

# Scale Aware Adaptation for Land-Cover Classification in Remote Sensing Imagery

Xueqing Deng<sup>1</sup>, Yi Zhu<sup>\*2</sup>, Yuxin Tian<sup>1</sup>, and Shawn Newsam<sup>1</sup>

<sup>1</sup>EECS, UC Merced {xdeng7, ytian8, snewsam}@ucmerced.edu

<sup>2</sup>Amazon Web Services {yzaws}@amazon.com

## Abstract

Land-cover classification using remote sensing imagery is an important Earth observation task. Recently, land cover classification has benefited from the development of fully connected neural networks for semantic segmentation. The benchmark datasets available for training deep segmentation models in remote sensing imagery tend to be small, however, often consisting of only a handful of images from a single location with a single scale. This limits the models' ability to generalize to other datasets. Domain adaptation has been proposed to improve the models' generalization but we find these approaches are not effective for dealing with the scale variation commonly found between remote sensing image collections. We therefore propose a scale aware adversarial learning framework to perform joint cross-location and cross-scale land-cover classification. The framework has a dual discriminator architecture with a standard feature discriminator as well as a novel scale discriminator. We also introduce a scale attention module which produces scale-enhanced features. Experimental results show that the proposed framework outperforms state-of-the-art domain adaptation methods by a large margin. The open-sourced codes are available on Github: [https://github.com/xdeng7/scale-aware\\_da](https://github.com/xdeng7/scale-aware_da).

## 1. Introduction

High-resolution remote sensing imagery is becoming increasingly available due to the number and revisit rate of traditional satellite and airborne image capture platforms as well as the advent of newer platforms like drones. This imagery provides convenient and large-scale coverage and so is being applied to a number of societally important problems such as land cover segmentation [29], traffic monitoring [34], urban planning [5], vehicle detection [11], building extraction [51], geolocization[42] etc. While remote

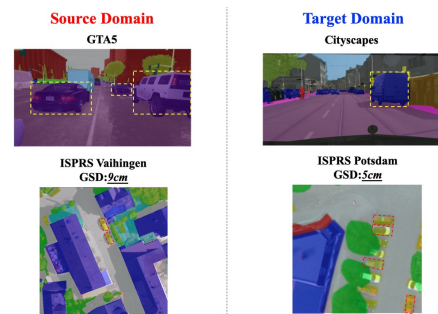


Figure 1: Object sizes vary both within and between ground-level image datasets such as in the driving scenes at the top. Segmentation models trained on one dataset are already scale-invariant and so standard domain adaptation techniques are not designed to explicitly account for scale variation. In contrast, object sizes usually do not vary within RS image datasets since all the images have the same ground sample distance (GSD). We propose a novel framework that explicitly accounts for scale when adapting between RS image datasets with different scales such as at the bottom.

sensing (RS) image analysis has benefited from advances in deep learning in the computer vision community, there often remains unique challenges that limit the straightforward application of standard approaches to the RS case.

Another issue that limits the performance of automated RS image analysis, particularly deep learning approaches, is that the availability of the annotated ground truth data needed for training has not kept pace with the imagery (or its diversity). As an example, current semantic segmentation datasets for land cover classification, which are very labor intensive to produce, contain limited labeled samples from only a few locations. The ISPRS Vaihingen dataset [1] contains just 33 labeled tiles with 6 semantic classes. The recent SkyScapes [4] dataset has more classes with 30 but still contains only 16 images. DeepGlobe [15] is the largest collection, containing 1146 images with 7 classes. Deep

<sup>\*</sup>Work was done prior to joining Amazon.

learning models trained on these (relatively) small datasets have difficulty generalizing to other image sets, i.e. large-scale WorldView imagery. They suffer from the so-called domain shift problem [12, 43]. One of the great benefits of using RS imagery is its fast revisit time and large coverage for Earth observation. However, this scalability is limited by domain shift problems.

Domain shift in RS imagery can happen along a number of dimensions including differences in sensor, spectra, resolution, etc. which have been widely explored [7, 47, 8]. There has been much less work, somewhat surprisingly, on the cross-location domain shift problem in which a model trained on one location that has ground truth data, the source location, is applied to another location without ground truth data, the target location. The work that has been done simply applies standard domain adaptation techniques [7, 50, 16]. However, none of this work explicitly considers the important notion of scale which we believe is often integral to the cross-location problem especially for tasks where there is limited training data like semantic segmentation in RS imagery. We therefore propose a novel scale adaptive framework for improved cross-location domain adaptation. The framework is adversarial and includes separate feature and scale discriminators as well as a scale attention module.

Domain shift caused by location also exists in other datasets such as road scene imagery. However, we found that scale is more important for overhead (RS) than (most) ground-level imagery. A simple visual investigation illustrates why scale is more important. Fig. 1 contains a pair of images from different driving scene datasets and a pair of images from different RS image datasets. (The ground truth segmentation masks are overlaid on the images.) The objects in the road scenes vary in size both within a dataset (even within an image) as well as between datasets. A deep learning model trained on one dataset learns scale-invariant features so is better able to generalize to the other dataset. However, in the RS imagery, due to the planar nature of the scene and the nadir viewpoint, the size of the objects does not vary (much) within a dataset if all the images have the same resolution or ground sample distance (the physical size of a pixel on the ground typically indicated in meters or similar) which is usually the case. The RS imagery in the source domain in Fig. 1 has a GSD of 9cm so all the cars measure around  $17 \times 22$  pixels assuming  $1.5 \times 2$  meter cars. The GSD of the target domain is just 5cm so all the cars are proportionally larger and measure around  $30 \times 40$  pixels. A model trained on one dataset will not learn the scale-invariant features needed to label the differently sized cars in the other dataset. Data augmentation during training through random scaling is not an effective solution. Additional examples of showing that scale is more important for RS than regular imagery can be found in the supplementary

material.

