

Benchmarking Segmentation Models with Mask-Preserved Attribute Editing

Zijin Yin¹ Kongming Liang^{1*} Bing Li² Zhanyu Ma¹ Jun Guo¹

¹ Beijing University of Posts and Telecommunications

² King Abdullah University of Science and Technology

¹{yinzijin2017, liangkongming, mazhanyu, guojun}@bupt.edu.cn ²bing.li@kaust.edu.sa

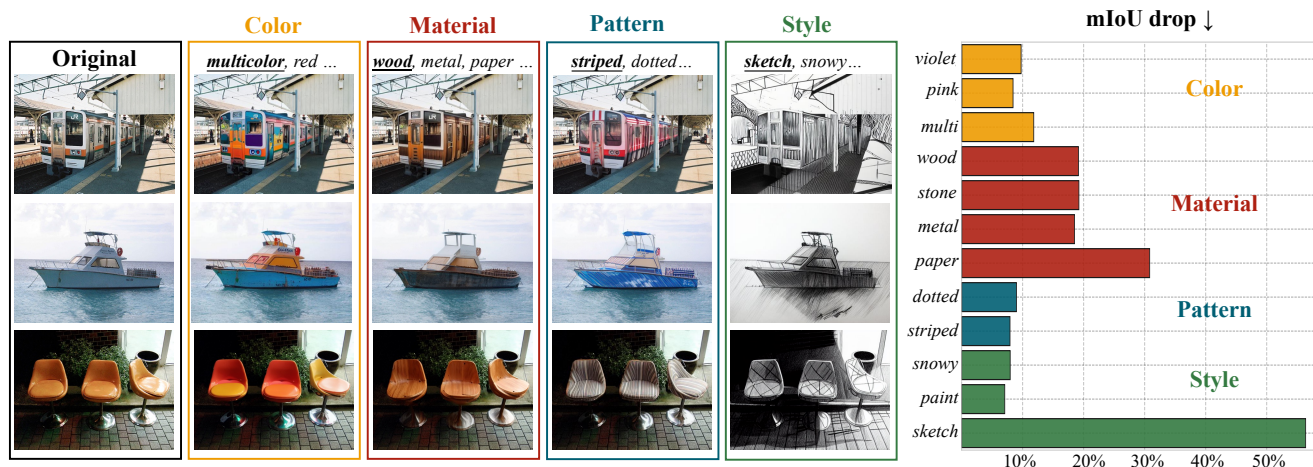


Figure 1. The illustration of the motivation of our work. **Left:** Our mask-preserved attribute editing pipeline generates testing images with various attribute changes for evaluating the robustness of segmentation methods to attribute variations. **Right:** The average performance drop of segmentation models on our generated data, shows the sensitivity to different types of attribute variations.

Abstract

When deploying segmentation models in practice, it is critical to evaluate their behaviors in varied and complex scenes. Different from the previous evaluation paradigms only in consideration of global attribute variations (e.g. adverse weather), we investigate both local and global attribute variations for robustness evaluation. To achieve this, we construct a mask-preserved attribute editing pipeline to edit visual attributes of real images with precise control of structural information. Therefore, the original segmentation labels can be reused for the edited images. Using our pipeline, we construct a benchmark covering both object and image attributes (e.g. color, material, pattern, style). We evaluate a broad variety of semantic segmentation models, spanning from conventional close-set models to recent open-vocabulary large models on their robustness to different types of variations. We find that both local and global attribute variations affect segmentation performances, and the sensitivity of models diverges across different variation types. We argue that local attributes have the same importance as global attributes, and should be considered in

the robustness evaluation of segmentation models. Code: <https://github.com/PRIS-CV/Pascal-EA>.

1. Introduction

In the last few years, deep-learning-based image segmentation models have witnessed remarkable progress in addressing complex visual scenes, enabling ubiquitous applications ranging from autonomous driving [18] to medical image analysis [44]. In order to ensure their reliability and robustness in real-world scenarios, it is desirable to evaluate them in varied and complex scenes in advance.

As to varied real-world scenarios, segmentation methods are required to be robust to various shifts in terms of local attribute variations (e.g. object color, material, pattern) and global attribute changes (e.g. image style). For example, a train can be painted with different colors and patterns, and a boat is composed of various materials (see Figure 1). Moreover, segmentation methods may encounter images with global attribute changes such as variations in weather conditions and image styles in the real world. Nevertheless, the problem is: “How sensitive are existing segmentation models to local and global attribute changes?”

* Corresponding author.

Regrettably, this research avenue remains under-explored.

The main challenge of robustness evaluation is the lack of high-quality test data with abundant local and global variations. In addition, even if the data can be collected, the annotation cost of the segmentation masks is quite high. The ACDC dataset [53] manually collects samples with adverse weather in city streets. Sun *et al.* [55] used simulation software to build a synthetic dataset, while Multi-weather Cityscapes [46] employed style-transfer models [71] to generate data. However, they only evaluate segmentation models under global variations, such as weather and different scenes of city streets. Moreover, due to the limited capabilities of the simulation software itself and style-transfer models, their data are not consistent with realistic images.

In this paper, we provide a mask-preserved attribute editing pipeline for the robustness evaluation of segmentation models. The generated data can cover diverse variations including both object and image attributes. In particular, inspired by the impressive generation performance of text-guided image editing methods [5, 28, 58], we leverage a pre-trained diffusion model and text instructions to transfer a testing image into images with various attributes via editing. To avoid manually annotating labels for edited images, we need to maintain the structural information of the image, such that the ground-truth segmentation mask of an original image is used as that of its edited ones. However, typical methods often improperly modify the structure of irrelevant regions (*e.g.* removing an object from the background) during attribute editing. To address this issue, we propose mask-guided attention in the diffusion model [50] to consistently edit target attributes at the object level and preserve the structure of an image. Thanks to the proposed module, our tool can edit images with attribute variations in a tuning-free manner, while avoiding manually annotating ground-truth segmentation masks.

