

Intra-Source Style Augmentation for Improved Domain Generalization

Yumeng Li^{1,2} Dan Zhang^{1,4} Margret Keuper^{2,3} Anna Khoreva^{1,4}

¹Bosch Center for AI ²University of Siegen ³MPI for Informatics ⁴University of Tübingen

{yumeng.li, dan.zhang2, anna.khoreva}@de.bosch.com margret.keuper@uni-siegen.de

Abstract

The generalization with respect to domain shifts, as they frequently appear in applications such as autonomous driving, is one of the remaining big challenges for deep learning models. Therefore, we propose an intra-source style augmentation (ISSA) method to improve domain generalization in semantic segmentation. Our method is based on a novel masked noise encoder for StyleGAN2 inversion. The model learns to faithfully reconstruct the image preserving its semantic layout through noise prediction. Random masking of the estimated noise enables the style mixing capability of our model, i.e. it allows to alter the global appearance without affecting the semantic layout of an image. Using the proposed masked noise encoder to randomize style and content combinations in the training set, ISSA effectively increases the diversity of training data and reduces spurious correlation. As a result, we achieve up to 12.4% mIoU improvements on driving-scene semantic segmentation under different types of data shifts, i.e., changing geographic locations, adverse weather conditions, and day to night. ISSA is model-agnostic and straightforwardly applicable with CNNs and Transformers. It is also complementary to other domain generalization techniques, e.g., it improves the recent state-of-the-art solution RobustNet by 3% mIoU in Cityscapes to Dark Zürich.

1. Introduction

The varying environment with potentially diverse illumination and adverse weather conditions makes challenging the deployment of deep learning models in an open-world [60, 82]. Therefore, improving the generalization capability of neural networks is crucial for safety-critical applications such as autonomous driving (see for example Fig. 1). While generally the target domains can be inaccessible or unpredictable at training time, it is important to train a generalizable model, based on the known (source) domain, which may offer only a limited or biased view of the real world [7, 61].

Diversity of the training data is considered to play an

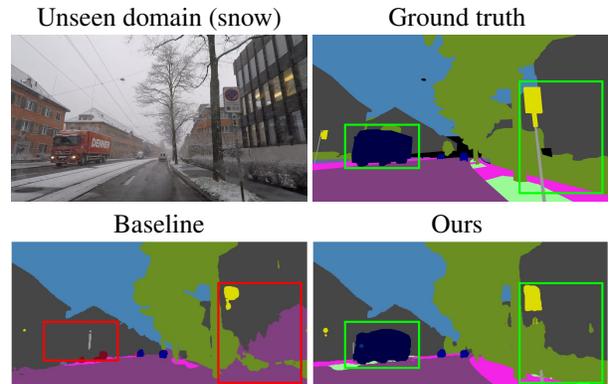


Figure 1. Semantic segmentation results of HRNet [69] on unseen domain (snow), trained on Cityscapes [14] and tested on ACDC [60]. The model trained with our ISSA can successfully segment the truck, while the baseline model fails completely.

important role for domain generalization, including natural distribution shifts [63]. Many existing works assume that multiple source domains are accessible during training [4, 30, 34, 44, 46, 47, 85]. For instance, Li *et al.* [44] divide source domains into meta-source and meta-target to simulate domain shift for learning; Hu *et al.* [30] propose multi-domain discriminant analysis to learn a domain-invariant feature transformation. However, for pixel-level prediction tasks such as semantic segmentation, collecting diverse training data involves a tedious and costly annotation process [8]. Therefore, improving generalization from a *single source domain* is exceptionally compelling, particularly for semantic segmentation.

One pragmatic way to improve data diversity is by applying data augmentation. It has been widely adopted in solving different tasks, such as image classification [26, 28, 67, 80, 86], GAN training with limited data [33, 37], or pose estimation [6, 55, 68]. One line of data augmentation techniques focuses on increasing the content diversity in the training set, such as geometric transformation (e.g., cropping or flipping), CutOut [17], and CutMix [79]. However, CutOut and CutMix are ineffective on natural domain shifts as reported in [63]. Style augmentation, on the other hand, only modifies the style - the non-semantic appearance such as texture and color of the image [20] - while preserv-

ing the semantic content. By diversifying the style and content combinations, style augmentation can reduce overfitting to the style-content correlation in the training set, improving robustness against domain shifts. Hendrycks corruptions [25] provide a wide range of synthetic styles, including weather conditions. However, they are not always realistic looking, thus being still far from resembling natural data shifts. In this work, we propose an intra-source style augmentation (ISSA) strategy for semantic segmentation, aiming to improve the style diversity in the training set without extra labeling effort or using extra data sources.

Our augmentation technique is based on the inversion of StyleGAN2 [39], which is the state-of-the-art unconditional Generative Adversarial Network (GAN) and thus ensures high quality and realism of synthetic samples. GAN inversion allows to encode a given image to latent variables, and thus facilitates faithful reconstruction with style mixing capability. To realize ISSA, we learn to separate semantic content from style information based on a single source domain. This allows to alter the style of an image while leaving the content unchanged. Specifically, we make use of the styles extracted within the source domain and mix them up. Thus, we can increase the data diversity and alleviate the spurious correlation in the given training data.

The faithful reconstruction of images with complex structures such as driving scenes is non-trivial. Prior methods [3, 18, 57, 58, 76] are mainly tested on simple single-object-centric datasets, e.g., CelebA-HQ [36], FFHQ [38], or LSUN [78]. As shown in [2], extending the native latent space of StyleGAN2 with a stochastic noise space can lead to improved inversion quality. However, all style *and* content information will be embedded in the noise map, leaving the latent codes inactive in this setting. Therefore, to enable the precise reconstruction of complex driving scenes as well as style mixing, we propose a masked noise encoder for StyleGAN2. The proposed random masking regularization on the noise map encourages the generator to rely on the latent prediction for reconstruction. Thus, it allows to effectively separate content and style information and facilitates realistic style mixing, as shown in Fig. 2.