We therefore develop a novel domain adaptation framework that explicitly accounts for scale changes in order to improve generalization in cross-location semantic segmentation for RS imagery. The proposed framework contains dual adversarial discriminators including a standard feature discriminator and a novel scale discriminator as well as a scale attention module. There are separate discriminators for feature and scale adaptation between the source and target domains. The scale attention module selectively weights concatenated multi-scale features to obtain scale-enhanced features. Our novel contributions include:

- We establish and demonstrate that explicitly accounting for scale variation is integral to RS image domain adaptation yet no existing work does this. We propose an innovative scale adaptive framework for cross-location semantic segmentation which includes a novel scale discriminator and a scale attention module to facilitate training. We demonstrate that scale-aware adaptation results in significant performance gains.
- We show our proposed approach outperforms state-of-the-art domain adaptation methods on several cross-dataset segmentation tasks such as Potsdam  $\leftrightarrow$  Vaihingen and DeepGlobe  $\leftrightarrow$  Vaihingen.

## 2. Related Work

**Semantic segmentation** There is a large body of literature on semantic segmentation [9, 10, 20, 56]. Here, we review only the most related work in terms of handling scale [54, 19, 31, 26]. Fully Convolutional Network (FCN) based methods [33] have made great progress in semantic segmentation. To capture multi-scale information, DeepLab networks [9], PSPNet [54] and CaseNet [26] develop mechanisms to integrate multi-scale features. Benefiting from the above exploration on regular images, a number of works [32, 37, 3, 55, 18, 17] have applied the techniques to pixel-wise land-cover classification in RS imagery. These methods focus on multi-scale analysis but not explicitly on scale adaptation as proposed in this paper.

**Domain adaptation** Numerous domain adaptation methods have been developed to improve image classification by aligning the feature distributions between the source and the target domains [45, 22, 39, 2]. Domain adaptation for segmentation has recently started to receive more attention due to the expense of performing pixel-wise annotation. Many approaches have been proposed for road scene segmentation [43, 13, 23, 24, 44, 48, 27, 58, 14]. Since scale is much less significant for road scene images, these studies focus mainly on adapting texture, appearance, etc. There has been some work on using Generative Adversarial Networks (GAN) for domain adaptation in RS image segmentation [6, 50]. However, these approaches just apply standard

computer vision methods without considering challenges specific to RS imagery. We instead propose a framework that addresses the importance of scale when adapting between domains.

**Attention** Attention was originally proposed to exploit long-range dependencies in machine translation [46]. It has since been adopted in a number of computer vision tasks [53, 49, 52, 57]. Self-attention has been used as a non-local operation to learn positional relations in video classification [49], to learn a better image generator [52], as well as to learn both channel and spatial attention for scene segmentation [21] and land cover segmentation [38]. In order to augment the extraction of multi-scale information, we propose a scale attention module (channel attention) which improves the scale adaptation.

### 3. Problem Formulation

We formulate our problem as cross-scale and cross-location domain adaptation for semantic segmentation in RS imagery (pixel-wise land-cover classification). We assume the domain shift is caused by different scales and locations between the source and target datasets. We recognize, though, that different locations do not necessarily have different scales. Our framework is unified in that it can deal with domain shift caused by scale or by location or by both.

We denote the different locations as  $\mathcal{S}$  and  $\mathcal{T}$  and the different scales as  $\theta$  and  $\sigma$ . We denote an image  $x$  from source location  $\mathcal{S}$  with scale  $\theta$  as  $x_{\mathcal{S}^\theta} \in X\{\text{location} = \mathcal{S}, \text{scale} = \theta\}$  and its label as  $y_{\mathcal{S}^\theta} \in Y\{\text{location} = \mathcal{S}, \text{scale} = \theta\}$ .  $X$  and  $Y$  represent all the images and labels in one domain. Our goal is to adapt, in an unsupervised manner, a segmentation model  $\mathcal{G}$  trained using images  $X_{\mathcal{S}^\theta}$  and labels  $Y_{\mathcal{S}^\theta}$  from source domain  $\mathcal{S}^\theta$  to perform segmentation in target domain  $\mathcal{T}^\sigma$  which has only images  $X_{\mathcal{T}^\sigma}$  to produce predicted segmentation masks  $\hat{Y}_{\mathcal{T}^\sigma}$ .

## 4. Methods

In this section, we describe our scale aware domain adaptation framework. We first revisit conventional domain adaptation methods which only have feature discriminators. We then describe our new scale discriminator for addressing the scale variation between domains. Finally, we describe our scale attention module for obtaining scale-enhanced features for improved segmentation.

### 4.1. Domain Adaptation Revisited

Recent domain adaptation methods for semantic segmentation are adversarially based [43, 23]. An image from either the source or target domain is input to the segmentation network. The resulting feature maps or segmentation predictions are then fed to a discriminator which tries to determine the domain of the input. The goal of the segmentation network now is to not only produce an accurate seg-

mentation of the source domain image (for which we have labels) but to also fool the discriminator. This forces the internal representations of the model to align between domains so it can better exploit its supervised training in the source domain when it is applied to the target domain.

