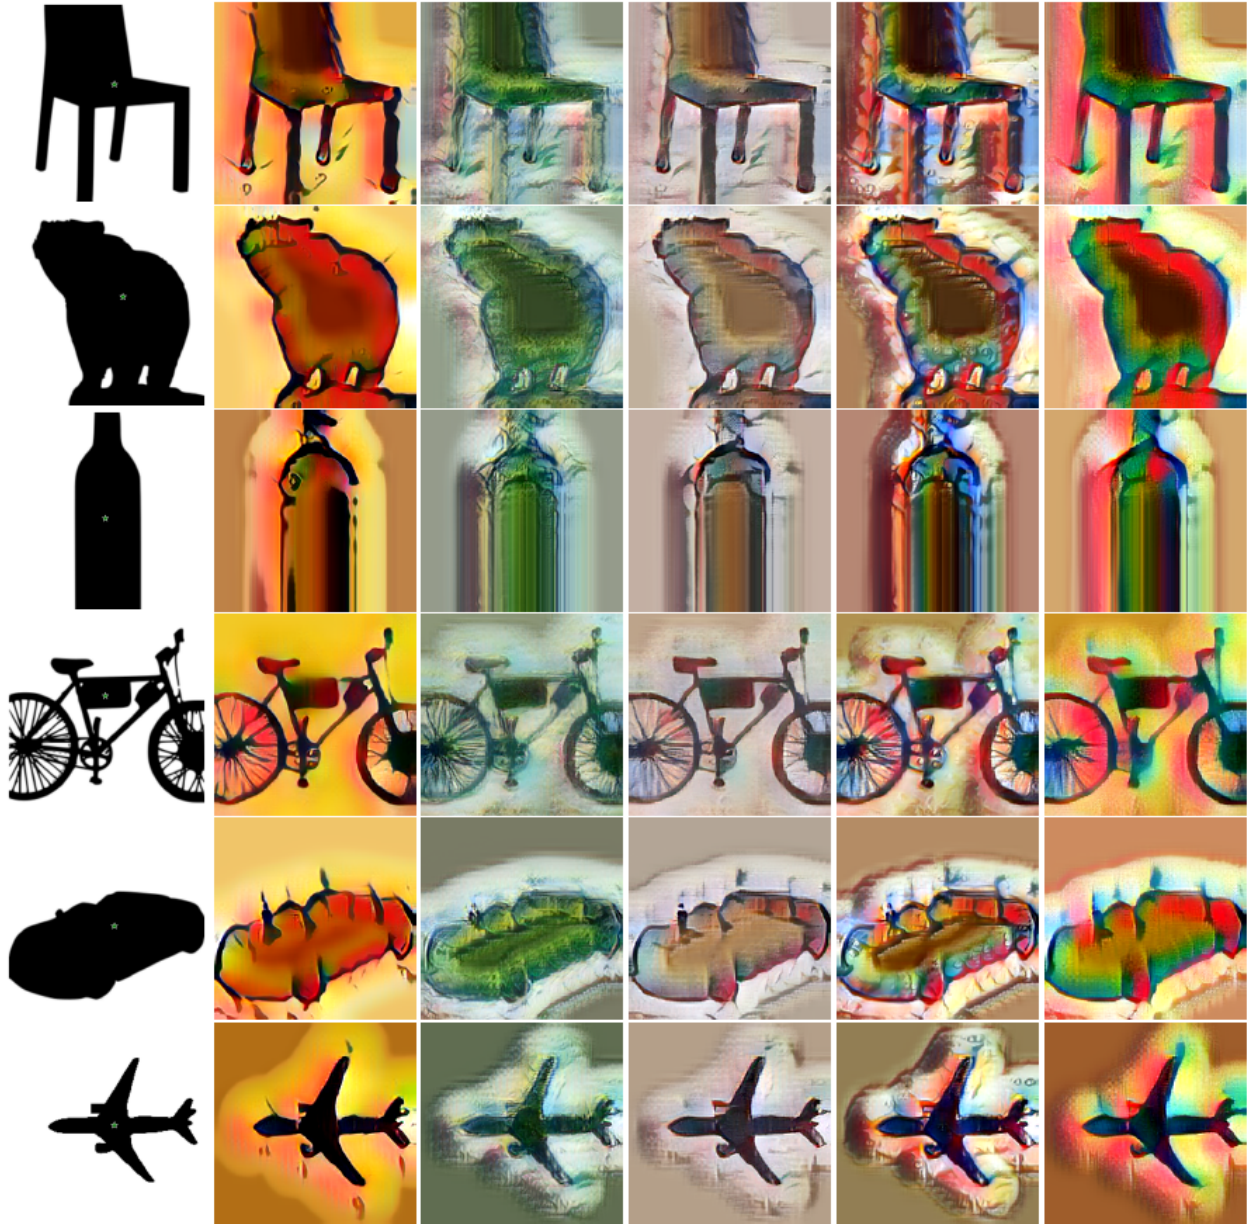

Figure 10: The synthetic images after style transfer. Figure (a) refers to the original content image with the location of the point prompt marked in a green star. Figures (b) to (f) refer to the synthetic images of the content image with styles such as cartoon, naturalism, painting, photorealism, and suprematism, respectively.

