

FSDR: Frequency Space Domain Randomization for Domain Generalization

Jiaxing Huang, Dayan Guan, Aoran Xiao, Shijian Lu*

School of Computer Science Engineering, Nanyang Technological University

{Jiaxing.Huang, Dayan.Guan, Aoran.Xiao, Shijian.Lu}@ntu.edu.sg

Abstract

Domain generalization aims to learn a generalizable model from a ‘known’ source domain for various ‘unknown’ target domains. It has been studied widely by domain randomization that transfers source images to different styles in spatial space for learning domain-agnostic features. However, most existing randomization methods use GANs that often lack of controls and even alter semantic structures of images undesirably. Inspired by the idea of JPEG that converts spatial images into multiple frequency components (FCs), we propose Frequency Space Domain Randomization (FSDR) that randomizes images in frequency space by keeping domain-invariant FCs (DIFs) and randomizing domain-variant FCs (DVF) only. FSDR has two unique features: 1) it decomposes images into DIFs and DVFs which allows explicit access and manipulation of them and more controllable randomization; 2) it has minimal effects on semantic structures of images and domain-invariant features. We examined domain variance and invariance property of FCs statistically and designed a network that can identify and fuse DIFs and DVFs dynamically through iterative learning. Extensive experiments over multiple domain generalizable segmentation tasks show that FSDR achieves superior segmentation and its performance is even on par with domain adaptation methods that access target data in training.

1. Introduction

Semantic segmentation has been a longstanding challenge in computer vision research, which aims to assign a class label to each and every pixel of an image. Deep learning based methods [4, 38, 79, 84] have achieved great successes at the price of large-scale densely-annotated training data [8] that are usually prohibitively expensive and time-consuming to collect. One way of circumventing this constraint is to employ synthetic images with automatically generated labels [52, 51] in network training. However,

*Corresponding author.

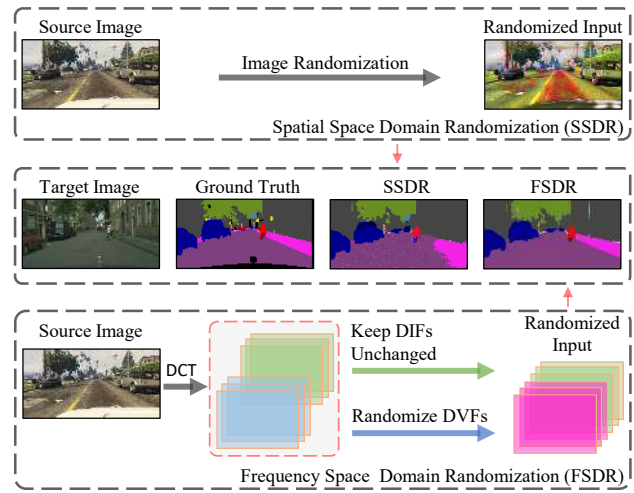


Figure 1. Our proposed frequency space domain randomization (FSDR) converts images into multiple frequency components (FCs) with discrete cosine transform (DCT) and identifies domain variant FCs (DVFs) and domain invariant FCs (DIFs). Such explicit isolation allows it to randomize DVFs while keeping DIFs unchanged in training which often leads to more generalizable models. Traditional spatial space domain randomization (SSDR) tends to produce sub-optimal segmentation as it randomizes images as a whole without isolating and preserving domain invariant features. Green, blue and red boxes denotes DIFs, DVFs and randomized DVFs, respectively (best viewed in color).

such models usually undergo a drastic performance drop while applied to real-world images [78] due to the domain bias and shift [55, 41, 67, 71, 56].

Unsupervised domain adaptation (UDA) has been studied widely to tackle the domain mismatch problem by learning domain invariant/aligned features from a labelled source domain and an unlabelled target domain [22, 29, 30, 58, 66, 69, 40, 72, 5, 6]. However, its training requires target-domain data which could be hard to collect during the training stage such as autonomous driving at various cities, robot exploration of various new environments, etc. Additionally, it is not scalable as it requires network re-training or fine-tuning for each new target domain. Domain generalization has attracted increasing attention as it learns domain

invariant features without requiring target-domain data in training [77, 44, 11, 36, 33, 34]. One widely adopted generalization strategy is domain randomization (DR) that learns domain-agnostic features by randomizing or stylizing source-domain images via adversarial perturbation, generative adversarial networks (GANs), etc. [59, 70, 77, 50]. However, most existing DR methods randomize the whole spectrum of images in the spatial space which tends to modify domain invariant features undesirably.

We propose an innovative frequency space domain randomization (FSDR) technique that transforms images into frequency space and performs domain generalization by identifying and randomizing domain-variant frequency components (DVF) while keeping domain-invariant frequency components (DIF) unchanged. FSDR thus overcomes the constraints of most existing domain randomization methods which work over the full spectrum of images in the spatial space and tend to modify domain-invariant features undesirably as illustrated in Fig. 1. We explored two different approaches for domain randomization in the frequency space. The first is spectrum analysis based FSDR (FSDR-SA) that identifies DIFs and DVFs through empirical studies. The second is spectrum learning based FSDR (FSDR-SL) that identifies DIFs and DVFs through dynamic and iterative learning processes. Extensive experiments show that FSDR improves the model generalization greatly. Additionally, FSDR is complementary to spatial-space domain generalization and the combination of the two improves model generalization consistently.

The contributions of this work can be summarized in three aspects. *First*, we propose an innovative frequency space domain randomization technique that transforms images into frequency space and achieves domain randomization by changing DVFs only while keeping DIFs unchanged. *Second*, we design two randomization approaches in the frequency space that identify DVFs and DIFs effectively through empirical experiments and dynamic learning, respectively. *Third*, extensive experiments over multiple domain generalization tasks show that our proposed frequency space domain randomization technique achieves superior semantic segmentation consistently.

2. Related Works

Domain Generalization (DG) aims to learn a generalizable model from ‘known’ source domains for various ‘unknown’ target domains [44, 11]. Most existing DG methods can be broadly categorized into multi-source DG [44, 13, 12, 34, 16, 33, 43, 42, 59, 36, 2, 10, 35] and single-source DG [64, 70, 77, 50] both of which strive to learn domain-agnostic features by either learning a domain-invariant feature space [44, 13, 12, 16, 34, 59, 36, 10, 2, 82] or aggregating domain specific modules [33, 43, 42, 35]. Specifically, multi-source DG handles generalization by joint su-

pervised and unsupervised learning [2], domain perturbation [59], adversarial feature learning [36], meta-learning [34], new domain discovery [13] and network duplication [33, 35]. Single-source DG handles generalization by domain randomization that augment data [64, 70] or domains [77, 50]. Our work belongs to single-source DG, aiming to tackle a more challenging scenario in domain generalization when only one single source domain is available in training.