In summary, we make the following contributions:

- We propose a masked noise encoder for GAN inversion, which enables high quality reconstruction and style mixing of complex scene-centric datasets.
- We explore GAN inversion for intra-source data augmentation, which can improve generalization under natural distribution shifts on semantic segmentation.
- Extensive experiments demonstrate that our proposed augmentation method ISSA consistently promotes domain generalization performance on driving-scene semantic segmentation across different network architectures, achieving up to 12.4% mIoU improvement, even

with limited diversity in the source data and without access to the target domain.

2. Related Work

Domain Generalization. Domain generalization concerns the generalization ability of neural networks on a target domain that follows a different distribution than the source domain, and is inaccessible at training. Various approaches have been proposed, which employ data augmentation [31, 40, 48, 62, 86], domain alignment [30, 34, 46, 47, 85], meta-learning [4, 44, 45, 83], or ensemble learning [19, 52, 72]. While the majority focuses on image-level tasks, e.g., image classification or person re-identification, a few recent works [12, 41–43] investigate pixel-level prediction tasks such as semantic segmentation. RobustNet [12] proposes an instance selective whitening loss to the instance normalization, aiming to selectively remove information that causes a domain shift while maintaining discriminative features. [41] introduces a memory-guided meta-learning framework to capture co-occurring categorical knowledge across domains. [42, 43] make use of extra data in the wild.

Another line of work explores feature-level augmentation [48, 86]. MixStyle [86] and DSU [48] add perturbation at the normalization layer to simulate domain shifts at test time. However, this perturbation can potentially cause a distortion of the image content, which can be harmful for semantic segmentation (see Sec. 4.2). Moreover, these methods require a careful adaptation to the specific network architecture. In contrast, ISSA performs style mixing on the image-level, thus being model-agnostic, and can be applied as a complement to other methods in order to further increase the generalization performance.

Data Augmentation. Data augmentation techniques can diversify training samples by altering their style, content, or both, thus preventing overfitting and improving generalization. Mixup augmentations [16, 67, 80] linearly interpolate between two training samples and their labels, regularizing both style and content. Despite effectiveness shown on image-level classification tasks, they are not well suited for dense pixel-level prediction tasks. CutMix [79] cuts and pastes a random rectangular region of the input image into another image, thus increasing the content diversity. Geometric transformation, e.g., random scaling and horizontal flipping, can also serve this purpose. In contrast, Hendrycks corruptions [25] only affect the image appearance without modifying the content. Their generated images look artificial, being far from resembling natural data, and thus offer limited help against natural distribution shifts [63].

StyleMix [28] is conceptually closer to our method, which aims to decompose training images into content and style representations and then mix them up to generate more samples. Nonetheless, their AdaIN [32] based style mixing method cannot fulfill the pixel-wise label-preserving re-

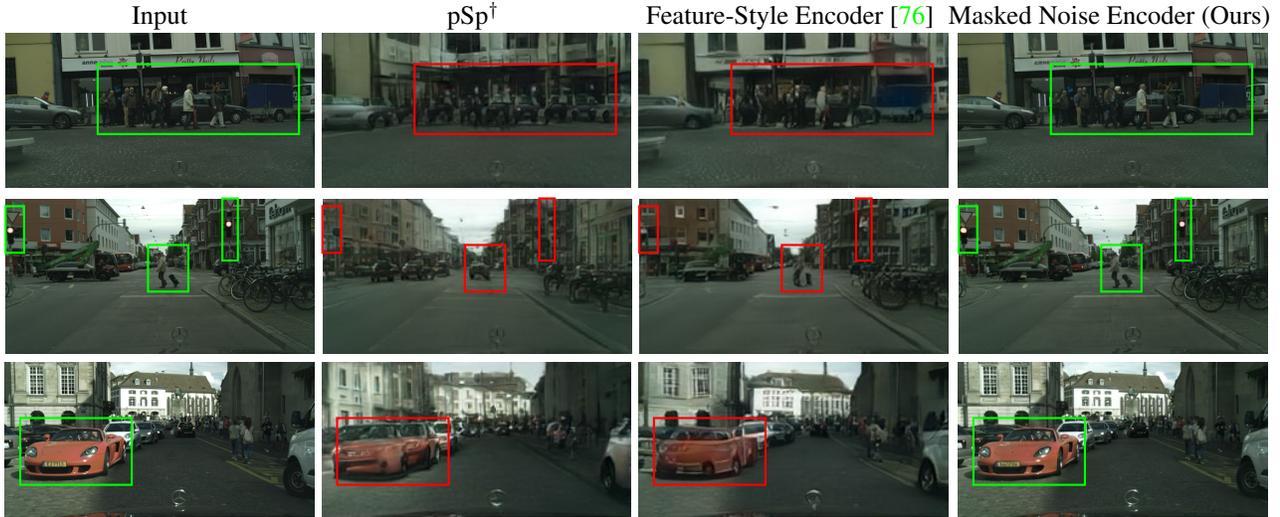


Figure 2. Qualitative results (best view in color and zoom in) of StyleGAN2 inversion methods on Cityscapes, i.e., pSp^\dagger , Feature-Style encoder [76] and our masked noise encoder. Note, pSp^\dagger is an improved version of pSp [57] introduced by us. pSp^\dagger can reconstruct the rough layout of the scene but still struggles to preserve details. The Feature-Style encoder shows a better reconstruction quality, yet it cannot faithfully reconstruct small objects (e.g. pedestrian), and some objects (e.g. the vehicle, bicycle) are rather blurry. Our masked noise encoder has highest image fidelity, preserving finer details in the inverted image. More visual examples, including the original pSp results, can be found in Fig. S.2 in the supp. material.

quirement (see Fig. 7). Our ISSA is also a style-based data augmentation technique. Benefiting from the usage of a state-of-the-art GAN, it can generate natural looking samples, altering only the style of the original images while preserving their content and, thus, enabling the re-use of the ground truth label maps.