The segmentation network is now updated using two losses, the segmentation loss when the input is from the source domain and the adversarial loss when the input is from the target domain. Given any segmentation network  $\mathcal{G}$  (e.g., [54, 10]), source image  $x_{\mathcal{S}^\theta}$  will result in prediction  $p_{\mathcal{S}^\theta} = \mathcal{G}(x_{\mathcal{S}^\theta})$  and target image  $x_{\mathcal{T}^\sigma}$  will result in prediction  $p_{\mathcal{T}^\sigma} = \mathcal{G}(x_{\mathcal{T}^\sigma})$ . Note that we choose to use the low-dimensional softmax output predictions instead of the feature maps following [43] since they contain rich spatial information shared between locations which makes it easier for the network to adapt. When the input is from the source domain, the multi-class cross-entropy segmentation loss

$$\mathcal{L}_{seg}(x_{\mathcal{S}^\theta}) = -y_{\mathcal{S}^\theta} \log(p_{\mathcal{S}^\theta}), \quad (1)$$

is computed where  $y_{\mathcal{S}^\theta}$  is the annotated label. A segmentation loss cannot be computed for a target image since its label is not available. So, in order to adapt the segmentation network to the target domain, a feature discriminator  $\mathcal{D}_{feat}$  is added and an adversarial loss is calculated

$$\mathcal{L}_{adv\_feat}(x_{\mathcal{T}^\sigma}) = -\log(\mathcal{D}_{feat}(p_{\mathcal{T}^\sigma})). \quad (2)$$

This is a binary cross-entropy loss, designed to fool the discriminator by forcing the space of target predictions  $p_{\mathcal{T}^\sigma}$  to match the space of source predictions  $p_{\mathcal{S}^\theta}$ . At this point, the overall objective for updating the segmentation network is a combination of  $\mathcal{L}_{seg}$  and  $\mathcal{L}_{adv\_feat}$

$$\mathcal{L}(x_{\mathcal{S}^\theta}, x_{\mathcal{T}^\sigma}) = \mathcal{L}_{seg} + \lambda_f \mathcal{L}_{adv\_feat}, \quad (3)$$

where  $\lambda_f$  is the weight of the adversarial loss.

The feature discriminator  $\mathcal{D}_{feat}$  is updated using its own adversarial loss

$$\mathcal{L}_{\mathcal{D}_{feat}}(p) = -(1 - z) \log(\mathcal{D}_{feat}(p)) + z \log(\mathcal{D}_{feat}(p)) \quad (4)$$

where  $z = 0$  if the output prediction map  $p$  is from the target location  $\mathcal{T}$ , and  $z = 1$  if it is from the source location  $\mathcal{S}$ . The segmentation network and the feature discriminator are optimized in an alternating manner. When one is being updated, the other is frozen.

### 4.2. Scale Discriminator

The standard domain adaption framework above achieves decent results for cross-location segmentation when the source and target domain have similar scale, i.e.,  $x_{\mathcal{S}^\theta}$  and  $x_{\mathcal{T}^\theta}$ . However, it does not do well when the scale varies, i.e.,  $x_{\mathcal{S}^\theta}$  and  $x_{\mathcal{T}^\sigma}$ . When the scale of target dataset is different from the source, the performance of a model trained with just a feature discriminator decreases by 20% in the RS case (details in the supplementary material).

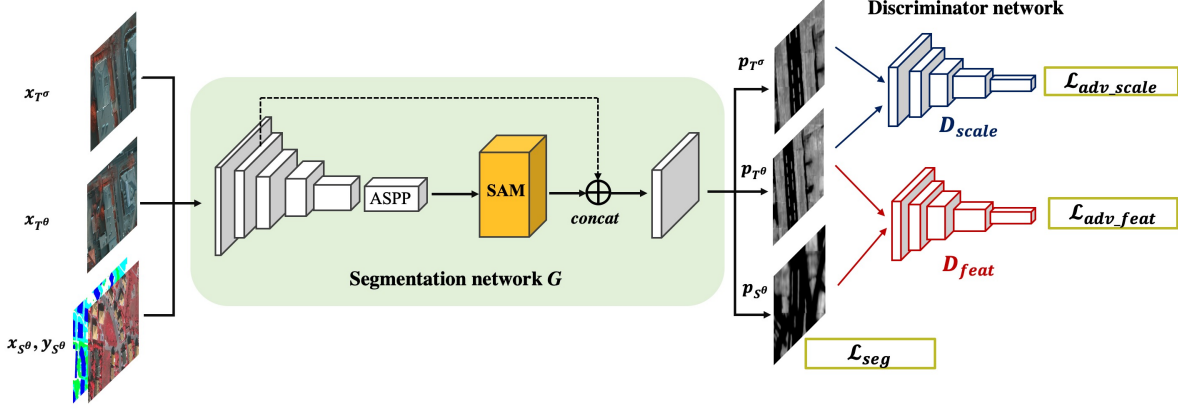


Figure 2: Our proposed scale aware adaptation framework which contains two adversarial learning components, one for feature adaptation and another for scale adaptation. We use DeepLabV3+ as our segmentation network. ASPP is an Atrous Spatial Pyramid Pooling module. We also propose a scale attention module (SAM) to produce scale-enhanced features. See the text for more details.

We therefore propose a dual discriminator network which includes the standard feature discriminator as well as a new scale discriminator. We split the adaptation task into two sub-tasks, one that focuses on (cross-location) feature adaptation and another that focuses on scale adaptation. The scale discriminator has the same network structure as the feature discriminator. The framework is shown in Fig. 2.