Domain Randomization (DR) is the common strategy in domain generalization [64, 53, 49, 63, 65, 59, 70, 77, 50] especially when a single source domain is available [64, 70, 77, 50]. Existing methods aim to enrich the variation of synthetic images in a source domain by gradient-based domain perturbation [59], adversarial data augmentation [70], GANs-based domain augmentation [77] or adversarial domain augmentation [50]. However, these methods randomize the whole spectrum of images in the spatial space which could modify domain invariant features undesirably. Our method transforms images into frequency space and divides images into DIFs and DVFs explicitly. It randomizes images by modifying DVFs without changing DIFs which has minimal effects over (domain-invariant) semantic structures of images.

Domain Adaptation is closely relevant to domain generalization but it exploits (unlabelled) target-domain data during the training process. The existing domain adaptation methods can be broadly classified into three categories. The first is *adversarial training* based which employs adversarial learning to align source and target distributions in the feature, output or latent space [22, 39, 69, 40, 66, 7, 81, 54, 55, 72, 1, 68, 26, 18, 25, 31]. The second is *image translation* based which translates images from source domains to target domains to mitigate domain gaps [21, 58, 37, 80, 23, 32, 76, 75]. The third is *self-training* based which employs “pseudo label” to guide unsupervised learning over unlabelled target-domain data [88, 57, 83, 19, 86, 17, 24].

3. Method

This section presents our Frequency Space Domain Randomization (FSDR) technique, which consists of three major subsections on spectrum analysis that identifies DIFs and DVFs statically, spectrum analysis in FSDR that studies how DIFs and DVFs work for domain randomization, and spectrum learning in FSDR that shows how domain randomization can be achieved via spectrum learning.

3.1. Problem Definition

We focus on the problem of unsupervised domain generalization (UDG) in semantic segmentation. Given source-domain data $X_s \subset \mathbb{R}^{H \times W \times 3}$ with C-class pixel-level segmentation labels $Y_s \subset (1, C)^{H \times W}$, our goal is to learn a semantic segmentation model G that well performs on un-

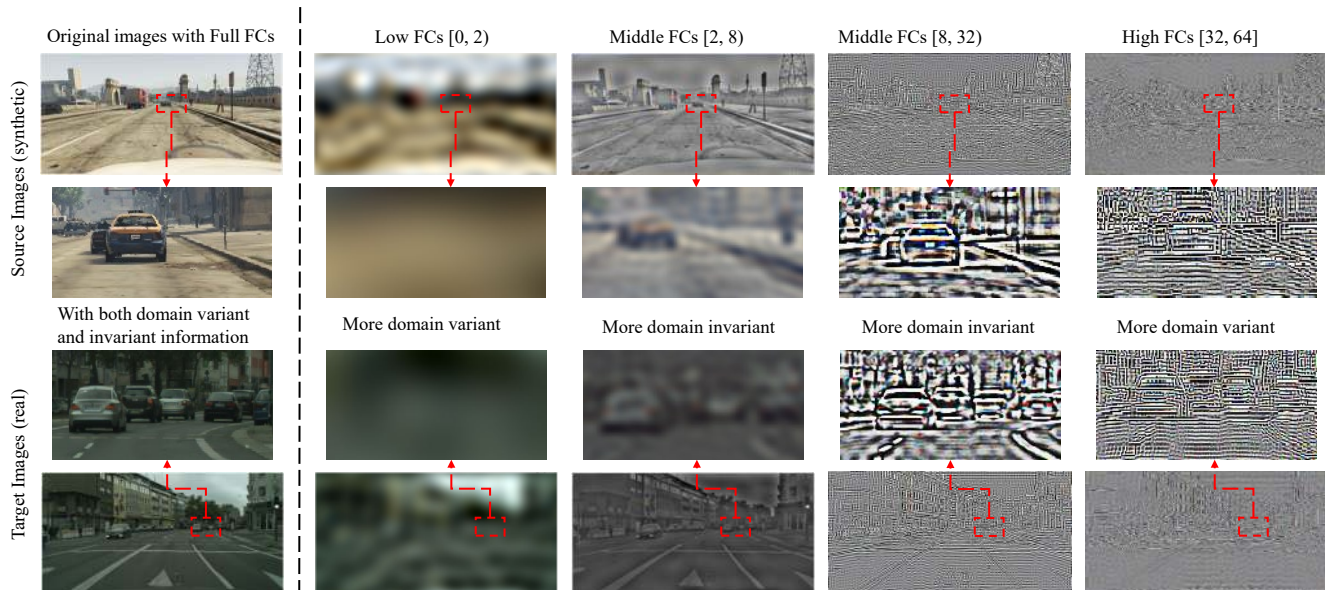


Figure 2. Visualization of spectrum decomposition of synthetic and real images: Column 1 shows *Source Images (synthetic)* and *Target Images (real)* that contain both domain variant and domain invariant information. Columns 2 and 4 show low-pass (i.e., [0, 2]) and high-pass (i.e., [32, 64]) filtered images that largely capture domain-variant information such as colors and styles. Columns 3 and 4 show middle-pass (i.e., [2, 8] and [8, 32]) filtered images that largely capture domain invariant information such as structures and shapes. Images in Rows 2 and 3 are close-up views of highlighted regions in images in Rows 1 and 2.

seen target-domain data X_t . The baseline model is trained with the original source-domain data only:

$$\mathcal{L}_{orig} = l(G(x_s), y_s), \quad (1)$$

where $l(\cdot)$ denotes the standard cross entropy loss.

3.2. Spectrum Analysis

This subsection describes spectrum analysis for identifying DIFs and DVFs. For each source-domain image $x_s \in \mathbb{R}^{H \times W \times 3}$, we first convert it into frequency space with Discrete Cosine Transform (DCT) and then decompose the converted signals into 64 FCs $f_s \in \mathbb{R}^{H \times W \times 3 \times 64}$ with a band-pass filter \mathcal{B}_p :

$$f_s = \mathcal{B}_p(\mathcal{D}(x_s)) = \{f_s(0), f_s(1), f_s(3), \dots, f_s(63)\}, \quad (2)$$

where \mathcal{D} denotes DCT, $f_s(i) \in \mathbb{R}^{H \times W \times 3}$, $i = \{0, 1, \dots, 63\}$ denote the frequency-wise representation, and the definition of $\mathcal{B}_p(\cdot)$ is in appendix.