Based on our mask-preserved attribute editing pipeline, we construct a benchmark covering various object and image attribute changes, and we evaluate the robustness of existing segmentation methods. We find that, as same as global attributes, local attribute variations also affect segmentation performances. And sensitivity of segmentation methods varies across different attributes. For example, performance declines most on object material variations (See Figure 1). Moreover, the experimental results show open-vocabulary methods with stronger backbones and massive training data can not necessarily exhibit robustness, compared to conventional close-set methods. These findings suggest that object attribute variations have the same importance as image attribute variations to improve robustness.

The main contributions are summarized as follows:

- We provide a mask-preserved attribute editing pipeline that can change various attributes of real images without the requirement of re-collecting segmentation labels.

- We explore the robustness of existing segmentation models to both object and image attribute variations.
- We conduct extensive experiments and find that segmentation models exhibit varying sensitivity to different attribute variations.

2. Related Work

Semantic Segmentation. Conventional semantic segmentation methods [6, 62, 67] obtain significant performance across various benchmark datasets, but fail to generalize well on new environments. Recently, since distinguished performance on alignment between visual and language data from visual language models (VLMs) [33, 49], many open-vocabulary segmentation frameworks [40, 65, 73] are proposed by using language data as auxiliary weak supervision. They normally generate class-agnostic masks, and then use the label embeddings from a Visual Language Model (VLM) to classify the proposal masks. Several benchmarks are proposed to evaluate segmentation models against domain shifts caused by attribute variations. The ACDC benchmark [53] manually collects images and labels from multiple target domain sources. SHIFT [55] employs simulation software to automatically generate data with continuous and discrete scene shifts in city streets. Other works [20, 31, 46, 52, 55] synthesize images by training style transfer models [34, 71]. However, they only focus on global attribute variations, *e.g.* weathers [53] and scenes [55] in autonomous driving. Contrary to them, our research evaluates both local and global attribute variations which can depict real-world scenarios more comprehensively.

Model Diagnosis. Diagnosing the model’s behavior and explaining its mistake is critical for understanding and building robustness. In counterfactual explanation [24] literature, it is common to explore what differences to input image will flip the prediction of the model. Previous methods explain decisions by mining similar failure samples from datasets to provide interpretations [21, 23, 25, 60]. However, the size of datasets constrains the variation types of similar images and hence hinders their overall interpretability. Recently, generation-based approaches [2, 22, 30, 32, 35, 39, 43, 48, 59, 61, 69] achieve remarkable progress in explaining model failures. Luo *et al.* measure models’ sensitivity to pre-defined attributes changes by optimization in latent space of StyleGAN [34]. The closest works to ours are Prabhu *et al.* [48] and ImageNet-E [39], which employ recent diffusion-based image editing techniques and achieve curated control on the diversity of image editing. However, their approaches inevitably interfere with irrelevant regions when editing object attributes and can not be applied to image segmentation tasks. Contrary to them, our proposed approach utilizes object masks to prevent potential interference and generate more realistic variations.

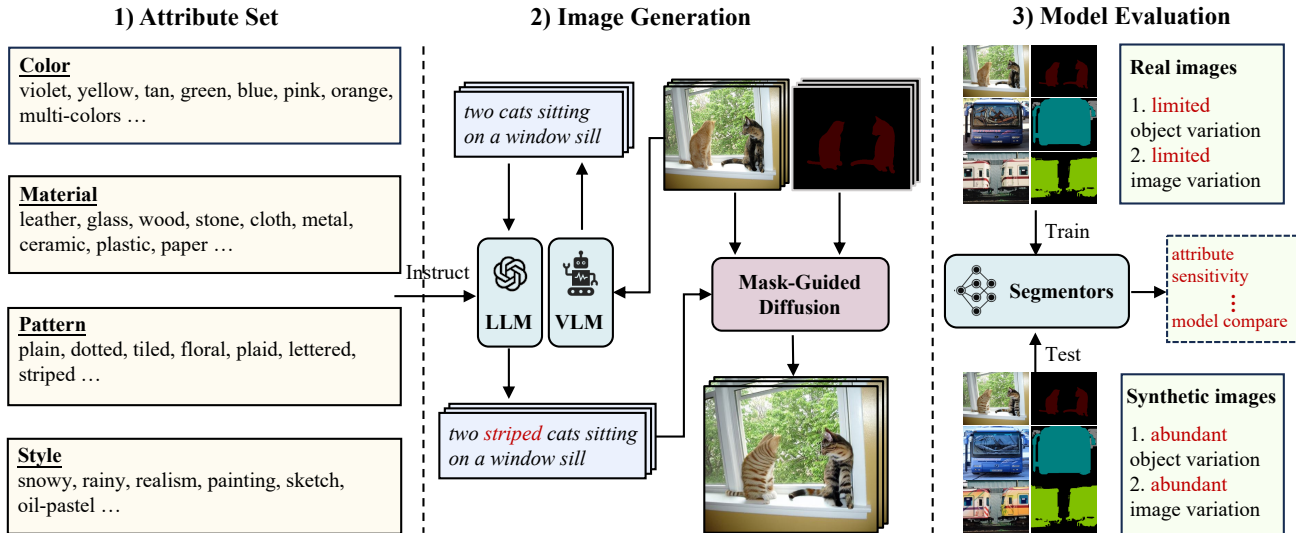


Figure 2. The illustration of our mask-preserved attribute editing pipeline. (1) We construct the Attribute Set which defines local and global variations. (2) We edit real images with different attribute variations with the collaboration of the large language model and diffusion model. (3) The robustness of segmentation models can be evaluated on the edited images against various types of attribute variations.