GAN Inversion. Showing good results, GAN inversion has been explored for many applications such as face editing [1, 2, 87], image restoration [54], and data augmentation [21, 53]. StyleGANs [37–39] are commonly used for inversion, as they demonstrate high synthesis quality and appealing editing capabilities. Nevertheless, there is a known distortion-editability trade-off [64]. Thus, it is crucial to achieve a curated performance for a specific use case.

GAN inversion approaches can be classified into three groups: optimization based methods [1, 2, 13, 15, 23, 35], encoder based models [5, 57, 64, 71, 76] methods, and hybrid approaches [3, 9, 18, 58]. Optimization methods generally have worse editability and need exhaustive optimization for each input. Thus, in this paper, we use an encoder based method for our style mixing purpose. The representative encoder based work pSp encoder [57] embeds the input image in the extended latent space \mathcal{W}^+ of StyleGAN. The e4e encoder [64] improves editability of pSp while trading off detail preservation. To improve reconstruction quality the Feature-Style encoder [76] further replaces the lower latent code prediction with a feature map prediction. Despite much progress, most prior work only showcases applications on single object-centric datasets, such as FFHQ [38], LSUN [78]. They still fail on more complex scenes, thus

restricting their application in practice. Our masked noise encoder can fulfil both the fidelity and the style mixing capability requirements, rendering itself well-suited for data augmentation for semantic segmentation. To the best of our knowledge, our approach is the first GAN inversion method which can be effectively applied as data augmentation for the semantic segmentation of complex scenes.

3. Method

We introduce our intra-source style augmentation (ISSA) in Sec. 3.1, which relies on GAN inversion that can offer faithful reconstruction and style mixing of images. To enable better style-content disentanglement, we propose a masked noise encoder for GAN inversion in Sec. 3.2. Its detailed training loss is described in Sec. 3.3.

3.1. Intra-Source Style Augmentation (ISSA)

The lack of data diversity and the existence of spurious correlation in the training set often lead to poor domain generalization. To mitigate them, ISSA aims at modifying styles of the training samples while preserving their semantic content. It employs GAN inversion to randomize the style-content combinations in the training set. In doing so, it diversifies the source training set and reduces spurious style-content correlations. Because the content of images is preserved and only the style is changed, the ground truth label maps can be re-used for training, without requiring any further annotation effort.

ISSA is built on top of an encoder-based GAN inversion pipeline given its fast inference. GANs, such as StyleGANs [37–39], have shown the capability of encoding rich

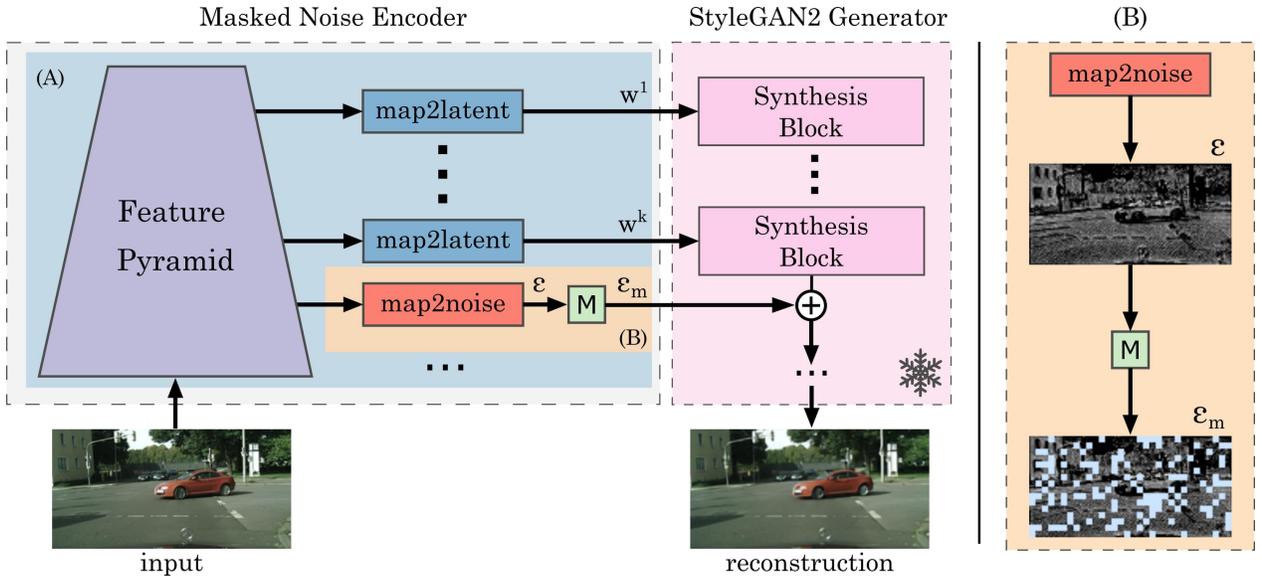


Figure 3. **Method overview.** Our encoder is built on top of the pSp encoder [57], shown in the blue area (A). It maps the input image to the extended latent space \mathcal{W}^+ of the pre-trained StyleGAN2 generator. To promote the reconstruction quality on complex scene-centric dataset, e.g., Cityscapes, our encoder additionally predicts the noise map at an intermediate scale, illustrated in the orange area (B). \boxed{M} stands for random noise masking, regularization for the encoder training. Without it, the noise map overtakes the latent codes in encoding the image style, so that the latter cannot make any perceivable changes on the reconstructed image, thus making style mixing impossible.

semantic and style information in intermediate features and latent spaces. For encoder-based GAN inversion, an encoder is trained to invert an input image back into the latent space of a pre-trained GAN generator. ISSA needs an encoder that can separately encode the style and content information of the input image. With such an encoder, it can synthesize new training samples with new style-content combinations, i.e., it can take the content and style codes from different training samples within the source domain and feed them to the pre-trained generator. Since ISSA modifies only the image style with this encoder, the new synthesized training samples already have a ground truth label map in place.

StyleGAN2 can synthesize scene-centric datasets like Cityscapes [14] and BDD100K [77]. However, existing GAN inversion encoders cannot provide the desired fidelity to enable ISSA to improve domain generalization of semantic segmentation via data augmentation. Loss of fine details or inauthentic reconstruction of small-scale objects would harm the model’s generalization ability. Therefore, we propose a novel encoder design to invert StyleGAN2, termed *masked noise encoder* (see Fig. 3).