We now have three kinds of input images instead of two. Besides the source  $x_{S^\sigma}$  and target  $x_{T^\sigma}$  images, we derive a resized version of the target image  $x_{T^\theta}$  via bilinear interpolation whose scale matches that of the source. These three types of inputs allow us to create two adversarial flows, one that considers images from the same location but different scales, and another that considers images with the same scale but from different locations. The key to our framework is the new, matched scale image  $x_{T^\theta}$ . It allows our network to focus on adapting features between the source and target locations without also having to adapt for scale.

The feature discriminator  $\mathcal{D}_{feat}$  is now updated using images with the same scale but from different locations using the adversarial loss

$$\mathcal{L}_{\mathcal{D}_{feat}}(p_\theta) = -(1-z) \log(\mathcal{D}_{feat}(p_\theta)) + z \log(\mathcal{D}_{feat}(p_\sigma)). \quad (5)$$

Note that the difference between (4) and (5) is that the output prediction  $p$  in (5) has the same scale  $\theta$  no matter whether it is from the source or target location.

The scale discriminator  $\mathcal{D}_{scale}$  is updated using images from the same location but with different scales using the adversarial loss

$$\mathcal{L}_{\mathcal{D}_{scale}}(p_\mathcal{T}) = -(1-z) \log(\mathcal{D}_{scale}(p_\mathcal{T})) + z \log(\mathcal{D}_{scale}(p_\mathcal{T})), \quad (6)$$

where  $z = 0$  if  $p_\mathcal{T}$  is has the target scale  $\sigma$ , and  $z = 1$  if it has the source scale  $\theta$ .

We now turn to the update of the segmentation network.

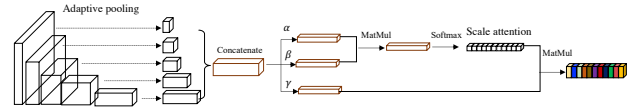


Figure 3: Proposed scale attention module (SAM). The input is feature maps from the segmentation network. The output is scale-enhanced features. MatMul: matrix multiplication.

Similar to (2), a feature adversarial loss is calculated using the feature discriminator to adapt the segmentation network to the target features

$$\mathcal{L}_{adv\_feat}(x_{T^\theta}) = -\log(\mathcal{D}_{feat}(p_{T^\theta})). \quad (7)$$

In order to adapt the segmentation network to the target scale, a scale adversarial loss is computed as

$$\mathcal{L}_{adv\_scale}(x_{T^\sigma}) = -\log(\mathcal{D}_{scale}(p_{T^\sigma})). \quad (8)$$

The overall objective for updating the segmentation network is the sum of the three losses,

$$\mathcal{L}(x_{S^\theta}, x_{T^\sigma}) = \mathcal{L}_{seg} + \lambda_f \mathcal{L}_{adv\_feat} + \lambda_s \mathcal{L}_{adv\_scale}. \quad (9)$$

Here,  $\lambda_s$  and  $\lambda_f$  are the hyperparameters for loss weights.

### 4.3. Scale Attention Module

The ASPP module in the DeepLab networks has demonstrated its effectiveness at handling multi-scale information for semantic segmentation. However, the input to the ASPP module is the low resolution feature maps which do not contain rich spatial information. We therefore adopt a self-attention mechanism to learn scale-enhanced feature maps for improved domain adaptation.

We develop a scale attention module (SAM) to re-weight the multi-scale concatenated features as shown in Fig. 3. The input to our SAM consists of five feature maps extracted from different layers of an encoder network composed of a DeepLabV3+ model with a ResNet101 back-

Table 1: Comparing our approach which incorporates scale adaptation to resampling the target imagery with and without standard domain adaptation. Underline indicates the test set is resampled to match the training set. mIoU: higher is better.

Method	Training set	GSD	Test set	GSD	mIoU
No domain adaptation	Potsdam	5cm	Vaihingen	9cm	32.62
No domain adaptation		5cm		<u>5cm</u>	30.85
No domain adaptation		<u>9cm</u>		9cm	31.74
Standard domain adaptation		5cm		9cm	40.74
Standard domain adaptation		5cm		<u>5cm</u>	41.77
Standard domain adaptation		<u>9cm</u>		9cm	43.09
<i>Our approach</i>		5cm		9cm	<b>47.66</b>

bone. These feature maps are the outputs of each residual group and the final ASPP module. Adaptive pooling is used to transform the maps to the same spatial dimension. They are then concatenated into a single multi-scale feature map for performing self-attention.

The concatenated feature map  $f \in \mathbb{R}^{H \times W \times C}$ , where  $H, W, C$  denote the height, width and number of channels, is first reshaped to  $\mathbb{R}^{N \times C}$ , where  $N = H \times W$ . The scale attention  $\mathcal{A}(f)$  is then computed as

$$\mathcal{A}(f) = \text{softmax}(\alpha(f)^T \beta(f)). \quad (10)$$

Here,  $\alpha$  and  $\beta$  are two  $1 \times 1$  convolutional layers and  $T$  indicates the transpose operation. The scale attention measures the impact of each channel based on various scales on all the other channels.

The final re-weighted feature map  $\mathcal{O}(f)$  is computed using the scale-based attention weights through

$$\mathcal{O}(f) = \mathcal{A}(f) \gamma(f), \quad (11)$$

where  $\gamma$  is another  $1 \times 1$  convolutional layer to transform the input. Finally, we reshape  $\mathcal{O}(f)$  back to the original dimension  $\mathbb{R}^{H \times W \times C}$  and feed it to the segmentation head for the final prediction. Note that we do not use the residual connection in our self-attention module since we only want the scale enhanced feature map.