3. Method

The overall pipeline of our method is illustrated in Figure 2. We first construct the Attribute Set which defines both local and global variations. Then we edit images with attribute variations via the collaboration of a large language model and diffusion model. Based on our design, we can generate samples with diverse objects and styles. Finally, we evaluate the most representative segmentation models and investigate their robustness to different attribute variations. The details of each component are introduced as follows.

3.1. Attribute Set

Generally, attributes [17, 41] are characteristic properties of objects and images, including color, shape, texture, and style. A visual instance can be characterized by the composition of various attributes. For example, a cat may possess *brown color*, *soft texture*, *dotted skin*, *feline shape*. In the existing visual datasets, some attributes, *e.g.* the shape of cats, almost remain unchanged across diverse environments, while other attributes, *e.g.* textures, materials, colors, have better diversity within a specific class [19]. To generate abundant evaluation samples, we construct an attribute set that encompasses all types of variations at both the object and image levels.

To trustfully evaluate the segmentation models, we investigate the attributes that tend to vary in the real world. In summary, we manually select attributes under four categories: (a) Local instance color, *e.g.* blue, red. (b) Local instance material, *e.g.* metal, wood, stone. (c) Local appearance pattern, *e.g.* plain, dotted, lettered. (d) Global image

style, *e.g.* photo, painting, different weathers. Some examples are shown in Figure 2. By changing a single attribute and fixing the others, we can generate sufficient test samples with abundant variations.

3.2. Image Generation

Since the generated images are used to evaluate segmentation models, they should have the following properties:

- **Credibility**: generated samples should be credible, adhering to real-world conditions, which implies that any modifications should be realistic.
- **Fidelity**: generated images should not disrupt irrelevant information and contravene layout defined by original semantic segmentation labels. (Since acquiring segmentation labels is laborious, we must ensure their correctness after generation.)

To ensure the above properties in our generated test samples, we design a mask-preserved attribute editing pipeline. The procedural details are illustrated as follows.

Text Manipulation. Since existing generative models are usually instructed by language, we first obtain the textual description of a real test image and incorporate different attributes linguistically. In this way, our method can be user-friendly. Specifically, we use a pre-trained large vision-language model (BLIP-2 [38]) to acquire the textual description. To perform meaningful modification to text, we divided a text description into several editable components based on linguistic formal: domain, subjects and their adjectives (attributes), actions (verbs), and backgrounds (objects). For example, in “a photo of a white horse on the grass”, the domain is “a photo”, the subject is “horse” and

its adjective is “white”, the action is “on” and the background is “grass”. Then we employ a language model (GPT-3.5 turbo) [4] to generate textual variations by instructions built using considered attributes. An example is change ‘a horse on the grass’ to ‘a black horse on the grass’. Please refer to our supplementary material for more prompts and examples. Finally, we use the manipulated text descriptions as input prompts to guide image editing.

Mask-Guided Diffusion. There are several efforts [3, 45] to edit specific image areas by the guidance of object mask. However, most of them directly generate visual content according to the input text without injecting the features from the original images. Therefore, they can only achieve limited performance in object attribute manipulation. Based on this observation, we utilize the setting suggested by [5, 28, 47, 58], where features extracted from the real image are directly injected into the diffusion process of the generated image. However, we found that they can hardly change one specific portion of the original images without affecting others. For instance, manipulating the color of an object will marginally change the details of the adjacent background (refer to Figure 5). This heavily diminishes the assessment reliability of generated test images, as the segmentation performance on whole images is susceptible to any pixel perturbation [1]. From our perspective, the main reason is the attention maps in the diffusion process can only represent the rough spatial layout of objects but fail to accurately depict the detailed object localization. One trivial solution is directly replacing the background with pixel features from original images in each time step [12]. However, this induces a mass of artifacts in the object boundary.

Inspired by the previous work [36], we propose the Mask-Guided Attention which utilizes the object segmentation masks to rectify attention maps in the diffusion process. We denote a real image as $I \in \mathbb{R}^{3 \times H \times W}$ with height H and width W and its semantic layout label as L . Given the text guidance P , the objective of our method is to generate image I^* which complies with P , and strictly preserves the semantic layout defined by L . Concretely, supposed that a binary object mask $S \in \mathbb{R}^{H \times W}$ is extracted from the semantic map L , we rectify the attention map as follows:

$$A = \text{softmax} \left(\frac{QK^\top + R \odot \phi(S')}{\sqrt{d}} \right), \quad (1)$$

where the query-key pair conditioning map $R \in \mathbb{R}^{|Q| \times |K|}$ defines whether to rectify the attention score for a particular pair, we set different configurations of R in self and cross attention layers. $\phi(\cdot)$ is a flatten function, and S' is a rectification map confining the region of attention scores which is defined as:

$$S'(x, y) = \begin{cases} 0 & S(x, y) = 1 \\ -\infty & \text{otherwise.} \end{cases} \quad (2)$$

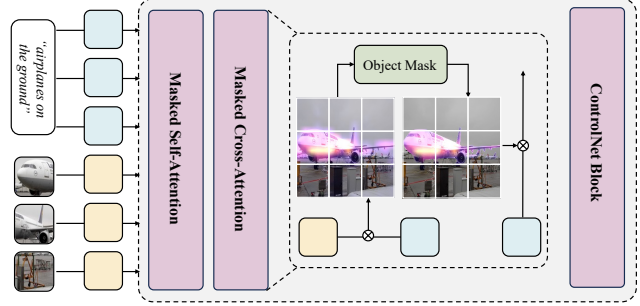


Figure 3. The illustration of block in our method. Our attention mechanism utilizes an object mask to rectify attention maps in self-attention and cross-attention layers. We adopt ControlNet block [70] to further restrict the semantic layout.

In the cross-attention layer, the semantic knowledge of textual features is incorporated into intermediate visual features in which objects appearance and image layouts gradually emerge. The attention maps between each text token and visual token determine the intensity and area of influence of corresponding text words. We rectify the corresponding region and intensity of cross-attention scores of text words, the conditioning map R at location (i, j) is

$$R(i, j) = \begin{cases} c & i \in \{K\}, \forall j \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where K is the set of token indexes corresponding to the object mask S and c indicates the strengthens or weakens the extent to which text tokens affect the resulting image.