3.2. Masked Noise Encoder

We build our encoder upon the pSp encoder [57]. It employs a feature pyramid [50] to extract multi-scale features from a given image, see Fig. 3-(A). We improve over pSp by identifying in which latent space to embed the input image for the high-quality reconstruction of the images with complex street scenes. Further, we propose a novel training scheme to enable the style-content disentanglement of the

encoder, thus improving its style mixing capability.

Extended Latent Space. The StyleGAN2 generator takes the latent code $w \in \mathcal{W}$ generated by an MLP network and randomly sampled additive Gaussian noise maps $\{\epsilon\}$ as inputs for image synthesis. As pointed out in [1], it is suboptimal to embed a real image into the original latent space \mathcal{W} of StyleGAN2, due to the gap between the real and synthetic data distributions. A common practice is to map the input image into the extended latent space \mathcal{W}^+ . The multi-scale features of the pSp feature pyramid are respectively mapped to the latent codes $\{w^k\}$ at the corresponding scales of the StyleGAN2 generator, i.e., map2latent in Fig. 3-(A).

Additive Noise Map. The latent codes $\{w^k\}$ from the extended latent space \mathcal{W}^+ alone are not expressive enough to reconstruct images with diverse semantic layouts such as Cityscapes [14] as shown in Fig. 2-(pSp[†]). The latent codes of StyleGAN2 are one-dimensional vectors that modulate the feature vectors at different spatial positions identically. Therefore, they cannot precisely encode the semantic layout information, which is spatially varying. To address this issue, our encoder additionally predicts the additive noise map ϵ of the StyleGAN2 at an intermediate scale, i.e., map2noise in Fig. 3-(B).

Random Noise Masking. While offering high-quality reconstruction, the additive noise map can be too expressive so that it encodes nearly all perceivable details of the input image. This results in a poor style-content disentanglement and can damage the style mixing capability of the encoder (see Fig. 4). To avoid this undesired effect, we propose to

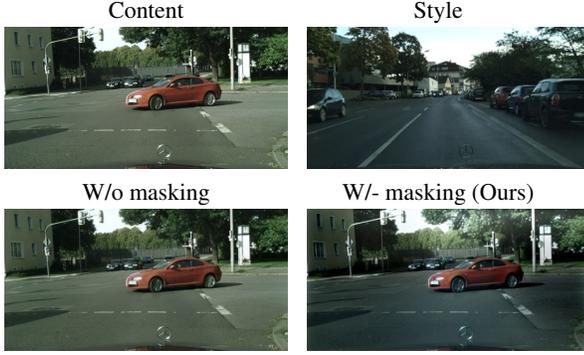


Figure 4. Style mixing effect enabled by random noise masking (best view in color). Despite the good reconstruction quality, the encoder trained without masking cannot change the style of the given Content image. In contrast, the encoder trained with masking can modify it using the style from the given Style image.

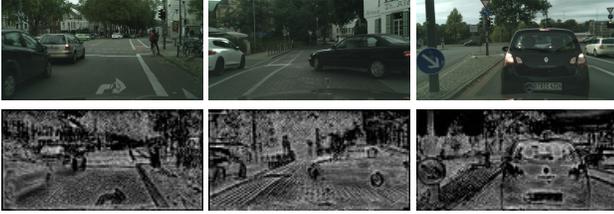


Figure 5. Noise map visualization of our masked noise encoder. The noise map encodes the semantic content of the image.

regularize the noise prediction of the encoder by random masking of the noise map. Note that the random masking as a regularization technique has also been successfully used in reconstruction-based self-supervised learning [24, 74]. In particular, we spatially divide the noise map into non-overlapping $P \times P$ patches, see **[M]** in Fig. 3-(B). Based on a pre-defined ratio ρ , a subset of patches is randomly selected and replaced by patches of unit Gaussian random variables $\epsilon \sim N(0, 1)$ of the same size. $N(0, 1)$ is the prior distribution of the noise map at training the StyleGAN2 generator. We call this encoder *masked noise encoder* as it is trained with random masking to predict the noise map.

The proposed random masking reduces the encoding capacity of the noise map, hence encouraging the encoder to jointly exploit the latent codes $\{w^k\}$ for reconstruction. Fig. 4 visualizes the style mixing effect. The encoder takes the noise map ϵ_c and latent codes $\{w_s^k\}$ from the content image and style image, respectively. Then, they are fed into StyleGAN2 to synthesize a new image, i.e., $G(w_s^k, \epsilon_c)$. If the encoder is not trained with random masking, the new image does not have any perceptible difference with the content image. This means the latent codes $\{w^k\}$ encode negligible information of the image. In contrast, when being trained with masking, the encoder creates a novel image that takes the content and style from two different images. This observation confirms the enabling role of masking for content and style disentanglement, and thus the improved

style mixing capability. The noise map no longer encodes all perceptible information of the image, including style and content. In effect, the latent codes $\{w^k\}$ play a more active role in controlling the style. In Fig. 5, we further visualize the noise map of the masked noise encoder and observe that it captures well the semantic content of the scene.

3.3. Encoder Training Loss

Mathematically, the proposed StyleGAN2 inversion with the masked noised encoder E^M can be formulated as

$$\begin{aligned} \{w^1, \dots, w^K, \epsilon\} &= E^M(x); \\ x^* &= G \circ E^M(x) = G(w^1, \dots, w^K, \epsilon). \end{aligned} \quad (1)$$

The masked noise encoder E^M maps the given image x onto the latent codes $\{w^k\}$ and the noise map ϵ . The StyleGAN2 generator G takes both $\{w^k\}$ and ϵ as the input and generates x^* . Ideally, x^* should be identical to x , i.e., a perfect reconstruction.