We identify DIFs and DVFs in f_s by a set of control experiments as shown in Table 1. For each source image, we first filter out FCs with indexes between certain lower/upper threshold (under ‘Rejected bands’ in Table 1) with a band reject filter \mathcal{B}_r and then train models with remaining FCs. The band reject filter $f'_s = \mathcal{B}_r(f_s; I)$ can be defined by:

$$f'_s(i) = \begin{cases} 0, & \text{if } I(i) = 0, \\ f_s(i), & \text{otherwise,} \end{cases} \quad (3)$$

Table 1. We analyze and identify domain variant and invariant frequency components (FCs) by training models with certain FCs on source domain (synthetic), and testing with target domain images (real) in the classification task. The ‘source acc.’ means the test accuracy on the source SYNTHIA dataset while the ‘target acc.’ represents the test accuracy on the target ImageNet dataset.

Band-reject Spectrum analysis		
Rejected band	Source acc.	Target acc.
Null (with full FCs)	95.5	65.2
[0, 1]	95.1	68.6
[1, 2]	95.3	67.1
[2, 4]	95.4	62.3
[4, 8]	95.4	62.7
[8, 16]	95.5	64.6
[16, 32]	95.6	64.9
[32, 64]	95.9	67.4

where I is a 64-dimensional binary mask vector whose values are 1 for the preserved components and 0 for the discarded.

We then apply the trained model to the target images to examine the domain invariance and generalization of the filtered source-domain FCs. Specifically, improved (or degraded) performance over the target data means that the removed FCs are domain variant (or invariant) and removing them prevents learning domain variant (or invariant) features and improves (or degrades) generalization. With such

spectrum analysis experiments with different filter masks I , DIFs and DVFs $f = \{f_{var}, f_{invar}\}$ can be identified and recorded by a binary mask vector I^{SA} as follows:

$$\begin{aligned} x_s^f &= \mathcal{D}^{-1}(\mathcal{B}_p(\mathcal{D}(x_s))), \\ G^f &= \arg \min_{\theta} l(G(x_s^f; \theta), y_s), \\ x_s^{f'} &= \mathcal{D}^{-1}(\mathcal{B}_r(\mathcal{B}_p(\mathcal{D}(x_s)); I)), \\ G^{f'} &= \arg \min_{\theta} l(G(x_s^{f'}; \theta), y_s), \\ I^{SA}(I) &= \begin{cases} 0, & \text{if } \alpha(G^{f'}(x_t^f), y_t) > \alpha(G^f(x_t^f), y_t), \\ 1, & \text{if } \alpha(G^{f'}(x_t^f), y_t) < \alpha(G^f(x_t^f), y_t), \end{cases} \end{aligned}$$

where $\alpha()$ denotes a binary function that evaluates the prediction accuracy (it returns 1.0 for correct prediction and 0.0 otherwise), x_s^f/x_t^f is the input with full FCs from the source/target domain, $x_s^{f'}/x_t^{f'}$ is the filtered input (*i.e.*, $f'_s = \mathcal{B}_r(f_s; I)$) from source domain, y_s/y_t is the source/target ground truth, and I^{SA} is a 64-dimensional binary vector telling whether a FC is domain invariant (*i.e.*, $I^{SA} = 1$) or variant (*i.e.*, $I^{SA} = 0$).

We evaluated the spectrum analysis over a synthetic-to-real domain generalizable image classification task (*i.e.*, cropped SYNTHIA to ImageNet). Table 1 shows experimental results.

We can observe that removing high-frequency and low-frequency components both improve the model generalization clearly. In addition, we visualize the spectrum decomposition of source (*i.e.*, GTA) and target (*i.e.*, Cityscapes) images in Fig. 2. We can observe that low-frequency and high-frequency components capture more domain variant features (in between GTA and Cityscapes) as compared with middle-frequency components. Note several prior works exploited different FC properties successfully in other tasks such as data compression [74] and supervised learning [73].

3.3. Spectrum Analysis in Domain Randomization

This subsection describes the spectrum analysis based frequency space domain randomization (FSDR-SA) by using the “ I^{SA} ”. Different from existing domain randomization [77] that employs GANs for image stylization, we adopt histogram matching [48] for online image translation. Specifically, we adjust the frequency space coefficient histogram of source images to be similar to that of the reference images by matching their cumulative density function. With “ I^{SA} ”, this histogram matching based randomization applies to DVFs of source images only without affecting DIFs. It adds little extra parameters and computation, and is much more efficient than GAN based translation. The top part of Figure 3 illustrates how FSDR-SA works by simply replacing I^{SL} with I^{SA} .

Given a source image $x_s \in X_s$, the corresponding pixel-level label $y_s \in Y_s$, the domain invariant FCs mask vector I^{SA} and a ImageNet image $x_{img} \in X_{ImageNet}$ as randomization reference, we first transform x_s and x_{img} into frequency space and decompose them by $f_s = \mathcal{B}_p(\mathcal{D}(x_s))$, $f_{img} = \mathcal{B}_p(\mathcal{D}(x_{img}))$. We then randomize DVFs of f_s with that of f_{img} by adjusting the histogram of $f_s(1 - I^{SA})$ to match that of $f_{img}(1 - I^{SA})$. The FSDR-SA function $\mathcal{R}^{SA}(x_s; x_{ref})$ can be defined by:

$$\mathcal{R}^{SA}(x_s; x_{img}) = \mathcal{D}^{-1}(\mathcal{H}_m(f_s(1 - I^{SA}), f_{img}(1 - I^{SA}))), \quad (4)$$

where $\mathcal{H}_m()$ is the histogram matching function that adjusts the histogram of first input to match that of second.

The training loss of the FSDR-SA can be defined by:

$$\mathcal{L}_{SA} = l(G(\mathcal{R}^{SA}(x_s; x_{img})), y_s) \quad (5)$$

3.4. FSDR based on Spectrum Learning

This section describes our spectrum learning based frequency space domain randomization (FSDR-SL) that identifies DIFs and DVFs via iterative learning, which enables dynamical and adaptive FSDR. We implement FSDR-SL with entropy [60] that was widely adopted in different tasks such as semi-supervised learning [85, 15, 62], clustering [27, 28], domain adaptation [72, 87], etc.

‘Entropy’ works by measuring class overlap [15, 85, 72, 87], *i.e.*, the prediction entropy decreases as classes overlap increases [3, 45]. Leveraging this property, FSDR-SL identifies DIFs and DVFs according to the prediction entropy of reference images. Specifically, FSDR-SL is trained by using the decomposed multi-channel FCs of source images. If the trained model produces low entropy (*i.e.*, high confidence) predictions for a real target image, it indicates that the activated FCs of the target image are predictive with good semantic invariance across domains. In such cases, the employed FCs are identified as DIFs and randomization will be applied to other FCs to encourage the network to generate and learn invariant features during the iterative training process. Otherwise, no actions are taken as it is not clear whether non-activated FCs are semantically variant/uncorrelated or invariant/correlated. Note the image translation process is the same as that in FSDR-SA.