In the self-attention layer, the intermediate visual features compute affinities with each other, allowing to retain the coherent spatial layout and shape details. We leverage the object mask S to constrain interactions between different regions. For example, we retain the attention scores between tokens within the same regions and suppress that between tokens within different regions. Concretely, we define the conditioning map R at location (i, j) as:

$$R(i, j) = \begin{cases} 1 & \phi(S)(i) = 1 \text{ and } \phi(S)(j) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

We incorporate the designed attention mechanism into decoding blocks of U-Net [51] in the Latent Diffusion Model [50]. In global style editing, we denote S as a matrix of ones where every entry is equal to one.

In order to further control that generated images rigorously adhere to their original semantic layout, we integrate existing controllable generative module ControlNet [70] into each blocks. The whole designed block is illustrated in Figure 3. In this way, our pipeline can obtain samples with abundant diversity, without changing their original semantic segmentation labels. And we can hence expand existing benchmarks requiring no laborious label collection.

Table 1. Quantitative results of all segmentation models under different attribute variations in our Pascal-EA. The results are reported using mIoU (\uparrow), and the best results of close-set methods and open-vocabulary methods are separately bold. We replace the original validation set with the reconstructed images obtained by our pipeline (denoted as Val).

Method	Backbone	Val	Color			Material				Pattern		Style		Overall	
			violet	pink	multi-color	wood	stone	metal	paper	dotted	striped	snowy	painting		sketch
DeepLabV3+ [7]	RN-50	73.65	66.46	67.29	63.34	52.96	54.10	55.07	41.98	62.85	63.59	65.14	61.39	15.33	55.79 (-17.86)
OCRNet [67]	HR-48	75.32	67.78	69.98	65.65	56.79	56.76	58.69	46.42	65.41	67.41	67.71	65.45	22.06	59.18 (-16.14)
Segmenter [54]	ViT-B	80.66	74.14	75.2	72.77	65.84	63.11	67.15	54.89	74.48	75.26	75.22	75.86	31.19	67.09 (-13.57)
Segformer [64]	MIT-B3	79.45	71.06	72.37	69.01	60.55	61.6	63.02	51.88	70.8	72.02	71.91	71.23	21.45	63.08(-16.38)
Mask2Former [8]	Swin-B	90.81	80.10	80.15	72.92	61.14	62.92	62.77	49.48	77.29	77.78	76.21	77.06	21.82	66.64 (-24.17)
CATSeg [9]	Swin-B	95.59	88.06	88.51	87.28	82.93	81.96	82.26	74.79	88.63	90.30	90.06	92.95	48.06	82.98 (-12.61)
OVSeg [40]	Swin-B	92.83	83.23	85.84	85.17	79.08	77.01	79.70	70.93	86.12	86.06	86.93	89.74	42.79	79.38 (-13.45)
ODISE [65]	SD-1.5	91.54	71.96	72.05	70.05	70.23	73.46	64.50	54.59	81.55	84.33	82.61	85.97	26.32	69.79 (-21.75)
X-Decoder [72]	Focal-T	90.42	81.70	83.89	79.36	70.59	69.64	72.72	52.28	83.58	83.85	83.22	87.05	34.99	73.57 (-16.85)
SEEM [73]	Focal-T	89.21	80.30	82.74	78.55	70.79	69.56	71.37	56.94	81.76	82.42	83.71	84.91	34.45	73.13 (-16.09)
SEEM [73]	SAM-B	89.10	77.14	79.48	72.20	61.41	63.37	64.98	50.95	77.47	77.80	76.21	77.06	28.18	67.19 (-21.91)
Overall		86.21	76.54	77.95	74.21	66.57	66.68	67.48	55.00	77.27	78.26	78.08	78.97	29.69	
			-9.68	-8.27	-12.01	-19.65	-19.54	-18.74	-31.22	-8.95	-7.96	-8.14	-7.25	-56.53	

4. Experiments

In this section, we first introduce the experimental setups. Then we benchmark various segmentation models on our edited images and discuss on the results. At last, we assess the quality of our edited images.

4.1. Experimental Setup

Benchmark. We construct our benchmark **Pascal-EA (Editable Attributes)** by editing images in the validation set of Pascal VOC dataset [16]. Most samples in Pascal VOC [16] are object-centric, making it more suitable for exploring the impact of local variables on segmentation performance. Considering potential errors in generated images, we manually discarded ones with noticeable errors to ensure the quality of the benchmark. Since we want to stress-test models, we manually choose the following uncommon attributes to construct the benchmark. (1) *object colors*: violet, pink, and multicolor, which provide a deviation from objects’ common colors. (2) *object material*: wood, stone, metal and paper. (3) *object pattern*: dotted and striped, which have unusual appearance compared to common objects. (4) *image style*: snow, painting, and sketch which provides blurring circumstances, abstract descriptions of realistic objects, and absence of color and texture visual cues, respectively. Please refer to our supplementary material for qualitative results.

Target Models. We conduct evaluation experiments on various architectures, including close-set and open-vocabulary methods. For close-set methods, we adopt convolution-based models DeepLabV3+ [7], OCRNet [67], and transformer-based models Segmenter [54], Segformer [64], and the universal image segmentation architecture Mask2Former[8]. For open-vocabulary methods, we adopt transformer-based approaches CATSeg [9] and OVSeg [40], the typical diffusion-based method ODISE [65], and two

SOTA generalized frameworks X-Decoder [72] and SEEM [73]. We implement target models with MMsegmentation [10] and Detectron2 [63]. All models are trained in the training set of Pascal VOC [16], and we use the released official weights and original recipes to compare at their best.

Metric. We report the results of mIoU as the metric.