When training the masked noise encoder E^M to reconstruct x , the original noise map ϵ is masked before being fed into the pre-trained G

$$\epsilon_M = (1 - M_{noise}) \odot \epsilon + M_{noise} \odot \epsilon, \quad (2)$$

$$\tilde{x} = G(w^1, \dots, w^K, \epsilon_M), \quad (3)$$

where M_{noise} is the random binary mask, \odot indicates the Hadamard product, and \tilde{x} denotes the reconstructed image with the masked noise ϵ_M . The training loss for the encoder is given as

$$\mathcal{L} = \mathcal{L}_{mse} + \lambda_1 \mathcal{L}_{lips} + \lambda_2 \mathcal{L}_{adv} + \lambda_3 \mathcal{L}_{reg}, \quad (4)$$

where $\{\lambda_i\}$ are weighting factors. The first three terms are the pixel-wise MSE loss, learned perceptual image patch similarity (LPIPS) [81] loss and adversarial loss [22],

$$\mathcal{L}_{mse} = \|(1 - M_{img}) \odot (x - \tilde{x})\|_2, \quad (5)$$

$$\mathcal{L}_{lips} = \|(1 - M_{feat}) \odot (\text{VGG}(x) - \text{VGG}(\tilde{x}))\|_2, \quad (6)$$

$$\mathcal{L}_{adv} = -\log D(G(E^M(x))). \quad (7)$$

which are the common reconstruction losses for encoder training [57, 87]. Note that masking removes the information of the given image x at certain spatial positions, the reconstruction requirement on these positions should then be relaxed. M_{img} and M_{feat} are obtained by up- and down-sampling the noise mask M_{noise} to the image size and the feature size of the VGG-based feature extractor. The adversarial loss is obtained by formulating the encoder training as an adversarial game with a discriminator D that is trained to distinguish between reconstructed and real images.

The last regularization term is defined as

$$\mathcal{L}_{reg} = \|\epsilon\|_1 + \|E_w^M(G(w_{gt}, \epsilon)) - w_{gt}\|_2. \quad (8)$$



Figure 6. Visual examples of style mixing on BDD100K (best view in color) enabled by our masked noise encoder. By combining the latent codes $\{w_s^k\}$ of I_s and the noise map ϵ_c of I_c , the synthesized images $G(w_s^k, \epsilon_c)$ preserve the content of I_c with a new style resembling I_s .

The L1 norm helps to induce sparse noise prediction. It is complementary to random masking, reducing the capacity of the noise map. The second term is obtained by using the ground truth latent codes w_{gt} of synthesized images $G(w_{gt}, \epsilon)$ to train the latent code prediction $E_w^M(\cdot)$ [76]. It guides the encoder to stay close to the original latent space of the generator, speeding up the convergence.

4. Experiments

Sec. 4.1 and Sec. 4.2 respectively report our experiments on the StyleGAN2 inversion and domain generalization of semantic segmentation.

Datasets. We conduct extensive experiments using the following driving scene datasets: Cityscapes (CS) [14], BDD100K (BDD) [77], ACDC [60] and Dark Zürich (DarkZ) [59]. Cityscapes is collected from different cities primarily in Germany, under good/medium weather conditions during daytime. BDD100K is a driving-scene dataset collected in the US, representing a geographic location shift from Cityscapes. Besides, it also includes more diverse scenes (e.g., city streets, residential areas, and highways) and different weather conditions captured at different times of the day. Both ACDC and Dark Zürich are collected in Switzerland. ACDC contains four adverse weather conditions (rain, fog, snow, night) and Dark Zürich contains night scenes. The default setting is to use Cityscapes as the source training data, whereas the validation sets of the other datasets represent unseen target domains with different types of natural shifts, i.e., used only for testing. In the supp. material, we also report the numbers where BDD100K is used as the source set and the remaining datasets are treated as unseen domains. In both cases, we consider a *single source domain* for training.

Training details. We experiment with two image resolu-

tions: 128×256 and 256×512 . The StyleGAN2 [37] model is first trained to *unconditionally* synthesize images and then fixed during the encoder training. To invert the pre-trained StyleGAN2 generator, the masked noise encoder predicts both latent codes in the extended \mathcal{W}^+ space and the additive noise map. In accordance with the StyleGAN2 generator, \mathcal{W}^+ space consists of 14 and 16 latent code vectors for the input resolution 128×256 and 256×512 , respectively. The additive noise map is always at the intermediate feature space with one fourth of the input resolution. We use the same encoder architecture, optimizer, and learning rate scheduling as pSp [57]. Our encoder is trained with the loss function defined in Eq. (4) with $\lambda_1 = 10$ and $\lambda_2 = \lambda_3 = 0.1$. For our random noise masking, we use a patch size P of 4 with a masking ratio $\rho = 25\%$. A detailed ablation study on the noise map and a computational complexity analysis of the encoder can be found in S.1.

We use the trained masked noise encoder to perform ISSA as described in Sec. 3.1. We experiment with several architectures for semantic segmentation, i.e., HRNet [69], SegFormer [73], and DeepLab v2/v3+ [10, 11]. The baseline segmentation models are trained with their default configurations and using the standard augmentation, i.e., random scaling and horizontal flipping.

4.1. Masked Noise Encoder

Reconstruction quality. Table 1 shows that our masked noise encoder considerably outperforms two strong StyleGAN2 inversion baselines pSp [57] and Feature-Style encoder [76] in all three evaluation metrics. The achieved low values of MSE, LPIPS [81] and FID [27] indicate its high-quality reconstruction. Both the masked noise encoder and the Feature-Style encoder adopt the adversarial loss \mathcal{L}_{adv} and regularization using synthesized images with ground truth latent codes w_{gt} . Therefore, we also add them

Method	MSE ↓	LPIPS ↓	FID ↓
pSp [57]	0.078	0.348	130.62
pSp [†] [57]	0.049	0.339	14.60
Feature-Style [76]	0.025	0.220	7.14
Ours	0.011	0.124	3.94

Table 1. Reconstruction quality on Cityscapes at the resolution 128×256 . MSE, LPIPS [81] and FID [27] respectively measure the pixel-wise reconstruction difference, perceptual difference, and distribution difference between the real and reconstructed images. The proposed masked noise encoder (Ours) consistently outperforms pSp, pSp[†] and the feature-style encoder. Note, pSp[†] is introduced by us, by training pSp with an additional discriminator and incorporating synthesized images for better initialization.