The idea of FSDR-SL is quite similar to that of self-training that either takes low-entropy/high-confidence predictions [88, 87]) as pseudo labels or directly minimizes the entropy of high-entropy/low-confidence predictions [72]. FSDR-SL reduces the overall prediction entropy by preserving low-entropy FCs while randomizing the rest FCs in domain randomization. Specifically, we first transform a batch of ImageNet images $x_{img}^b \in X_{ImageNet}$ into spatial representation of FCs $\mathcal{D}^{-1}(\mathcal{B}_p(\mathcal{D}(x_{img}^b)))$ as shown in the bottom part of Fig. 3. We then feed them to the segmentation model G to identify DIFs and DVFs based on FCs ac-

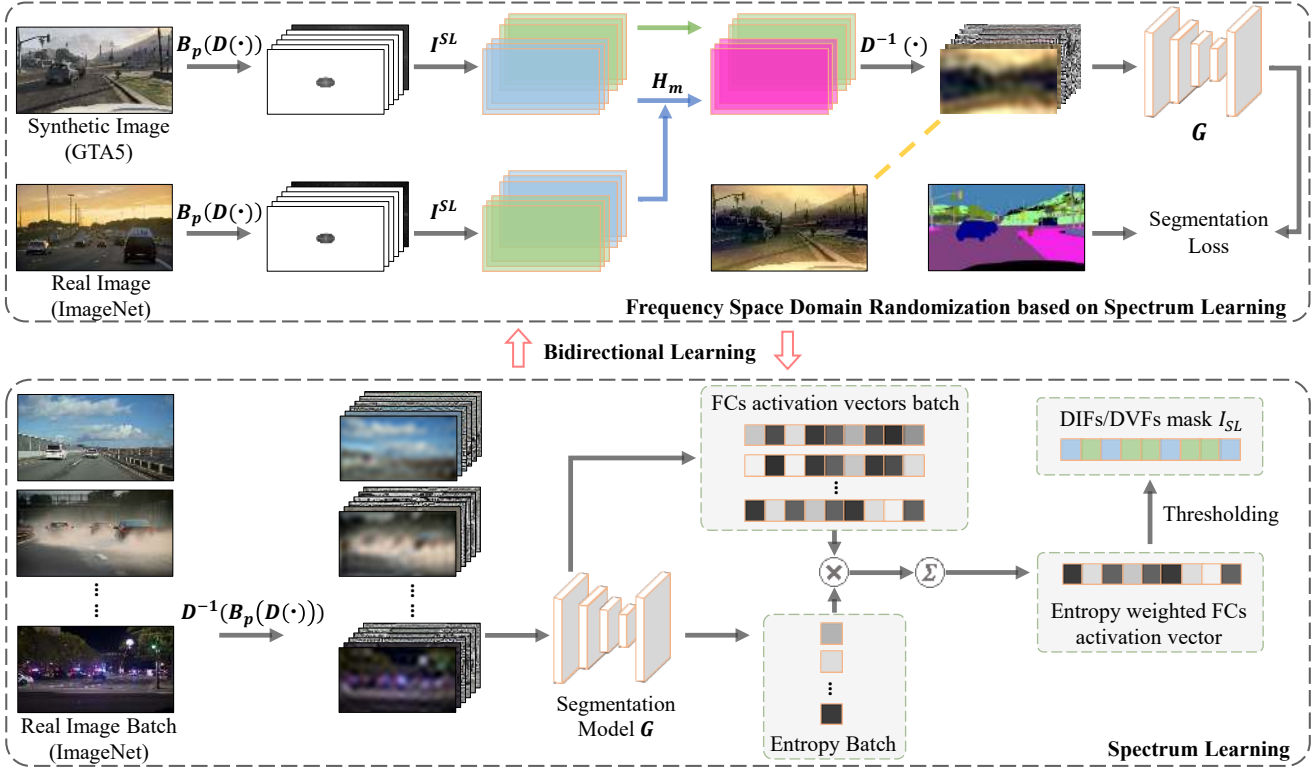


Figure 3. Overview of spectrum learning based frequency space domain randomization (FSDR-SL): FSDR-SL is a bidirectional learning framework that consists of two alternative learning processes, namely, Spectrum Learning (bottom part) and Frequency Space Domain Randomization (top part). In Spectrum Learning, we first decompose a batch of real images into spatial representation of FCs via $D^{-1}(\mathcal{B}_p(\mathcal{D}(\cdot)))$ and then feed them to the segmentation model G to identify DIFs and DVFs based on FCs activation vectors and corresponding prediction entropy. The learnt information is recorded in a binary mask vector I^{SL} . In FSDR, we first transform synthetic and real images into frequency space via $\mathcal{B}_p(\mathcal{D}(\cdot))$ and use the learnt I^{SL} to randomize DVFs of the synthetic image conditioned on the reference real image via histogram matching (*i.e.*, \mathcal{H}_m). We then employ $D^{-1}(\cdot)$ to transform the randomized synthetic images into spatial representation and use them to optimize the segmentation model via segmentation loss. Note the reference real image in FSDR is sampled from the *Real Image Batch* in Spectrum Learning, which forms a bidirectional learning framework. $\mathcal{D}(\cdot)$ and $\mathcal{D}^{-1}(\cdot)$ denotes discrete cosine transform (DCT) and inverse DCT; $\mathcal{B}_p(\cdot)$ represents a band-pass filter that decomposes the input into multiple FCs; Green boxes denote DIFs and green arrow represents keeping DIFs unchanged; Blue boxes denote DVFs and blue arrows represent randomizing DVFs; Gray arrows denote regular data flow; Yellow dash line indicates full-spectrum visualization of the randomized image.

tivation vectors and corresponding prediction entropy. The learnt information is recorded in a binary mask vector I^{SL} as follows:

$$E^b, A^b \leftarrow G(\mathcal{D}^{-1}(\mathcal{B}_p(\mathcal{D}(x_{img}^b))))),$$

$$I^{SL} = RS\left(\sum_{b=1}^B (-A^b E^b); p\right),$$

where E^b denotes a batch of averaged prediction entropy, A^b denotes a batch of averaged input FC activation vectors, B denotes the batch size, RS denotes rank and select FCs with top p portion entropy-weighted activation values as domain invariant FCs; I^{SL} is a 192 length binary vector whose values record whether each FC is domain invariant (*i.e.*, $I^{SL} = 1$) or variant (*i.e.*, $I^{SL} = 0$).