4.2. Evaluation Results

Table 1 presents segmentation performances under different attribute variations. Initially, our generated images, featuring unusual visual cues, reduce performance across all models to varying extents. However, it is noteworthy that the models exhibit a greater vulnerability to material alteration than to adjustments in other local attributes, such as color and pattern. For instance, under the material-based variations, performance decline ranged from 18.74% to 31.22%, whereas under the color-based variations, the decline ranged from 8.27% to 12.01%. Changes in the object’s appearance pattern have less impact on the models’ performance compared to color and material. Since the choice of material significantly influences the visual texture of an object’s appearance, we speculate that segmentation models are significantly sensitive to texture, which aligns with observations found in classifiers [22]. Furthermore, results reveal a notable disparity in model performance when comparing sketch images to other global styles. For instance, performances decrease 56.53% under sketch style variation, while dropping 7.25% and 8.14% under painting and snow scenarios. We argue that sketch images introduce substantial perturbations to model performance as they represent visual content solely through line-based representations without any texture and color information.

To compare the overall robustness of models, we also compute the average performance decline under all attributes in the last column of Table 1. Firstly, in close-set

Table 2. mIoU (\uparrow) of four augmentation algorithms on our Pascal-EA. The best results are in bold.

Method	Val	Color			Material				Pattern		Style		
		violet	pink	multi-color	wood	stone	metal	paper	dotted	striped	snowy	painting	sketch
Mask2Former [8]	90.81	80.10	80.15	72.92	61.14	62.92	62.77	49.48	77.29	77.78	76.21	77.06	21.82
CutOut [15]	91.44	80.68	82.02	73.44	62.23	64.70	60.55	50.28	75.32	78.32	75.95	75.41	19.55
CutMix [68]	92.01	80.79	83.45	75.53	64.81	66.68	62.87	49.83	77.84	78.73	76.03	76.48	21.62
Digital Corruption [26]	91.19	73.29	75.47	69.63	57.96	55.51	56.40	45.01	75.16	75.79	79.22	77.98	22.59
AugMix [27]	91.88	75.38	76.49	70.44	60.29	58.92	58.73	48.84	75.33	76.01	80.46	80.92	23.05

methods, we observe that Segmenter [54] exhibits best performances under material variations while Mask2Former [8] performs best in others. And we notice that the overall performance decline of DeepLabV3+ [7] and OCRNet [67] is even less than that observed in transformer architectures such as Segformer [64] and Mask2former [8]. This implies that recent transformer-based models have greater segmentation accuracy than CNN-based methods, but they do not necessarily show an improvement in robustness. Previous studies [14] have shown that recent transformers are remarkably more robust than baseline models as the domain gap is large, which is opposite to our findings. We think the reason may be our edited images have a smaller domain gap compared to [14]. Secondly, in open-vocabulary frameworks, we notice ODISE [65] exhibits considerably inferior performance compared to CATSeg [9] and OVSeg [40]. Since they have only a discrepancy in mask proposal backbone, we argue that Stable Diffusion [50] possesses comparatively worse robustness than specialized backbone [42] in segmentation. Besides, two holistic-trained frameworks X-Decoder [72] and SEEM [73] are worse than other approaches. We speculate the reason relies on the utilization of the Vision-Language Model (VLM) which classifies proposed masks. These two frameworks retrain a vision-language alignment space from scratch, and other approaches leverage off-the-shelf CLIP [49] trained with much more abundant data. Finally, we also evaluate the performance of the recent Segment-Any-Thing (SAM) [37] backbone using the SEEM [73] framework. SEEM method with different backbone exhibits similar performances in the in-distribution set while gaining considerable disparity in our generated benchmarks, in which the SAM fails to show better robustness. In summary, all phenomena reveal that better in-distribution performances do not indicate better performance under our variations, and stronger backbones and more training data do not necessarily bring robustness improvement. Please refer to our supplementary materials for more results and analysis.

4.3. Discussion

Effectiveness of augmentation techniques. We also conduct experiments to explore the effectiveness of existing data augmentation algorithms. We select several representa-




Figure 4. Average mIoU (\uparrow) of all segmentation methods under the combination of two different attribute variations.

tive methods: CutOut [15], CutMix [68], AugMix [27] and Digital Corruption [26], and use vanilla Mask2Former [8] as a baseline model. The results are presented in Table 2. We can obtain several observations: (1) All augmentation methods improve in-distribution performances. (2) While CutOut [15] and CutMix [68] exhibit more robustness on object-level variations of color, material, and pattern, they can not improve performances on image-level style variations. Compared to the previous two, Digital Corruption [26] and AugMix [27] exhibit opposite characteristics. We speculate that the reason behind these divergent behaviors is local object variation edits the visual content information whereas global variation tends to edit visual styles. Since Digital Corruption [26] and AugMix [27] change the style information of images by methods such as brightness and sharpness, the model trained on it can only exhibit advancement on global style variations. However, CutOut [15] and CutMix [68] directly change the visual content of specific regions, they can act better under object variations.

Multiple Attribute Variation. We conduct a statistical analysis of the average model performances when exposed to samples with a combination of two distinct attribute variations. The results are presented in Figure 4. It is evident that when two types of variation are combined, the overall performance gains more deterioration compared to that under a single one. This observation indicates that addressing multiple attribute variations poses a greater challenge than dealing with a single in isolation.

Variation Degree. The extent to which edited images deviated from the original ones can be adjusted by coefficient c in the Equation 3, higher value of c means greater variation. We adjust different values of c in the image editing process to investigate the impact of variation degree on seg-

Table 3. mIoU (\uparrow) of five segmentation methods under three adverse weather conditions with different variation degrees. Segmentation performances gradually deteriorate as shift increases.