to train pSp and note this version as pSp[†]. While pSp[†] improves over pSp in MSE and FID, it still underperforms compared to the others. This confirms that inverting into the extended latent space \mathcal{W}^+ only allows limited reconstruction quality on Cityscapes. The Feature-Style encoder [76] replaces the prediction of the low level latent codes with feature prediction, which results in better reconstruction without severely harming style editability. However, its reconstruction on Cityscapes is still not satisfying and underperforms to our masked noise encoder. As noted in [76], the feature size of the Feature-Style encoder is restricted. Using a larger feature map to improve reconstruction quality can only be done as a replacement of more latent code predictions. Consequently, it largely reduces the expressiveness of the latent embedding and leads to extremely poor editability, being no longer suitable for downstream applications, e.g., style mixing data augmentation.

The visual comparison across pSp[†], the Feature-Style encoder and our masked noise encoder is shown in Fig. 2 and is aligned with the quantitative results in Table 1. pSp[†] has overall poor reconstruction quality. The Feature-Style encoder cannot faithfully reconstruct small objects and restore fine details. In comparison, our masked noise encoder offers high-quality reconstruction, preserving the semantic layout and fine details of each class. Having a high-quality reconstruction is an important requirement for using the encoder for data augmentation. Unfortunately, neither pSp[†] nor the Feature-Style encoder achieve satisfactory reconstruction quality. For instance, they both fail at capturing the red traffic light in Fig. 2. Using such images for data augmentation can confuse the semantic segmentation model, leading to performance degradation.

Ablation on the masking effect. In Fig. 4 and Fig. 6, we visually observe that random masking offers a stronger perceivable style mixing effect compared to the model trained without masking. Next, we test the effect of masking on improving the domain generalization for the semantic segmentation task. In particular, we employ the encoder that is trained with and without masking to perform

Method	CS	ACDC	BDD	DarkZ
Baseline	70.47	41.48	45.66	15.25
ISSA w/o masking	69.68	44.63	46.45	17.36
ISSA w/- masking	69.48	47.43	47.87	26.10

Table 2. The effect of random noise masking on improving domain generalization via ISSA. We report the mean Intersection over Union (mIoU) of HRNet [69] trained on Cityscapes at the resolution 256×512 . BDD100K (BDD), ACDC, and Dark Zürich (DarkZ) represent different domain shifts from Cityscapes.



Figure 7. Comparison of StyleMix [28] and ISSA. StyleMix has rather low fidelity, while ISSA can preserve more details.

ISSA. In Table 2, while slightly degrading the source domain performance of the baseline model on Cityscapes, ISSA improves the domain generalization performance on BDD100K, ACDC and Dark Zürich. As ISSA with masked noise encoder is more effective at diversifying the training set and reducing the style-content correlation, it achieves more pronounced gains in Table 2, e.g., more than 10% improvement in mIoU from Cityscapes to Dark Zürich.

4.2. Domain Generalization

Comparison with data augmentation methods. Table 3 reports the mIoU scores of Cityscapes to ACDC domain generalization using two semantic segmentation models, i.e., HRNet [69] and SegFormer [73]. ISSA is compared with three representative data augmentations methods, i.e., CutMix [79], Hendrycks’s weather corruptions [25], and StyleMix [28]. Remarkably, our ISSA is the top performing method, consistently improving mIoU in both models and across all four different scenarios of ACDC, i.e., rain, fog, snow and night. Compared to HRNet, SegFormer is more robust against the considered domain shifts.

In contrast to the others, CutMix mixes up the content rather than the style. It improves the in-distribution performance on Cityscapes, but this gain does not extend to domain generalization. Hendrycks’s weather corruptions can be seen as the synthetic version of Cityscapes under the rain, fog, and snow weather conditions. While already mimicking ACDC at training, it can still degrade ACDC-snow by

Method	HRNet [69]						SegFormer [73]					
	CS	Rain	Fog	Snow	Night	Avg.	CS	Rain	Fog	Snow	Night	Avg.
Baseline	70.47	44.15	58.68	44.20	18.90	41.48	67.90	50.22	60.52	48.86	28.56	47.04
CutMix [79]	72.68	<u>42.48</u>	<u>58.63</u>	44.50	<u>17.07</u>	<u>40.67</u>	69.23	<u>49.53</u>	61.58	<u>47.42</u>	<u>27.77</u>	<u>46.57</u>
Weather [25]	69.25	50.78	60.82	<u>38.34</u>	22.82	43.19	67.41	54.02	64.74	49.57	28.50	49.21
StyleMix [28]	57.40	<u>40.59</u>	<u>49.11</u>	<u>39.14</u>	19.34	<u>37.04</u>	65.30	53.54	63.86	49.98	28.93	49.08
ISSA (Ours)	70.30	<u>50.62</u>	66.09	53.30	30.18	50.05	67.52	55.91	67.46	53.19	33.23	52.45
Oracle	70.29	65.67	75.22	72.34	50.39	65.90	68.24	63.67	74.10	67.97	48.79	63.56

Table 3. Comparison of data augmentation for improving domain generalization, i.e., from Cityscapes (train) to ACDC (unseen). The mean Intersection over Union (mIoU) is reported on Cityscapes (CS), four individual scenarios of ACDC (Rain, Fog, Snow and Night) and the whole ACDC (Avg.). Oracle indicates the supervised training on both Cityscapes and ACDC, serving as an upper bound on ACDC for the other methods. Note, it is not supposed to be an upper bound on Cityscapes. Underline denotes worse results than the baseline on ACDC. ISSA performs the best and consistently improves the mIoU in all four scenarios of ACDC using both HRNet and SegFormer.