Given a source image $x_s \in X_s$, the corresponding pixel-

level label $y_s \in Y_s$ and a randomly picked ImageNet image $x_{img} \in X_{ImageNet}$ as randomization reference, we first transform x_s and x_{img} to frequency space and decompose them to $f_s = \mathcal{B}_p(\mathcal{D}(x_s))$ and $f_{img} = \mathcal{B}_p(\mathcal{D}(x_{img}))$. The DVFs of f_s (*i.e.*, $\{f_s(i) | I^{SL}(i) = 0\}$) can then be randomized with that of f_{img} via histogram matching. The FSDR-SL function $\mathcal{R}_{inter}^{SL}(x_s; x_{img})$ can be defined by:

$$\mathcal{R}^{SL}(x_s; x_{img}) = \mathcal{D}^{-1}(\mathcal{H}_m(f_s(1-I^{SL}), f_{img}(1-I^{SL}))), \quad (6)$$

where $\mathcal{H}_m(\cdot)$ denotes the histogram matching function as used in Eq. 4. The FSDR-SL training loss can thus be defined as follows:

$$\mathcal{L}_{SL} = l(G(\mathcal{R}^{SL}(x_s; x_{img})), y_s) \quad (7)$$

Note the reference real image in FSDR is sampled from the Real Image Batch that have been performed Spectrum

Learning. Once it is run out, the training goes back to Spectrum Learning process. FSDR-SL thus forms a bidirectional learning framework as shown in Figure 3, which consists of two alternative steps: 1) Perform spectrum learning with Real Image Batch on current model; 2) Perform FSDR with spectrum learnt reference images and update model.

While combining the spectrum analysis and spectrum learning, the overall training objective of the proposed FSDR can be defined as follows:

$$\min_{\theta}(\mathcal{L}_{orig} + \mathcal{L}_{SA} + \mathcal{L}_{SL}) \quad (8)$$

4. Experiments

This section presents evaluations of our proposed FSDR including datasets and implementation details, comparisons with the state-of-the-art, ablation studies, and discussion, more details to be described in the ensuing subsections.

4.1. Datasets

We evaluate FSDR over two challenging unsupervised domain generalization tasks $GTA5 \rightarrow \{\text{Cityscapes, BDD, Mapillary}\}$ and $SYNTHIA \rightarrow \{\text{Cityscapes, BDD, Mapillary}\}$ that involve two synthetic source datasets and three real target datasets. GTA5 consists of 24,966 high-resolution synthetic images, which shares 19 classes with Cityscapes, BDD, and Mapillary. SYNTHIA consists of 9,400 synthetic images, which shares 16 classes with the three target datasets. Cityscapes, BDD, and Mapillary consist of 2975, 7000, and 18000 real-world training images and 500, 1000, and 2000 validation images. Note we did not use any target data in training, but just a small subset of ImageNet [9] as references for “stylizing/randomizing” the source-domain images as in [77].

4.2. Implementation Details

We use ResNet101 [20] or VGG16 [61] (pre-trained using ImageNet [9]) with FCN-8s [38] as the segmentation model G . The optimizer is Adam with momentum of 0.9 and 0.99. The learning rate is $1e - 5$ initially and decreased with ‘step’ learning rate policy with step size of 5000 and $\gamma = 0.1$. For all backbones, we modify input channels from 3 (*i.e.*, RGB channels) to 192 (*i.e.*, the frequency-wise multi-channel spatial representation, where one full spectrum is decomposed into 64 channels) for “Spectrum Learning” that identifies DVFs and DIFs adaptively. During training, we train FSDR-SL and FSDR without using \mathcal{L}_{SL} in the first epoch (to avoid very noisy predictions at the initial training stage) and then use all losses in the ensuing epochs.

4.3. Ablation Studies

We examine different FSDR designs to find out how they contribute to the network generalization in semantic segmentation. As shown in Table 2, we trained four models

Table 2. Ablation study for the domain generalization task $GTA \rightarrow \{\text{Cityscapes, Mapillary and BDD}\}$ (using ResNet-101 as backbone) in mIoU. Losses \mathcal{L}_{orig} , \mathcal{L}_{SA} , and \mathcal{L}_{SL} are defined in Eq. 1, 5, and 7, respectively.

Method	\mathcal{L}_{orig}	\mathcal{L}_{SA}	\mathcal{L}_{SL}	mIoU		
				City.	Mapi.	BDD
Baseline	✓			33.4	27.9	27.3
FSDR-SA	✓	✓		42.1	39.2	37.8
FSDR-SL	✓		✓	43.6	42.1	40.1
FSDR	✓	✓	✓	44.8	43.4	41.2

over the UDG task $GTA5 \rightarrow \{\text{Cityscapes, BDD, Mapillary}\}$: 1) *Baseline* that is trained with \mathcal{L}_{orig} without randomization as shown in Section 3.1, 2) *FSDR-SA* that is trained with spectrum analysis based randomization using \mathcal{L}_{SA} and \mathcal{L}_{orig} , 3) *FSDR-SL* that is trained with spectrum learning based randomization using \mathcal{L}_{orig} and \mathcal{L}_{SL} , and 4) *FSDR* that is trained with \mathcal{L}_{orig} , \mathcal{L}_{SA} , and \mathcal{L}_{SL} .

We applied the four models to the validation data of Cityscapes, BDD, and Mapillary and Table 2 shows experimental results. It can be seen that *Baseline* trained with the GTA data only (*i.e.*, \mathcal{L}_{orig}) does not performs well due to domain bias. *FSDR-SA* and *FSDR-SL* outperform *Baseline* by large margins, demonstrating the importance of preserving domain invariant features in domain randomization for training domain generalizable models. Additionally, *FSDR-SL* outperforms *FSDR-SA* clearly which is largely due to the adaptive and iterative spectrum learning that enables the dynamic and adaptive FSDR and a bidirectional learning framework. Further, *FSDR* performs the best consistently. It shows that *FSDR-SA* and *FSDR-SL* are complementary where the static spectrum analysis forms certain bases for effective and stable spectrum leaning while training domain generalizable models.

4.4. Comparisons with the State-of-Art

We compared FSDR with a number of state-of-the-art UDG methods as shown in Tables 3 and 4 (highlighted by X). The comparisons were performed over two tasks $GTA5 \rightarrow \{\text{Cityscapes, BDD, Mapillary}\}$ and $SYNTHIA \rightarrow \{\text{Cityscapes, BDD, Mapillary}\}$ where two network backbones ResNet101 and VGG16 were evaluated for each task. As the two tables show, FSDR outperforms all state-of-the-art UDG methods clearly and consistently across both tasks and two different network backbones. The superior segmentation performance is largely attributed to the proposed frequency space domain randomization that identifies and keeps DIFs unchanged and randomizes DVFs only. Without identifying and preserving DIFs, state-of-the-art methods [77, 14] tend to over-randomize images which may degrade image structures and semantics and lead to sup-optimal network generalization. We provide the qualitative comparison in Figure 4.