Val	Material: Wood						
	Style:snow			Material: Wood			
	$c = 1.0$	$c = 2.0$	$c = 3.0$	$c = 1.0$	$c = 2.0$	$c = 3.0$	
DeepLabV3+[6]	73.65	65.14	63.79	59.20	52.96	49.38	45.82
OCRNet [67]	75.32	67.71	63.55	60.29	56.79	51.67	46.23
Segmenter [54]	80.66	75.22	68.03	62.09	65.84	60.28	56.89
Segformer [64]	79.45	71.91	67.85	60.88	60.55	55.79	49.75
Mask2Former [8]	90.81	76.21	70.24	63.49	61.14	57.34	51.63

mentation performance. Experimental results are presented in Table 3. We notice a gradual degradation in all models’ performances as the degree of variation intensifies across all scenarios. This reveals our images do not make models collapse like the adversarial attack method [1], and hence have the capability of continuous stress-testing models.

Practical application. Our pipeline could also be utilized as a data augmentation approach to improve robustness of segmentation methods under hard samples. Please refer to our supplementary material for results details.

4.4. Image Quality Assessment

In this section, we assess the quality of our generated images by comparing them with previous diffusion-based image editing approaches, and popular benchmarks of segmentation robustness evaluation.

Comparison with diffusion-based image editing methods. Since our proposed pipeline serves as an evaluator for segmentation models, the reliability of edited images is our primary concern. We focus on object internal structure maintenance and no disruption to irrelevant information. Thus, we adopt LPIPS and DINO Dist [57], which measure structural distances, as our main metrics. We compare our proposed method with state-of-the-art image editing algorithms PnP [58], Prompt-to-Prompt [28]. Since there are several efforts to utilize object masks as guidance, such as DiffEdit [12] and Blended Latent Diffusion [3], we take them as additional comparisons. We also conducted an ablation study to demonstrate the effectiveness of the Mask-Guided Attention mechanism and ControlNet. Evaluations are performed on tasks of random replacing object color and material using Pascal VOC dataset [16].

The results are reported in Table 4. Firstly, previous approaches have the same performances of high scores in DINO Dist and LPIPS, but we argue that the reasons behind them are different. Since DiffEdit [12] and Blended-LDM [3] use a mask to specify the edited region, they can ensure that irrelevant background information is not affected, but they do not achieve fine-grained control on object inner structures, e.g. editing color will change other core properties. PnP [58] acts differently: it achieves precise control of the object’s inner structural consistency but will disturb details of the adjacent background. Therefore, these ap-

Table 4. Quantitative results of different diffusion models in editing color and material attributes of objects on Pascal VOC [16].

Method	Material		Color	
	DINO Dist (\downarrow)	LPIPS (\downarrow)	DINO Dist (\downarrow)	LPIPS (\downarrow)
Blended-LDM [3]	0.081	0.397	0.079	0.375
DiffEdit [12]	0.068	0.313	0.053	0.321
Prompt-to-Prompt [28]	0.044	0.367	0.030	0.472
PnP (Baseline) [58]	0.052	0.319	0.047	0.408
PnP w/ ControlNet [70]	0.010	0.284	0.010	0.365
PnP w/ MGA	0.017	0.250	0.012	0.309
PnP w/ Both (Ours)	0.003	0.185	0.001	0.156

proaches could impede the reliability of generated images for evaluation. Secondly, we observe that since our method induces extra spatial constraints in the diffusion process, we achieve the best performance. As shown in Figure 5, we can faithfully change object local attributes without affecting other information. Finally, we also compare our methods with baseline in complex scenes, qualitative results are shown in Figure 5. Our proposed method achieves improvement in semantic structure preservation under both object-centric images and complex scenes.

In summary, we prove that existing editing approaches are not adequate as an evaluator for segmentation models, and demonstrate the reliability of our proposed pipeline.

Comparison with other benchmarks. Existing benchmarks [46, 52, 52, 53, 55] for semantic segmentation robustness heavily focuses on adverse weather conditions of autonomous driving environments. They utilize generative models [31, 71] or simulator [52, 55] to synthesize samples, to prevent manually collecting samples and labels. In this way, we argue that some properties of synthetic images could impede the reliability of evaluation results. We compare our edited images with previous benchmarks in terms of image reality, to demonstrate the reliability of our pipeline. We adopt (1) CLIP Acc [29], which calculates the percentage of instances where the target image has a higher similarity to the target text than to the source text, (2) Fréchet Inception Distance (FID), which measures similarity between the distribution of real images and the distribution of generated images.

We separately generate images for four adverse kinds of weather (snow, rain, fog, and night) using the Cityscapes dataset [11] as source images. And we adopt three existing synthetic benchmarks Multi-weather [46], Fog Cityscapes [52], and SHIFT [55] as comparative methods. The FID score is computed by corresponding real images in ACDC [53]. The quantitative results are illustrated in Table 5. In most weathers, our generated data are more realistic and distributional closer to real images than previous synthetic benchmarks. At night, our images are drastically greater than simulation-based benchmark SHIFT [55], but worse than Multi-weather [46] which utilizes a set of GAN and CycleGAN models [13, 71] to transfer styles. We think this is because diffusion models struggle in manipulating bright-

Table 5. Comparison of our generated images with previous benchmarks under four adverse weather conditions, where generative and simulator refer to that images are generated by generative methods and simulator, respectively.