We emphasize that our proposed SAM computes reweighted features along channels from different scales. The difference between the se-layer in Squeeze-and-Excitation Networks [25] and channel attention [21] is that the former uses only single-scale features while channel attention uses residuals to keep the original features. The goal of SAM is to enhance the features to include different scales in order to help the adversarial training when the scale discriminator is added.

## 5. Experiments

### 5.1. Datasets

We evaluate our methods on three semantic segmentation datasets, two from the ISPRS 2D Semantic Labeling Challenge [1] and a third from the DeepGlobe land cover classification challenge [15].

**ISPRS 2D Semantic Labeling Challenge** This challenge

includes two datasets, Vaihingen and Potsdam, both labeled with six classes: impervious surface, building, low vegetation, tree, car and clutter/background. The Vaihingen set contains 33 image tiles with size  $2494 \times 2064$ , of which 16 are fully annotated with class labels. The spatial resolution is 9 cm. We select five images for validation (IDs: 11, 15, 28, 30 and 34) and the remaining 11 for training, following [36, 41]. The Potsdam set contains 38 tiles with size  $6000 \times 6000$ , of which 24 are annotated. The spatial resolution is 5cm. We select seven images for validation (IDs: 2\_11, 2\_12, 4\_10, 5\_11, 6\_7, 7\_8 and 7\_10) and the remaining 17 for training, again following [36, 41].

**DeepGlobe Land Cover Classification Challenge** This challenge introduces the first public dataset offering high-resolution sub-meter satellite imagery focusing on rural areas [15]. It contains 1146 satellite images of size  $2448 \times 2448$ , split into training/validation/test sets with 803/171/172 images. The images are from the DigitalGlobe Vivid+ dataset and have a pixel resolution of 50 cm. The classes include urban, agriculture, rangeland, forest, water, barren and unknown. The DeepGlobe dataset is more challenging due to its large coverage and dense annotations.

### 5.2. Implementation Details

We implement our framework using the PyTorch toolbox [40] on a Titan V GPU with 12 GB memory.

**Segmentation network** We choose the state-of-the-art semantic segmentation network DeepLabV3+ [10] with a ResNet101 backbone as our model architecture. The segmentation network  $\mathcal{G}$  is trained using a Stochastic Gradient Descent (SGD) optimizer with Nesterov acceleration where the momentum is 0.9 and the weight decay is  $10^{-4}$ . The initial learning rate is set to  $2.5 \times 10^{-4}$  and is decayed using a polynomial decay policy with a power of 0.9.

**Adversarial discriminator** We design our scale and feature discriminators using a fully convolutional network architecture following [43]. The discriminator consists of 5 convolutional layers with  $4 \times 4$  kernels and a stride of 2 as well as a padding of 1. Each convolutional layer is followed by a leaky\_ReLU activation unit with a negative slope of 0.2. The channel number for each layer is set to 64, 128, 256, 512 and 1. The input to both discriminators is the predicted

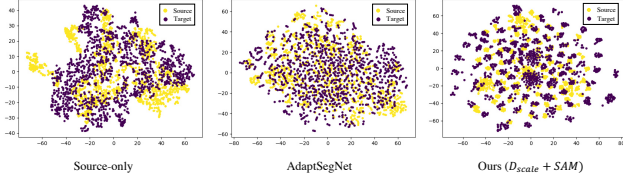


Figure 4: Visualization of the feature distributions using t-SNE. From left to right: before adaptation, domain adaptation NOT considering scale, domain adaptation considering scale. Source domain: ISPRS Potsdam, target domain: ISPRS Vaihingen. Our approach results in distributions that are more aligned and distinct. t-SNE hyperparameters are consistent between visualizations. (Zoom in to view details)

segmentation maps. To train the discriminators, we use the Adam optimizer [28] with an initial learning rate of  $10^{-4}$  and default momentum. We adopt the same polynomial decay policy as with training the segmentation network. We set the adversarial loss weights,  $\lambda_f$  and  $\lambda_s$ , to 0.005.

**Evaluation metrics** We use segmentation metrics to evaluate land cover classification performance. In particular, we compute the per class Intersection over Union (IoU) and the mean over all classes (mIoU) as percentages (%) where higher is better. We also compute the IoU gap with a model trained using the labeled target images where lower is better. This model serves as the oracle and can be considered an upper limit on the performance.

## 6. Experimental analysis

We evaluate the proposed framework as follows. We first compare our approach to standard methods for addressing scale mismatch such as image resampling and data augmentation (details in the supplementary material). This confirms our motivation. We then perform an ablation study to measure the contributions of the components of our framework. Finally, we compare our method to state-of-the-art to domain adaptation approaches.

### 6.1. General baselines

**Scale adaptation (ours) v.s. image resampling** A simple yet reasonable approach to address scale mismatch between source and target domains is simply to resize the resolution of one to match the other. We experimented with this using bilinear interpolation. (This assumes that the resolution of each is known which is our case.) It turns out, perhaps surprisingly, that this is not effective if there are other domain shifts such as location. Table 1 compares the results between resampling and our proposed method. The second and third rows show the results of training on 5cm Potsdam imagery and then testing on 9cm and 5cm (resampled) Vaihingen. The fourth through sixth rows incorporate standard non-scale aware domain adaptation, AdaptSegNet [43]. The results show some improvement, but they are still inferior to our method shown on row seven which achieves an mIoU

Table 2: Ablation study of adapting from Vaihingen scale-1 to Vaihingen scale-1/2.  $\mathcal{D}_{feat}$  and  $\mathcal{D}_{scale}$  indicate the feature and scale discriminators. SAM is the scale attention module. w/o DA means no domain adaptation.