A.2 Adversarial attack

In this section, we illustrate more results after FGSM and PGD attacks as shown in Figure 12. Similar results verify the vulnerability of SAM to both FGSM and PGD attacks.
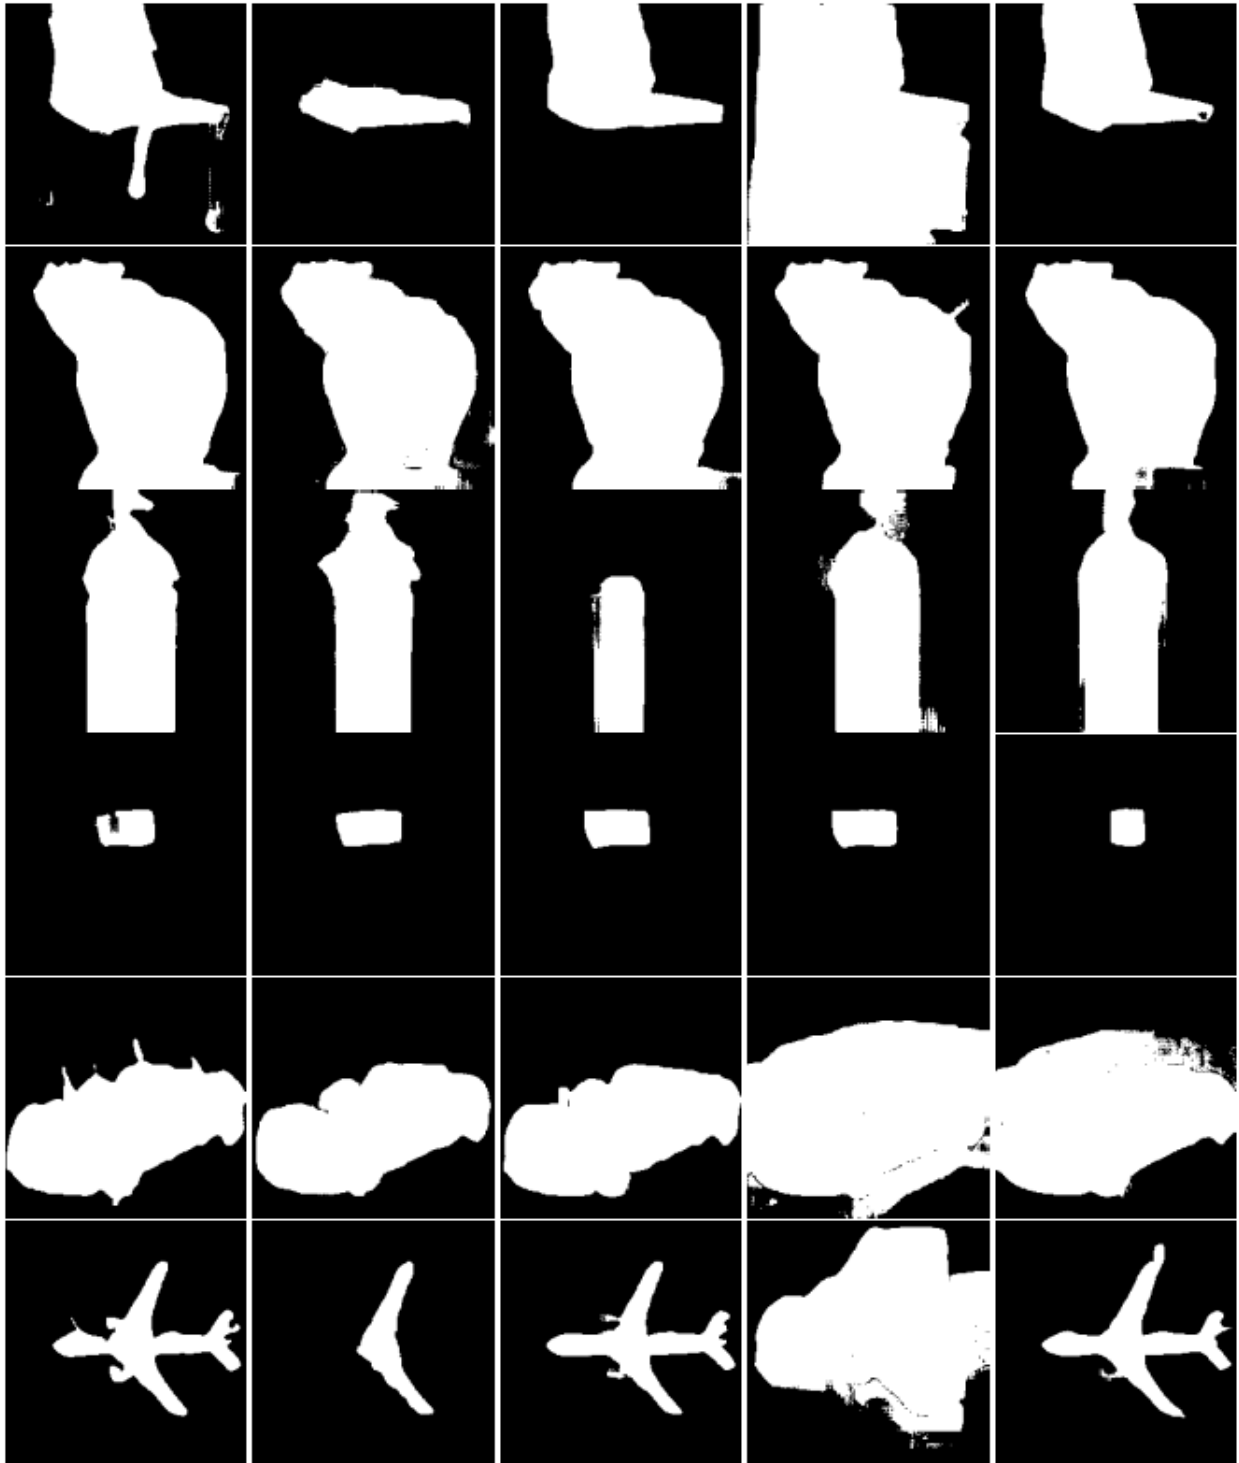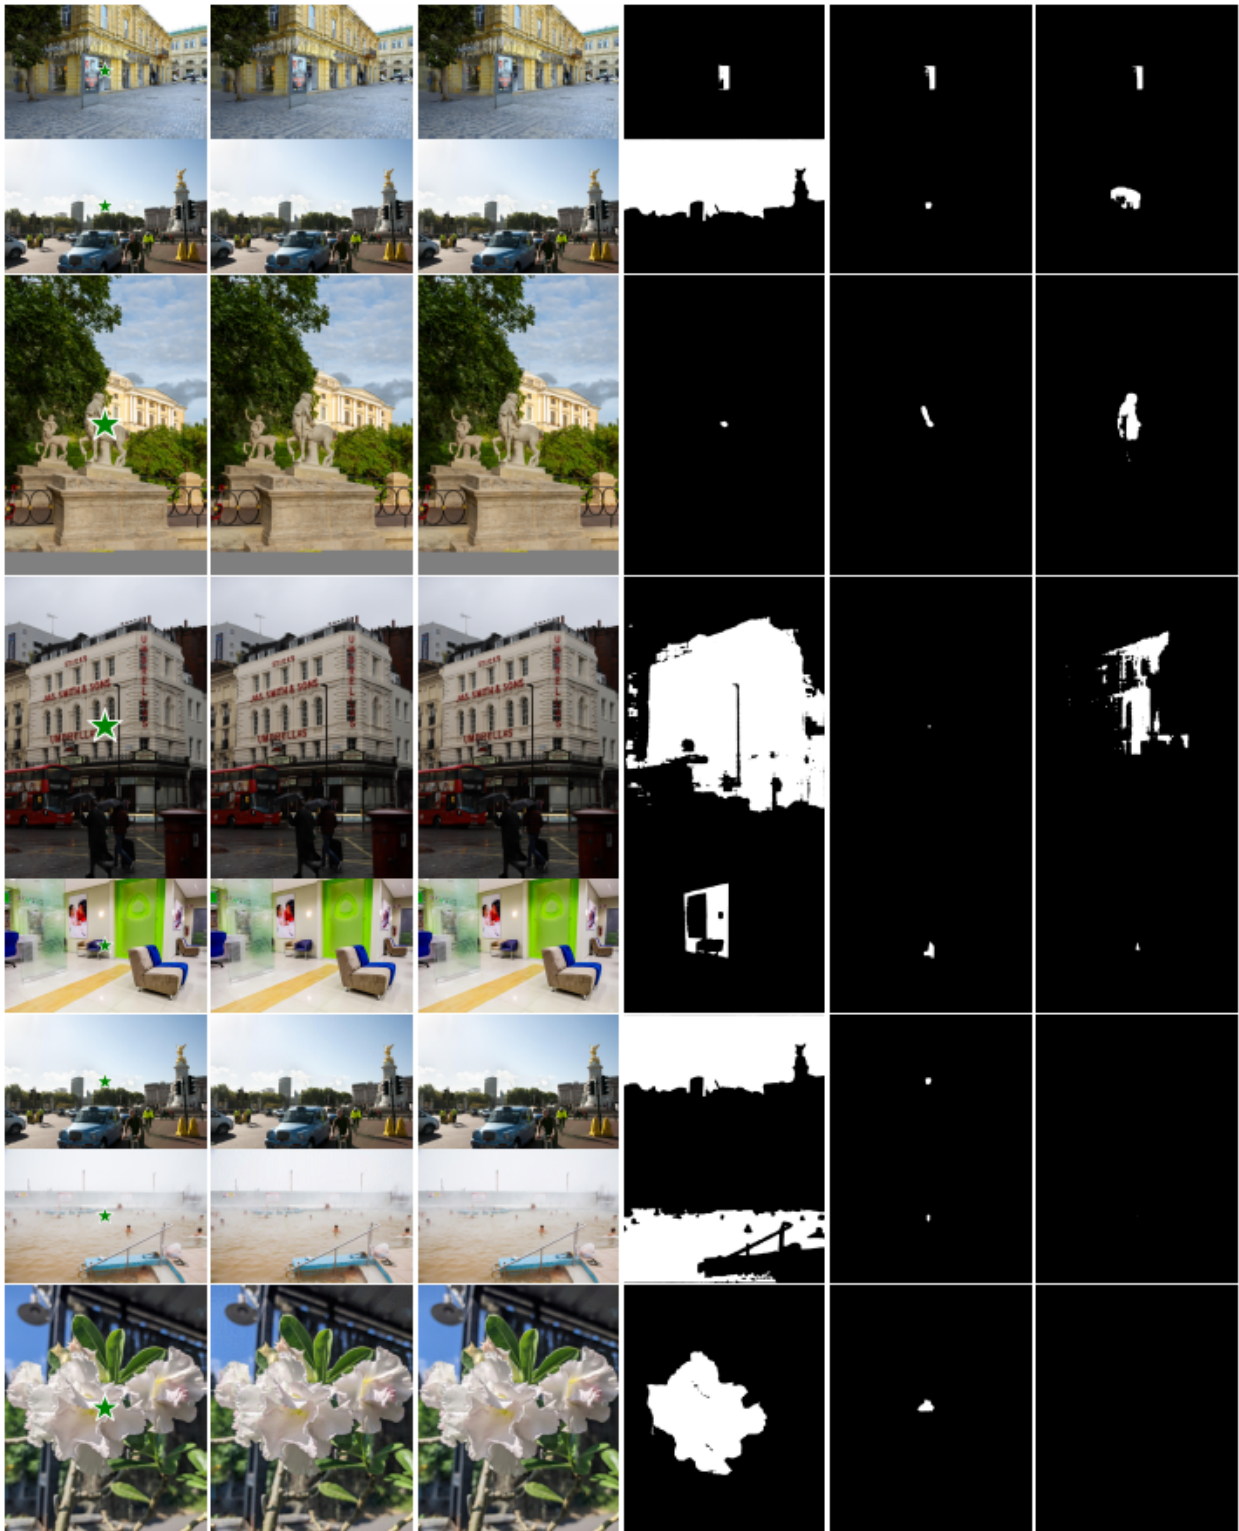
Figure 11: The mask of images after style transfer.

Figure 12: Robustness evaluation against FGSM and PGD-20 attacks with different values of δ. The δ values corresponding to the images from top to bottom are as follows: 0.1, 0.3, 0.5, 0.8, 1.0, 3.0, 5.0, and 8.0. The images from left to right are as follows: clean images, images after applying FGSM, images after applying PGD, masks of clean images, masks of images after applying FGSM, masks of images after applying PGD.