	Model	Snow		Rain		Fog		Night	
		CLIP Acc \uparrow	FID \downarrow	CLIP Acc \uparrow	FID \downarrow	CLIP Acc \uparrow	FID \downarrow	CLIP Acc \uparrow	FID \downarrow
ACDC [53]	real	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000
Multi-weather [46]	<i>generative</i>	0.163	197.48	0.852	189.85	-	-	0.906	154.30
Fog Cityscapes [52]	simulator	-	-	-	-	0.940	164.00	-	-
SHIFT [55]	simulator	-	-	0.980	242.28	0.955	272.58	0.914	274.45
Ours	<i>generative</i>	1.000	150.61	0.989	143.47	0.999	137.15	0.965	186.69

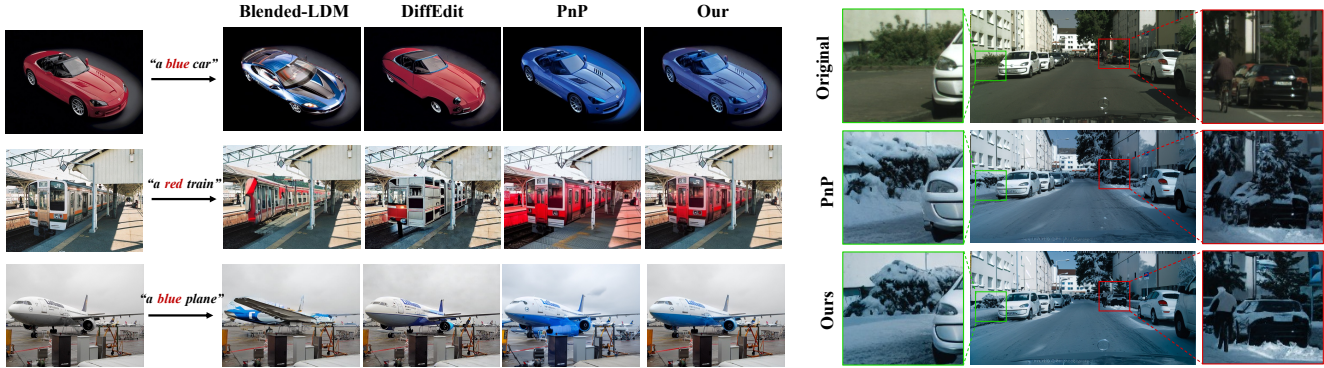


Figure 5. Qualitative comparison of edited images. **Left:** Manipulating object color in Pascal VOC [16]. **Right:** Changing to snowy day in Cityscapes [11]. Our method achieves the best performance in structure preservation and object consistency.

Table 6. Comparison of our pipeline with Stylized COCO [56] in local object and background variations. Our generated images have superior image reality and fidelity.

Benchmarks	Model	CLIP Acc \uparrow	DINO Dist \downarrow
Stylized-object [56]	AdaIN [31]	0.574	0.075
Ours-object	Diffusion	0.895	0.002
Stylized-background [56]	AdaIN [31]	0.678	0.084
Ours-background	Diffusion	0.965	0.004

ness and light information [66].

Furthermore, since previous Stylized COCO [56] generates local style variations on objects to study robustness in segmentation, we also compare our method with it. Following its setting, we generate images where objects and backgrounds have random art styles, respectively. From the results in Table 6, our method significantly surpasses it in image reality and structural preservation. In essence, we demonstrate that samples edited by our diffusion model have better image reality compared to conventional synthetic benchmarks for evaluating segmentation models. All experimental results can serve as evidence that our pipeline can be a better substitution for real images, especially since acquiring real samples and labels is laborious.

5. Conclusion

In this paper, we provide a pipeline that can precisely edit the visual attributes of real images and preserve their orig-

inal mask labels. With this pipeline, we construct a benchmark and evaluate the robustness of diverse segmentation models against different object and image attribute variations. Experimental results reveal that most models are vulnerable to object attribute changes. Meanwhile, advanced models with stronger backbones and massive training data do not necessarily show better robustness. Our experiments suggest object attributes should be taken into account for improving segmentation robustness. We also demonstrate the quality of our edited images and their reliability as test samples by comparing them with popular synthetic benchmarks and existing image editing methods.

Our work has some limitations. First, due to the failure modes of the diffusion model, it is difficult to edit the attributes of the person. Second, since inherent biases in diffusion models, altering an attribute could occasionally change other appearance information. For example, sometimes editing an object’s material to wood may deviate its original color to brown. We would like to explore how to avoid the spurious editing problem in future work.

Acknowledgements The authors thank Xinran Wang for his valuable feedback and discussion. This work was supported in part by National Natural Science Foundation of China (NSFC) No. 62106022, 62225601, U23B2052, in part by Beijing Natural Science Foundation Project No. Z200002 and in part by Youth Innovative Research Team of BUPT No. 2023QNTD02.

References

- [1] Anurag Arnab, Ondrej Miksik, and Philip HS Torr. On the robustness of semantic segmentation models to adversarial attacks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 888–897, 2018. 4, 7
- [2] Maximilian Augustin, Valentyn Boreiko, Francesco Croce, and Matthias Hein. Diffusion visual counterfactual explanations. *Advances in Neural Information Processing Systems*, 35:364–377, 2022. 2
- [3] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM Transactions on Graphics (TOG)*, 42(4):1–11, 2023. 4, 7
- [4] 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*, 33:1877–1901, 2020. 4
- [5] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. *arXiv preprint arXiv:2304.08465*, 2023. 2, 4
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2, 7
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 5, 6
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 5, 6, 7
- [9] Seokju Cho, Heeseong Shin, Sunghwan Hong, Seungjun An, Seungjun Lee, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2303.11797*, 2023. 5, 6
- [10] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 5
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 7, 8
- [12] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. 4, 7
- [13] Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance. In *European Conference on Computer Vision*, pages 88–105. Springer, 2022. 7
- [14] Pau de Jorge, Riccardo Volpi, Philip HS Torr, and Grégory Rogez. Reliability in semantic segmentation: Are we on the right track? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7173–7182, 2023. 6
- [15] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 6
- [16] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 5, 7, 8
- [17] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1778–1785. IEEE, 2009. 3
- [18] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020. 1
- [19] Ruth Fong and Andrea Vedaldi. Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8730–8738, 2018. 3
- [20] Gianni Franchi, Nacim Belkhir, Mai Lan Ha, Yufei Hu, Andrei Bursuc, Volker Blanz, and Angela Yao. Robust semantic segmentation with superpixel-mix. *arXiv preprint arXiv:2108.00968*, 2021. 2
- [21] Irena Gao, Gabriel Ilharco, Scott Lundberg, and Marco Tulio Ribeiro. Adaptive testing of computer vision models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4003–4014, 2023. 2
- [22] 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. *arXiv preprint arXiv:1811.12231*, 2018. 2, 5
- [23] Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee. Counterfactual visual explanations. In *International Conference on Machine Learning*, pages 2376–2384. PMLR, 2019. 2
- [24] Riccardo Guidotti. Counterfactual explanations and how to find them: literature review and benchmarking. *Data Mining and Knowledge Discovery*, pages 1–55, 2022. 2
- [25] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. Grounding visual explanations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 264–279, 2018. 2