Vaihingen scale-1 $\rightarrow$ scale-1/2					
Method	$\mathcal{D}_{feat}$	$\mathcal{D}_{scale}$	SAM	mIoU	IoU gap
w/o DA				50.21	17.43
			✓	<b>53.75</b>	13.89
w/ DA	✓			52.37	15.27
	✓	✓		54.75	12.90
	✓		✓	56.02	11.62
	✓	✓	✓	<b>57.29</b>	10.35
scale-1/2				67.54	0

of 47.66% *without resampling the target dataset*. This confirms that scale and other sources of domain mismatch are not necessarily independent and should be handled jointly. Similar results for Potsdam  $\rightarrow$  Vaihingen can be found in the supplementary materials.

**Impact of  $\mathcal{D}_{scale}$**  We further analyze the proposed  $\mathcal{D}_{scale}$  by visualizing the feature representations using t-SNE [35]. As shown in Fig. 4 (left), given two datasets with different scales, the high-dimensional feature representations from the encoder are not aligned when there is no adaptation. The source-only model shows strong clustering of the source domain, but fails to have similar influence on the target domain. After applying domain adaptation i.e. AdaptSegNet (middle), the feature representations are projected into an overlapping space. This is due to the effectiveness of adversarial learning which forces the two distributions to be aligned. However, the class clusters are barely visible in the target domain and so the model has difficulty identifying the class boundaries. With our adaptation (right), the features are not only aligned but the class clusters are distinct. This further supports our claim that scale and other sources of domain mismatch are not necessarily independent and are best handled together.

To show the effectiveness of the proposed approach, we consider a simple scenario where two datasets are from the same location but with different scales. We 1) investigate how well standard domain adaptation can adapt for scale change even in isolation of location change, and 2) perform an ablation study of our framework. To do this, we use bilinear interpolation to resample images from ISPRS Vaihingen at 1/2-scale. The original images *Vaihingen scale-1* have a GSD of 9cm and serve as the source dataset. The resampled images *Vaihingen scale-1/2* have a GSD of 18cm and serve as the target. Table 2 shows that standard domain adaptation with only a feature discriminator improves by 2.16% mIoU over a non-adaptation baseline (50.21%  $\rightarrow$  52.37%). Oracle performance, where the model is trained on the resampled images *Vaihingen scale-1/2*, is 67.54%. The 15.27% IoU gap between standard domain adaptation and the or-



Table 3: Bi-directional domain adaptation: Potsdam  $\leftrightarrow$  Vaihingen. \* indicates our DeepLabV3+ implementation. mIoU: higher is better. IoU gap: lower is better. Target-only serves as the oracle representing an upper limit on the performance.

Potsdam $\rightarrow$ Vaihingen							
	Imp. Sur.	Build.	Low vege.	Tree	Car	mIoU	IoU gap
Source-only	22.85	52.57	21.56	46.72	19.39	32.62	38.41
ADDA*	42.93	50.91	27.02	30.18	10.09	32.23	38.80
CyCADA*	49.39	55.29	28.03	32.04	10.49	35.05	35.98
AdaptSegNet*	53.72	56.08	24.74	39.68	29.49	40.74	30.28
Ours ( $\mathcal{D}_{scale}$ )	55.24	58.23	26.64	<b>50.87</b>	36.45	45.49	25.54
Ours ( $\mathcal{D}_{scale}$ + SAM)	<b>55.22</b>	<b>64.46</b>	<b>31.34</b>	50.40	<b>39.86</b>	<b>47.66</b>	<b>23.37</b>
Target-only	77.75	86.32	59.81	72.81	58.44	71.03	0.00

Vaihingen $\rightarrow$ Potsdam							
	Imp. Sur.	Build.	Low vege.	Tree	Car	mIoU	IoU gap
Source-only	31.06	37.35	44.13	16.76	31.29	32.12	45.0
ADDA*	41.33	44.21	36.07	29.81	15.11	33.31	43.81
CyCADA*	39.02	42.35	35.09	27.89	10.25	30.92	46.20
AdaptSegNet*	45.81	41.97	46.08	35.35	37.42	41.33	35.79
Ours ( $\mathcal{D}_{scale}$ )	49.36	<b>47.08</b>	51.49	37.17	39.91	45.00	32.12
Ours ( $\mathcal{D}_{scale}$ + SAM)	<b>49.76</b>	46.82	<b>52.93</b>	<b>40.23</b>	<b>42.97</b>	<b>46.54</b>	<b>30.58</b>
Target-only	79.25	85.84	73.21	68.36	78.93	77.12	0.00

acle demonstrates the limited ability of standard domain adaptation to handle scale variation. Table 2 also shows the individual contributions of our proposed scale discriminator ( $\mathcal{D}_{scale}$ ) and scale attention module (SAM). Adding the scale discriminator results in an improvement of 2.38% (52.37%  $\rightarrow$  54.75%). Adding the scale attention module results in an improvement of 3.65% (52.37%  $\rightarrow$  56.02%). We find that  $\mathcal{D}_{scale}$  and SAM are complementary. Combining both results in the highest mIoU, 57.29%. We also observe that our proposed SAM, as a feature enhancement technique, provides benefits even without domain adaptation. Incorporating just SAM into the baseline results in an improvement of 3.54% mIoU (50.21%  $\rightarrow$  53.75%).

In summary, this set of experiments shows that scale is intricately linked to the features in CNNs and that scale-specific adaptation and enhancement is important and advantageous.