Table 3. Domain generalization performance for the task GTA \rightarrow {Cityscapes, Mapillary, and BDD} in mIoU. “w/ Tgt” labels the methods that train models with (✓ *i.e.*, domain adaptation) and without (✗ *i.e.*, domain generalization) accessing target-domain data in Cityscapes.

Net.	Method	w/ Tgt	C	M	B	Mean
Res Net 101	CBST [88]	✓	44.9	40.3	40.5	41.9
	AdaSeg. [66]	✓	41.4	38.3	36.2	38.6
	MinEnt [72]	✓	42.3	38.5	34.4	38.4
	FDA [76]	✓	45.0	39.6	38.1	40.9
	IDA [46]	✓	45.1	39.4	37.6	40.7
	CrCDA [26]	✓	43.7	39.3	37.3	40.1
	IBN-Net [47]	✗	40.3	35.9	35.6	37.2
	DRPC [77]	✗	42.5	38.0	38.7	39.8
	Ours (FSDR)	✗	44.8	43.4	41.2	43.1
VGG 16	CBST [88]	✓	38.1	34.6	33.9	35.5
	AdaSeg. [66]	✓	35.0	32.6	31.3	33.0
	MinEnt [72]	✓	32.8	30.7	29.5	31.0
	FDA [76]	✓	37.9	33.8	32.1	34.6
	IDA [46]	✓	38.5	34.2	32.3	35.0
	CrCDA [26]	✓	36.1	32.6	31.8	33.5
	IBN-Net [47]	✗	34.8	31.0	30.4	32.0
	DRPC [77]	✗	36.1	32.3	31.6	33.3
	Ours (FSDR)	✗	38.3	37.6	34.4	37.1

We also compared FSDR with several state-of-the-art UDA methods as shown in Tables 3 and 4 (highlighted by ✓), where each method adapts to Cityscapes only and the adapted model is then evaluated over Cityscapes, Mapillary, and BDD. It can be seen that FSDR is even on par with the state-of-the-art UDA method which accesses target-domain data (*i.e.* Cityscapes) in training. In addition, FSDR performs more stably across the three target domains, whereas UDA methods [88, 72, 66, 46] demonstrate much larger mIoU drops for unseen (in training) Mapillary and BDD data. We conjecture that the UDA models may over-fit to the seen Cityscapes data which leads to poor generalization for unseen Mapillary and BDD data.

4.5. Discussion

FSDR is *complementary* with most existing domain adaptation and domain generalization networks, which can be easily incorporated into them with consistent performance boost but little extra parameters and computation. We evaluated this feature by incorporating FSDR into a number of domain adaptation and domain generalization networks as shown in Table 5. During training, we reconstruct the FSDR randomized multi-FC representation back to the full-spectrum representation with 3 channels (*i.e.* Ours*) for compatibility with the compared domain adaptation and generalization methods that work with full-

Table 4. Domain generalization performance for the task SYNTHIA \rightarrow {Cityscapes, Mapillary, and BDD} in mIoU. “w/ Tgt” labels the methods that train models with (✓ *i.e.*, domain adaptation) and without (✗ *i.e.*, domain generalization) accessing target-domain data in Cityscapes.

Net.	Method	w/ Tgt	C	M	B	Mean
Res Net 101	CBST [88]	✓	41.4	37.1	37.6	38.7
	AdaSeg. [66]	✓	38.2	36.1	35.3	36.5
	MinEnt [72]	✓	38.1	35.8	35.5	36.4
	FDA [76]	✓	41.2	36.1	36.4	37.9
	IDA [46]	✓	41.7	36.5	37.0	38.4
	CrCDA [26]	✓	39.0	36.4	36.7	37.4
	IBN-Net [47]	✗	37.5	33.7	33.0	34.7
	DRPC [77]	✗	37.6	34.1	34.3	35.4
	Ours (FSDR)	✗	40.8	39.6	37.4	39.3
VGG 16	CBST [88]	✓	38.2	33.5	32.2	34.6
	AdaSeg. [66]	✓	32.6	30.3	29.4	30.8
	MinEnt [72]	✓	31.4	29.8	28.9	30.0
	FDA [76]	✓	37.9	33.1	31.8	34.2
	IDA [46]	✓	38.6	34.2	32.7	35.1
	CrCDA [26]	✓	36.4	32.8	31.9	33.7
	IBN-Net [47]	✗	33.9	31.1	30.4	31.8
	DRPC [77]	✗	35.5	32.2	29.5	32.4
	Ours (FSDR)	✗	37.9	37.2	34.1	36.4

Table 5. FSDR is complementary with existing domain adaptation and generalization methods. For the task GTA \rightarrow {Cityscapes, Mapillary, and BDD}, including FSDR (Ours*) improves the domain adaptation and generalization performance consistently. “w/ Tgt” labels methods that train models with (✓ *i.e.* domain adaptation) or without (✗ *i.e.*, domain generalization) accessing target data in Cityscapes.

	w/ Tgt	Cityscapes		Mapillary		BDD	
		Base	+Ours*	Base	+Ours*	Base	+Ours*
Adapt-SegMap [66]	✓	41.4	45.6	38.3	43.9	36.2	41.9
MinEnt [72]	✓	42.3	45.7	38.5	43.7	34.4	41.7
CBST [88]	✓	44.9	46.8	40.3	44.3	40.5	42.8
FDA [76]	✓	45.0	46.1	39.6	44.2	38.1	42.1
IBN-Net [47]	✗	40.3	45.3	35.9	44.0	35.6	42.1
DRPC [77]	✗	42.5	45.8	38.0	44.2	38.7	42.6

spectrum three-channel images. As Table 5 shows, the incorporation of FSDR $\{\mathcal{L}_{SA}, \mathcal{L}_{SL}\}$ improves the semantic segmentation of state-of-the-art networks consistently. As the incorporation of FSDR just includes a few losses without changing network structures, the inference has little extra parameters and computation once the model is trained.