- [26] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. **6**
- [27] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. **6**
- [28] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. **2, 4, 7**
- [29] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. **7**
- [30] Phillip Howard, Avinash Madasu, Tiep Le, Gustavo Lujan Moreno, and Vasudev Lal. Probing intersectional biases in vision-language models with counterfactual examples. *arXiv preprint arXiv:2310.02988*, 2023. **2**
- [31] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. **2, 7, 8**
- [32] Guillaume Jeanneret, Loïc Simon, and Frédéric Jurie. Text-to-image models for counterfactual explanations: a black-box approach. *arXiv preprint arXiv:2309.07944*, 2023. **2**
- [33] 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 *International conference on machine learning*, pages 4904–4916. PMLR, 2021. **2**
- [34] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. **2**
- [35] Saeed Khorram and Li Fuxin. Cycle-consistent counterfactuals by latent transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10203–10212, 2022. **2**
- [36] Yunji Kim, Jiyoung Lee, Jin-Hwa Kim, Jung-Woo Ha, and Jun-Yan Zhu. Dense text-to-image generation with attention modulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7701–7711, 2023. **4**
- [37] 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. **6**
- [38] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. **3**
- [39] Xiaodan Li, Yuefeng Chen, Yao Zhu, Shuhui Wang, Rong Zhang, and Hui Xue. Imagenet-e: Benchmarking neural network robustness via attribute editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20371–20381, 2023. **2**
- [40] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. **2, 5, 6**
- [41] Kongming Liang, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Unifying visual attribute learning with object recognition in a multiplicative framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1747–1760, 2019. **3**
- [42] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. **6**
- [43] Jinqi Luo, Zhaoning Wang, Chen Henry Wu, Dong Huang, and Fernando De la Torre. Zero-shot model diagnosis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11631–11640, 2023. **2**
- [44] Maciej A Mazurowski, Haoyu Dong, Hanxue Gu, Jichen Yang, Nicholas Konz, and Yixin Zhang. Segment anything model for medical image analysis: an experimental study. *Medical Image Analysis*, 89:102918, 2023. **1**
- [45] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. **4**
- [46] Valentina Muşat, Ivan Fursa, Paul Newman, Fabio Cuzzolin, and Andrew Bradley. Multi-weather city: Adverse weather stacking for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2915, 2021. **2, 7, 8**
- [47] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. **4**
- [48] Viraj Prabhu, Sriram Yenamandra, Prithvijit Chattopadhyay, and Judy Hoffman. Lance: Stress-testing visual models by generating language-guided counterfactual images. *arXiv preprint arXiv:2305.19164*, 2023. **2**
- [49] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. **2, 6**
- [50] 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. **2, 4, 6**
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. **4**

- [52] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126:973–992, 2018. 2, 7, 8
- [53] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10765–10775, 2021. 2, 7, 8
- [54] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7262–7272, 2021. 5, 6, 7
- [55] Tao Sun, Mattia Segu, Janis Postels, Yuxuan Wang, Luc Van Gool, Bernt Schiele, Federico Tombari, and Fisher Yu. Shift: a synthetic driving dataset for continuous multi-task domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21371–21382, 2022. 2, 7, 8
- [56] Johannes Theodoridis, Jessica Hofmann, Johannes Maucher, and Andreas Schilling. Trapped in texture bias? a large scale comparison of deep instance segmentation. In *European Conference on Computer Vision*, pages 609–627. Springer, 2022. 8
- [57] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing vit features for semantic appearance transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10748–10757, 2022. 7
- [58] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 2, 4, 7
- [59] Joshua Vendrow, Saachi Jain, Logan Engstrom, and Alexander Madry. Dataset interfaces: Diagnosing model failures using controllable counterfactual generation. *arXiv preprint arXiv:2302.07865*, 2023. 2
- [60] Pei Wang and Nuno Vasconcelos. Scout: Self-aware discriminant counterfactual explanations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8981–8990, 2020. 2
- [61] Olivia Wiles, Isabela Albuquerque, and Sven Gowal. Discovering bugs in vision models using off-the-shelf image generation and captioning. *arXiv preprint arXiv:2208.08831*, 2022. 2
- [62] Huikai Wu, Junge Zhang, Kaiqi Huang, Kongming Liang, and Yizhou Yu. Fastfcn: Rethinking dilated convolution in the backbone for semantic segmentation. *arXiv preprint arXiv:1903.11816*, 2019. 2
- [63] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 5
- [64] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 5, 6, 7
- [65] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 2, 5, 6
- [66] Yuyang Yin, Dejia Xu, Chuangchuang Tan, Ping Liu, Yao Zhao, and Yunchao Wei. Cle diffusion: Controllable light enhancement diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8145–8156, 2023. 8
- [67] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, pages 173–190. Springer, 2020. 2, 5, 6, 7
- [68] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 6
- [69] Mehdi Zemni, Mickaël Chen, Éloi Zablocki, Hédi Ben-Younes, Patrick Pérez, and Matthieu Cord. Octet: Object-aware counterfactual explanations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15062–15071, 2023. 2
- [70] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 4, 7
- [71] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2, 7
- [72] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15116–15127, 2023. 5, 6
- [73] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023. 2, 5, 6