## 6.2. Comparison Study: Small Domain Gap

We now perform cross-location domain adaptation between Potsdam and Vaihingen. We consider this as a small domain gap scenario since both locations are in Germany and the difference in scale is moderate. We compare our approach to three recent state-of-the-art domain adaptation methods for semantic segmentation, ADDA [45], CyCADA [23] and AdaptSegNet [43]. For fair comparison, we implement our own versions (\*) with a DeepLabV3+ segmentation network. Table 3 contains the quantitative results and Fig. 5 shows the qualitative results. More qualitative results can be found in the supplementary materials.

We make several observations from Table 3. First, standard domain adaptation does not work well on the cross-

scale, cross-location problem. Take Vaihingen  $\rightarrow$  Potsdam for example. CyCADA\* performs even worse than the baseline, which is a model trained using the source dataset only without any domain adaptation. Second, our proposed scale adaptive framework achieves much higher performance in the cross-scale, cross-location scenario. With just our new scale discriminator, we improve over the previous best by 2.81% (42.19%  $\rightarrow$  45.00%). Adding the scale attention module further boosts our accuracy to 46.54%, which is 4.45% higher than AdaptSegNet\* and 15.62% higher than CyCADA\*.

**Visualization of SAM** To further show the effectiveness of SAM, we visualize in Fig. 5 the feature maps of the most weighted channel before (column 2) and after (column 3) applying SAM. We observe that SAM emphasizes feature maps at a more appropriate scale/size and results in better boundary estimates.

**Visual comparisons** Fig. 5 shows visualizations of segmentation prediction maps adapting Potsdam to Vaihingen (columns 4-7). We see that the cross-scale, cross-location segmentation problem is very challenging for RS imagery. The results without adaptation are poor. Most predictions are clutter/background (incorrect). With standard domain adaptation, AdaptSegNet\* generates improved but coarse predictions. Our framework results in more accurate predictions with crisper boundaries especially with SAM. Additional visualizations can be found in the supplementary materials.

## 6.3. Comparison Study: Large Domain Gap

We now perform cross-location domain adaptation between DeepGlobe and Vaihingen. We consider this as a

Table 4: Bi-directional domain adaptation: DeepGlobe  $\leftrightarrow$  Vaihingen. \* indicates our DeepLabV3+ implementation. mIoU: higher is better. IoU gap: lower is better. Target-only serves as the oracle representing an upper limit on the performance.

	DeepGlobe $\rightarrow$ Vaihingen				Vaihingen $\rightarrow$ DeepGlobe			
	Urban	R. land	Forest	mIoU	Urban	R. land	Forest	mIoU
Source-only	14.29	13.60	0.00	9.30	0.17	1.05	2.04	1.09
ADDA*	21.77	19.38	7.28	16.14	25.04	3.97	20.99	16.67
CyCADA*	20.38	20.39	3.48	14.75	24.37	3.29	24.30	17.32
AdaptSegNet*	22.28	17.71	5.81	15.27	26.82	3.98	31.04	20.61
Ours( $D_{scale}$ )	26.98	19.16	11.29	19.14	28.34	4.74	39.49	24.19
Ours( $D_{scale}$ +SAM)	<b>27.35</b>	<b>23.43</b>	<b>12.94</b>	<b>21.24</b>	<b>29.02</b>	<b>5.42</b>	<b>41.72</b>	<b>25.39</b>
Target-only	82.04	59.81	72.81	71.55	68.09	28.77	75.06	57.31

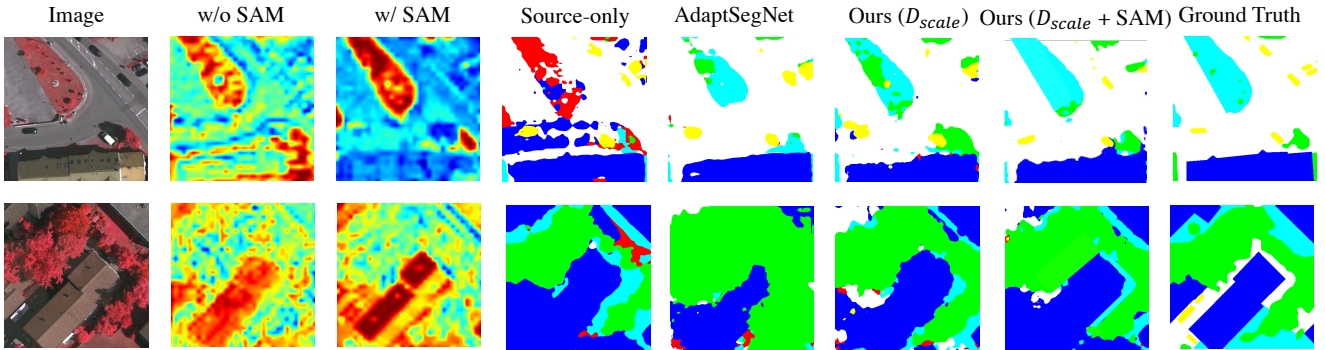


Figure 5: Visualization of Potsdam  $\rightarrow$  Vaihingen. The second and third columns are the visualizations of channel feature maps, blue: low value, red: high value. Our class predictions have crisper boundaries and are more similar to the ground truth than standard methods. White: road, blue: building, cyan: low vegetation, green: trees, yellow: cars, red: clutter.

large domain gap scenario since the location varies significantly and the difference in scale is large (GSDs of 50cm and 9cm). Also, DeepGlobe is satellite imagery while Potsdam and Vaihingen are aerial imagery which introduces increased sensor and spectrum variation. To align the Vaihingen and DeepGlobe segmentation classes, we merge impervious surface and building as urban, match low vegetation to rangeland, and match trees to forest. The results of bi-directional domain adaptation between DeepGlobe and Vaihingen are shown in Table 4. We make several observations. First, due to the large domain gap, the source-only model performs quite poorly. For example, the DeepGlobe  $\rightarrow$  Vaihingen mIoU is only 9.30%. Second, our scale adaptive framework again handles this now even more challenging scenario much better than standard domain adaptation. It achieves 5.97% higher mIoU than AdaptSegNet\* and 6.49% higher mIoU than CyCADA\*. Similar improvements can be observed for Vaihingen  $\rightarrow$  DeepGlobe.