The parameter p is important which controls the sensitivity of DIFs and DVFs identification/selection in Spectrum Learning. We studied the sensitivity of p by changing it from 0 to 1 with a step of 1/6, where ‘0’ means that all FCs are DVFs (*i.e.*, randomizing all FCs) and ‘1’ means that all FCs are DIFs (*i.e.*, no randomization). The task is

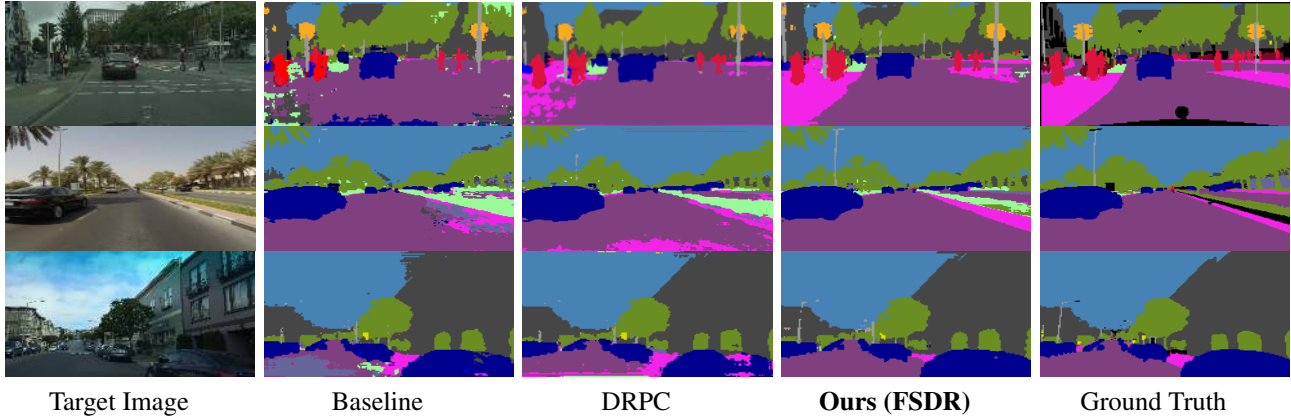


Figure 4. Qualitative illustration of domain generalizable semantic segmentation for GTA5 to Cityscapes (1st row), Mapillary (2nd row), and BDD (3rd row). FSDR preserves domain invariant information during domain randomization which produces better semantic segmentation especially around edge/class-transition area. As a comparison, DRPC [77] does not isolate and preserve domain invariant feature which leads to sub-optimal segmentation.

Table 6. The sensitivity of parameter p affects domain generalization: For the task GTA \rightarrow Cityscapes, the domain generalization performance varies with p as evaluated in mIoU.

Method	Proportion of preserved FCs						
	1	5/6	4/6	3/6	2/6	1/6	0
FSDR	33.4	43.8	44.3	44.8	44.2	41.0	38.3

domain generalizable semantic segmentation with GTA and Cityscapes as the source and target domains. Table 6 shows experimental results. It can be seen that FSDR are not sensitive to p while p lies between 1/6 and 5/6. In addition, the FSDR performance drops clearly while p is around either 0 or 1, demonstrating the necessity and importance of FC selection in training generalizable networks.

FSDR is also *generic* and can be easily adapted to other tasks. We evaluated this feature by conducting a preliminary UDG-based object detection test that generalizes from SYNTHIA to Cityscapes, Mapillary, and BDD. The experiment setup is the same as what we adopted in the earlier experiments except the change of tasks. Table 7 compares FSDR with a number of UDA and UDG based object detection methods (using ResNet101 as backbone). We can see that FSDR outperforms both domain adaptation and domain generalization methods consistently, demonstrating its genericity in different tasks.

5. Conclusion

This paper presents a frequency space domain randomization (FSDR) technique that randomizes images in frequency space by identifying and randomizing domain-variant FCs (DVF) while keeping domain-invariant FCs (DIF) unchanged. The proposed FSDR has two unique features: 1) it decomposes images into DIFs and DVFs which allows explicit access and manipulation of them and

Table 7. FSDR is generic and can work for other tasks like object detection: The detection task is SYNTHIA \rightarrow {Cityscapes, Mapillary, and BDD} as evaluated using metric mAP. “w/ Tgt” labels methods that train models with (\checkmark *i.e.*, domain adaptation) and without (\times *i.e.*, domain generalization) accessing target-domain data in Cityscapes.

Net.	Method	w/ Tgt	C	M	B	Mean
Res Net 101	Faster-RCNN	\times	24.3	20.8	20.1	21.7
	DA [6]	\checkmark	30.2	21.2	21.8	24.4
	MinEnt [72]	\checkmark	30.2	21.7	22.4	24.7
	CBST [88]	\checkmark	32.7	23.5	23.9	26.7
	FDA [76]	\checkmark	32.4	23.3	23.8	26.5
	IBN-Net [47]	\times	30.1	22.3	23.1	25.1
	Ours (FSDR)	\times	33.5	24.9	25.2	27.8

better control in image randomization; 2) it has minimal effects on image semantics and domain-invariant features. Specifically, we designed spectrum analysis based FSDR (FSDR-SA) and spectrum learning based FSDR (FSDR-SL) both of which can identify DIFs and DVFs effectively. FSDR achieves superior segmentation performance and can be easily incorporated into state-of-the-art domain adaptation and generalization networks with consistent improvement in domain generalization.

Acknowledgement. This research was conducted at Singtel Cognitive and Artificial Intelligence Lab for Enterprises (SCALE@NTU), which is a collaboration between Singapore Telecommunications Limited (Singtel) and Nanyang Technological University (NTU) that is funded by the Singapore Government through the Industry Alignment Fund - Industry Collaboration Projects Grant.

References

- [1] Yanpeng Cao, Dayan Guan, Weilin Huang, Jiangxin Yang, Yanlong Cao, and Yu Qiao. Pedestrian detection with unsupervised multispectral feature learning using deep neural networks. *information fusion*, 46:206–217, 2019. 2
- [2] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 2
- [3] Vittorio Castelli and Thomas M Cover. The relative value of labeled and unlabeled samples in pattern recognition with an unknown mixing parameter. *IEEE Transactions on information theory*, 42(6):2102–2117, 1996. 4
- [4] 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. 1
- [5] Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, and Kevin Chetty. Re-weighted adversarial adaptation network for unsupervised domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [6] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 1, 8
- [7] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7892–7901, 2018. 2
- [8] 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. 1
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [10] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *Advances in Neural Information Processing Systems*, pages 6450–6461, 2019. 2
- [11] Chuang Gan, Tianbao Yang, and Boqing Gong. Learning attributes equals multi-source domain generalization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 87–97, 2016. 2
- [12] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pages 2551–2559, 2015. 2
- [13] Boqing Gong, Kristen Grauman, and Fei Sha. Reshaping visual datasets for domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1286–1294, 2013. 2
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 6
- [15] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005. 4
- [16] Thomas Grubinger, Adriana Birlutiu, Holger Schöner, Thomas Natschläger, and Tom Heskes. Multi-domain transfer component analysis for domain generalization. *Neural processing letters*, 46(3):845–855, 2017. 2
- [17] Dayan Guan, Jiaying Huang, Shijian Lu, and Aoran Xiao. Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition*, 112:107764, 2021. 2
- [18] Dayan Guan, Jiaying Huang, Aoran Xiao, Shijian Lu, and Yanpeng Cao. Uncertainty-aware unsupervised domain adaptation in object detection. *arXiv preprint arXiv:2103.00236*, 2021. 2
- [19] Dayan Guan, Xing Luo, Yanpeng Cao, Jiangxin Yang, Yanlong Cao, George Vosselman, and Michael Ying Yang. Unsupervised domain adaptation for multispectral pedestrian detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [21] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, pages 1989–1998, 2018. 2
- [22] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. 1, 2
- [23] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1335–1344, 2018. 2
- [24] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Cross-view regularization for domain adaptive panoptic segmentation. *arXiv preprint arXiv:2103.02584*, 2021. 2
- [25] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Mlan: Multi-level adversarial network for domain adaptive semantic segmentation. *arXiv preprint arXiv:2103.12991*, 2021. 2
- [26] Jiaying Huang, Shijian Lu, Dayan Guan, and Xiaobing Zhang. Contextual-relation consistent domain adaptation for semantic segmentation. In *European Conference on Computer Vision*, pages 705–722. Springer, 2020. 2, 7