Method	CS	ACDC	BDD	DarkZ
Baseline [10]	61.73	30.86	34.30	11.62
MixStyle [86]	59.01	36.97	36.27	9.38
DSU [48]	59.59	38.31	35.53	12.29
ISSA (Ours)	62.20	43.21	42.60	21.56
MixStyle [86] + ISSA	60.17	41.81	42.17	20.56
DSU [48] + ISSA	60.20	43.31	42.24	24.63

Table 4. Comparison with feature-level augmentation methods on domain generalization performance of Cityscapes as the source. Following DSU [48], we conduct experiments using DeepLab v2 [10] as the baseline for fair comparison.

more than 5.8% in mIoU using HRNet. More results on the other corruption types can be found in the supp. material. StyleMix [28] also seeks to mix up styles. However, it does not work well for scene-centric datasets, such as Cityscapes. Its poor synthetic image quality (see Fig. 7) leads to the performance drop over the HRNet baseline in many cases, e.g., on Cityscapes to ACDC-fog from 58.68% to 49.11% mIoU.

Comparison with domain generalization techniques. We further compare ISSA with two advanced feature space style mixing methods designed to improve domain generalization performance: MixStyle [86] and DSU [48]. Both extract the style information at certain normalization layers of CNNs. MixStyle [86] mixes up styles by linearly interpolating the feature statistics, i.e., mean and variance, of different images, while DSU [48] models the feature statistics as a distribution and randomly draws samples from it.

We adopt the experimental setting of DSU with default hyperparameters, using DeepLab v2 [10] segmentation network with ResNet101 backbone. Table 4 shows that ISSA outperforms both MixStyle and DSU by a large margin. We also observe that there is a slight performance drop on the source domain (CS) when applying DSU and MixStyle. As they operate at the feature-level, there is no guarantee that the semantic content stays unchanged after the random perturbation of feature statistics. Thus, the changes in feature

Method	CS	ACDC	BDD	DarkZ
Baseline	69.01	44.23	43.27	16.03
RobustNet [12]	69.47	47.25	46.94	20.11
+ ISSA	69.45	47.55	48.44	23.09

Table 5. Combination of ISSA and RobustNet [12]. We adopt the experimental setting of RobustNet and use DeepLab v3+ [11] as the baseline. Our ISSA is complementary to RobustNet and further improves its generalization performance.

statistics might negatively affect the performance, as also indicated in [48]. Note that, in contrast, ISSA operates on the image space. Combining ISSA with MixStyle and DSU leads to a strong boost in performance of these methods.

Being model-agnostic, ISSA can be combined with other networks designed specifically for the domain generalization of semantic segmentation. To showcase its complementary nature, we add ISSA on top of RobustNet [12], which proposed a novel instance whitening loss to selectively remove domain-specific style information. Although color transformation has already been used for augmentation in RobustNet, ISSA can introduce more natural style shifts, thus helping to further remove style-specific biases. Table 5 verifies the effectiveness of ISSA. It brings extra gains for RobustNet, especially in the challenging day to night scenario, i.e., Cityscapes to Dark Zürich, boosting the performance from 20.11% to 23.09% in mIoU.

5. Conclusion

In this paper, we propose a GAN inversion based data augmentation method ISSA for learning domain generalized semantic segmentation using restricted training data from a single source domain. The key enabler for ISSA is the masked noise encoder, which is capable of preserving fine-grained content details and allows style mixing between images without affecting the semantic content. Extensive experimental results demonstrate the effectiveness of ISSA on domain generalization across different datasets and network architectures.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*, 2019. 3, 4
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *CVPR*, 2020. 2, 3
- [3] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *CVPR*, 2022. 2, 3
- [4] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, 2018. 1, 2
- [5] Christian Bartz, Joseph Bethge, Haojin Yang, and Christoph Meinel. One model to reconstruct them all: A novel way to use the stochastic noise in StyleGAN. In *BMVC*, 2021. 3
- [6] Yanrui Bin, Xuan Cao, Xinya Chen, Yanhao Ge, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, Changxin Gao, and Nong Sang. Adversarial semantic data augmentation for human pose estimation. In *ECCV*, 2020. 1
- [7] Simon Burton, Lydia Gauerhof, and Christian Heinzemann. Making the case for safety of machine learning in highly automated driving. In *SAFECOMP*, 2017. 1
- [8] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018. 1
- [9] Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *CVPR*, 2021. 3
- [10] 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. *TPAMI*, 2017. 6, 8, 15
- [11] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6, 8
- [12] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T. Kim, Seungryong Kim, and Jaegul Choo. RobustNet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *CVPR*, 2021. 2, 8
- [13] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In *CVPR*, 2020. 3
- [14] 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 *CVPR*, 2016. 1, 4, 6
- [15] Antonia Creswell and Anil Anthony Bharath. Inverting the generator of a generative adversarial network. *TNNLS*, 2018. 3
- [16] Ali Dabouei, Sobhan Soleymani, Fariborz Taherkhani, and Nasser M Nasrabadi. Supermix: Supervising the mixing data augmentation. In *CVPR*, 2021. 2
- [17] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint*, 2017. 1
- [18] Tan M Dinh, Anh Tuan Tran, Rang Nguyen, and Binh-Son Hua. Hyperinverter: Improving stylegan inversion via hypernetwork. In *CVPR*, 2022. 2, 3
- [19] Antonio D’Innocente and Barbara Caputo. Domain generalization with domain-specific aggregation modules. In *GCPR*, 2018. 2
- [20] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 1
- [21] Mayank Golhar, Taylor L Bobrow, Saowanee Ngamruengphong, and Nicholas J Durr. GAN Inversion for Data Augmentation to Improve Colonoscopy Lesion Classification. *arXiv preprint*, 2022. 3
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 5
- [23] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *CVPR*, 2020. 3
- [24] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 5
- [25] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2018. 2, 7, 8, 13, 14, 15
- [26] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2019. 1
- [27] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 6, 7
- [28] Minui Hong, Jinwoo Choi, and Gunhee Kim. StyleMix: Separating content and style for enhanced data augmentation. In *CVPR*, 2021. 1, 2, 7, 8, 13, 14, 15
- [29] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *CVPR*, 2022. 15
- [30] Shoubo Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. Domain generalization via multidomain discriminant analysis. In *UAI*, 2020. 1, 2
- [31] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *CVPR*, 2021. 2
- [32] Xun Huang and Serge Belongie. Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. In *ICCV*, 2017. 2
- [33] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Deceive d: Adaptive pseudo augmentation for gan training with limited data. *NeurIPS*, 2021. 1
- [34] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Feature alignment and restoration for domain generalization and adaptation. *arXiv preprint*, 2020. 1, 2
- [35] Kyoungkook Kang, Seongtae Kim, and Sunghyun Cho. Gan inversion for out-of-range images with geometric transformations. In *ICCV*, 2021. 3
- [36] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*, 2018. 2
- [37] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine,

- Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, 2020. 1, 3, 6
- [38] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 3
- [39] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2, 3
- [40] Rawal Khirodkar, Donghyun Yoo, and Kris Kitani. Domain randomization for scene-specific car detection and pose estimation. In *WACV*, 2019. 2
- [41] Jin Kim, Jiyoung Lee, Jungin Park, Dongbo Min, and Kwanghoon Sohn. Pin the memory: Learning to generalize semantic segmentation. In *CVPR*, 2022. 2
- [42] Namyup Kim, Taeyoung Son, Cuiling Lan, Wenjun Zeng, and Suha Kwak. WEDGE: Web-image assisted domain generalization for semantic segmentation. *arXiv preprint*, 2021. 2
- [43] Suhyeon Lee, Hongje Seong, Seongwon Lee, and Euntai Kim. WildNet: Learning domain generalized semantic segmentation from the wild. In *CVPR*, 2022. 2
- [44] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 1, 2
- [45] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy Hospedales. Episodic training for domain generalization. In *ICCV*, 2019. 2
- [46] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018. 1, 2
- [47] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex Kot. Domain generalization for medical imaging classification with linear-dependency regularization. In *NeurIPS*, 2020. 1, 2
- [48] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and LINGYU DUAN. Uncertainty Modeling for Out-of-Distribution Generalization. In *ICLR*, 2022. 2, 8
- [49] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *CVPR*, 2019. 15
- [50] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 4
- [51] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, 2019. 15
- [52] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Best sources forward: Domain generalization through source-specific nets. In *ICIP*, 2018. 2
- [53] Dinh Tan Nguyen, Cao Trung Tran, Trung Thanh Nguyen, Cao Bao Hoang, Van Phu Luu, Ba Ngoc Nguyen, and Pou Ian Cheong. Data augmentation for small face datasets and face verification by generative adversarial networks inversion. In *KSE*, 2021. 3
- [54] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *TPAMI*, 2021. 3
- [55] Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S Feris, and Dimitris Metaxas. Jointly optimize data augmentation and network training: Adversarial data augmentation in human pose estimation. In *CVPR*, 2018. 1
- [56] 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. 15
- [57] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *CVPR*, 2021. 2, 3, 4, 5, 6, 7, 12
- [58] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *arXiv preprint*, 2021. 2, 3, 13
- [59] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *ICCV*, 2019. 6
- [60] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Accd: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, 2021. 1, 6, 15
- [61] Sina Shafaei, Stefan Kugele, Mohd Hafeez Osman, and Alois Knoll. Uncertainty in machine learning: A safety perspective on autonomous driving. In *SAFECOMP*, 2018. 1
- [62] Nathan Somavarapu, Chih-Yao Ma, and Zsolt Kira. Frustratingly simple domain generalization via image stylization. *arXiv preprint*, 2020. 2
- [63] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. In *NeurIPS*, 2020. 1, 2
- [64] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *TOG*, 2021. 3
- [65] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018. 15
- [66] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *ICCV*, 2019. 15
- [67] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, 2019. 1, 2
- [68] Jiahang Wang, Sheng Jin, Wentao Liu, Weizhong Liu, Chen Qian, and Ping Luo. When human pose estimation meets robustness: Adversarial algorithms and benchmarks. In *CVPR*, 2021. 1
- [69] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2020. 1, 6, 7, 8, 13, 14, 15
- [70] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jin-

- jun Xiong, Wen-mei Hwu, Thomas S Huang, and Honghui Shi. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *CVPR*, 2020. 15
- [71] Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Weiming Zhang, Lu Yuan, Gang Hua, and Nenghai Yu. E2Style: Improve the efficiency and effectiveness of stylegan inversion. *TIP*, 2022. 3
- [72] Guile Wu and Shaogang Gong. Collaborative optimization and aggregation for decentralized domain generalization and adaptation. In *ICCV*, 2021. 2
- [73] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. 6, 7, 8, 13, 14, 15
- [74] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A simple framework for masked image modeling. In *CVPR*, 2022. 5
- [75] Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *CVPR*, 2020. 15
- [76] Xu Yao, Alasdair Newson, Yann Gousseau, and Pierre Hellier. Feature-Style Encoder for Style-Based GAN Inversion. *arXiv preprint*, 2022. 2, 3, 6, 7, 12
- [77] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020. 4, 6
- [78] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint*, 2015. 2, 3
- [79] 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 *ICCV*, 2019. 1, 2, 7, 8, 13, 14, 15
- [80] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond Empirical Risk Minimization. In *ICLR*, 2018. 1, 2
- [81] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 6, 7
- [82] Yuxiao Zhang, Alexander Carballo, Hanting Yang, and Kazuya Takeda. Autonomous Driving in Adverse Weather Conditions: A Survey. *arXiv preprint*, 2021. 1
- [83] Yuyang Zhao, Zhun Zhong, Fengxiang Yang, Zhiming Luo, Yaojin Lin, Shaozi Li, and Nicu Sebe. Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *CVPR*, 2021. 2
- [84] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *IJCV*, 2021. 15
- [85] Fan Zhou, Zhuqing Jiang, Changjian Shui, Boyu Wang, and Brahim Chaib-draa. Domain generalization with optimal transport and metric learning. *arXiv preprint*, 2020. 1, 2
- [86] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *ICLR*, 2021. 1, 2, 8
- [87] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020. 3, 5
- [88] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, 2019. 15