- [27] Himalaya Jain, Joaquin Zepeda, Patrick Pérez, and Rémi Gribonval. Subic: A supervised, structured binary code for image search. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 833–842, 2017. 4
- [28] Himalaya Jain, Joaquin Zepeda, Patrick Pérez, and Rémi Gribonval. Learning a complete image indexing pipeline. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4933–4941, 2018. 4
- [29] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019. 1
- [30] Guoliang Kang, Liang Zheng, Yan Yan, and Yi Yang. Deep adversarial attention alignment for unsupervised domain adaptation: the benefit of target expectation maximization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–416, 2018. 1
- [31] Minsu Kim, Sunghun Joung, Seungryong Kim, JungIn Park, Ig-Jae Kim, and Kwanghoon Sohn. Cross-domain grouping and alignment for domain adaptive semantic segmentation. *arXiv preprint arXiv:2012.08226*, 2020. 2
- [32] Marvin Klingner, Jan-Aike Termöhlen, Jacob Ritterbach, and Tim Fingscheidt. Unsupervised batchnorm adaptation (ubna): A domain adaptation method for semantic segmentation without using source domain representations. *arXiv preprint arXiv:2011.08502*, 2020. 2
- [33] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 2
- [34] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. *arXiv preprint arXiv:1710.03463*, 2017. 2
- [35] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1446–1455, 2019. 2
- [36] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. 2
- [37] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019. 2
- [38] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1, 6
- [39] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016. 2
- [40] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2507–2516, 2019. 1, 2
- [41] Yawei Luo, Zhedong Zheng, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Macro-micro adversarial network for human parsing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 418–434, 2018. 1
- [42] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Best sources forward: domain generalization through source-specific nets. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 1353–1357. IEEE, 2018. 2
- [43] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Robust place categorization with deep domain generalization. *IEEE Robotics and Automation Letters*, 3(3):2093–2100, 2018. 2
- [44] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18, 2013. 2
- [45] Terence J O’neill. Normal discrimination with unclassified observations. *Journal of the American Statistical Association*, 73(364):821–826, 1978. 4
- [46] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [47] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. 7, 8
- [48] Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987. 4
- [49] Aayush Prakash, Shaad Boochoon, Mark Brophy, David Acuna, Eric Cameracci, Gavriel State, Omer Shapira, and Stan Birchfield. Structured domain randomization: Bridging the reality gap by context-aware synthetic data. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7249–7255. IEEE, 2019. 2
- [50] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 2
- [51] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016. 1
- [52] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3234–3243, 2016. 1

- [53] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016. 2
- [54] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. *arXiv preprint arXiv:1711.01575*, 2017. 2
- [55] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018. 1, 2
- [56] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1
- [57] Fatemeh Sadat Saleh, Mohammad Sadegh Aliakbarian, Mathieu Salzmann, Lars Petersson, and Jose M Alvarez. Effective use of synthetic data for urban scene semantic segmentation. In *European Conference on Computer Vision*, pages 86–103. Springer, 2018. 2
- [58] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3752–3761, 2018. 1, 2
- [59] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 2
- [60] Claude Elwood Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948. 4
- [61] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [62] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint arXiv:1511.06390*, 2015. 4
- [63] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018. 2
- [64] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30. IEEE, 2017. 2
- [65] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 969–977, 2018. 2
- [66] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018. 1, 2, 7
- [67] Yi-Hsuan Tsai, Ming-Yu Liu, Deqing Sun, Ming-Hsuan Yang, and Jan Kautz. Learning binary residual representations for domain-specific video streaming. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 1
- [68] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1456–1465, 2019. 2
- [69] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017. 1, 2
- [70] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in neural information processing systems*, pages 5334–5344, 2018. 2
- [71] Tuan-Hung Vu, Wongun Choi, Samuel Schulter, and Manmohan Chandraker. Memory warps for learning long-term online video representations. *arXiv preprint arXiv:1803.10861*, 2018. 1
- [72] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019. 1, 2, 4, 7, 8
- [73] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8684–8694, 2020. 4
- [74] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1740–1749, 2020. 4
- [75] Jinyu Yang, Weizhi An, Sheng Wang, Xinliang Zhu, Chaochao Yan, and Junzhou Huang. Label-driven reconstruction for domain adaptation in semantic segmentation. In *European Conference on Computer Vision*, pages 480–498. Springer, 2020. 2
- [76] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 2, 7, 8
- [77] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2100–2110, 2019. 2, 4, 6, 7, 8

- [78] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016. [1](#)
- [79] Gongjie Zhang, Shijian Lu, and Wei Zhang. Cad-net: A context-aware detection network for objects in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 57(12):10015–10024, 2019. [1](#)
- [80] Xiaobing Zhang, Haigang Gong, Xili Dai, Fan Yang, Nianbo Liu, and Ming Liu. Understanding pictograph with facial features: End-to-end sentence-level lip reading of chinese. In *AAAI*, pages 9211–9218, 2019. [2](#)
- [81] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2020–2030, 2017. [2](#)
- [82] Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, and Dimitris N Metaxas. Knowledge as priors: Cross-modal knowledge generalization for datasets without superior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6528–6537, 2020. [2](#)
- [83] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 598–607, 2019. [2](#)
- [84] Hongyuan Zhu, Fanman Meng, Jianfei Cai, and Shijian Lu. Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation. *Journal of Visual Communication and Image Representation*, 34:12–27, 2016. [1](#)
- [85] Xiaojin Jerry Zhu. Semi-supervised learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2005. [4](#)
- [86] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5982–5991, 2019. [2](#)
- [87] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [4](#)
- [88] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 289–305, 2018. [2](#), [4](#), [7](#), [8](#)