#### 6.4. Discussion

**Limitations** One limitation of our approach is that the scale of both the source and target domains must be known (or at least the scale difference). This is usually not a problem in RS image analysis since this information is typically available. We have also investigated methods to automatically estimate the spatial resolution of RS imagery using

CNNs [30]. Another issue is that resampling images, using bilinear interpolation for example, is not equivalent to images whose scale varies due to having been acquired at different altitudes, etc. Despite this, our results show improved performance over standard domain adaptation. And, again, our primary goal is to illustrate the importance of scale in domain adaptation in RS segmentation.

## 7. Conclusion

We establish that scale is a significant factor for domain adaptation in RS image segmentation. We propose a scale adaptive adversarial learning framework to address the challenging cross-scale, cross-location problem. This framework includes a new scale discriminator that explicitly adapts for differences in scale. It also includes a scale attention module to produce scale-enhanced features. Our framework outperforms standard domain adaptation methods, achieving state-of-the-art results for small (Potsdam  $\leftrightarrow$  Vaihingen) and large (DeepGlobe  $\leftrightarrow$  Vaihingen) domain gap scenarios.

## 8. Acknowledgement

This work was funded in part by a National Science Foundation grant, #IIS-1747535. We gratefully acknowledge the support of NVIDIA Corporation through the donation of the GPU card used in this work.



## References

- [1] ISPRS 2D Semantic Labeling Challenge. <http://www2.isprs.org/commissions/comm3/wg4/semantic-labeling.html>.
- [2] N. Ammour, L. Bashmal, Y. Bazi, M. M. Al Rahhal, and M. Zuair. Asymmetric Adaptation of Deep Features for Cross-Domain Classification in Remote Sensing Imagery. *IEEE Geoscience and Remote Sensing Letters*, 2018.
- [3] N. Audebert, B. Saux, and S. Lefèvre. Beyond RGB: Very High Resolution Urban Remote Sensing with Multimodal Deep Networks. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2018.
- [4] S. Azimi, C. Henry, L. Sommer, A. Schumann, and E. Vig. SkyScapes Fine-Grained Semantic Understanding of Aerial Scenes. In *International Conference on Computer Vision (ICCV)*, 2019.
- [5] J. E. Ball, D. T. Anderson, and C. S. Chan. Comprehensive Survey of Deep Learning in 6 Remote Sensing: Theories, Tools, and Challenges for the Community. *Journal of Applied Remote Sensing*, 2017.
- [6] L. Bashmal, Y. Bazi, H. AlHichri, M. AlRahhal, N. Ammour, and N. Alajlan. Siamese-GAN: Learning Invariant Representations for Aerial Vehicle Image Categorization. *Remote Sensing*, 2018.
- [7] B. Benjdira, Y. Bazi, A. Koubaa, and K. Ouni. Unsupervised Domain Adaptation Using Generative Adversarial Networks for Semantic Segmentation of Aerial Images. *Remote Sensing*, 2019.
- [8] M. Bosch, G. Christie, and C. Gifford. Sensor Adaptation for Improved Semantic Segmentation of Overhead Imagery. In *Winter Conference on Applications of Computer Vision (WACV)*, 2019.
- [9] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2017.
- [10] L. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [11] X. Chen, S. Xiang, C. Liu, and C. Pan. Vehicle Detection in Satellite Images by Hybrid Deep Convolutional Neural Networks. *IEEE Geoscience and Remote Sensing Letters*, 2014.
- [12] Y. Chen, W. Chen, Y. Chen, B. Tsai, Y. Wang, and M. Sun. No More Discrimination: Cross City Adaptation of Road Scene Segmenters. In *International Conference on Computer Vision (ICCV)*, 2017.
- [13] Y. Chen, W. Li, and L. Van Gool. ROAD: Reality Oriented Adaptation for Semantic Segmentation of Urban Scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [14] Y. Chen, Y. Lin, M. Yang, and J. Huang. CrDoCo: Pixel-Level Domain Transfer With Cross-Domain Consistency. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] I. Demir, K. Koperski, D. Lindenbaum, G. Pang, J. Huang, S. Basu, F. Hughes, D. Tuia, and R. Raska. DeepGlobe 2018: A Challenge to Parse the Earth through Satellite Images. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018.
- [16] X. Deng, H. L. Yang, N. Makkar, and D. Lunga. Large Scale Unsupervised Domain Adaptation of Segmentation Networks with Adversarial Learning. In *2019 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2019.
- [17] X. Deng, Y. Zhu, and S. Newsam. What Is It Like Down There? Generating Dense Ground-Level Views and Image Features from Overhead Imagery Using Conditional Generative Adversarial Networks. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2018.
- [18] X. Deng, Y. Zhu, Y. Tian, and S. Newsam. Generalizing Deep Models for Overhead Image Segmentation Through Getis-Ord Gi\* Pooling. *arXiv preprint arXiv:1912.10667*, 2019.
- [19] H. Ding, X. Jiang, B. Shuai, A. Qun Liu, and G. Wang. Context Contrasted Feature and Gated Multi-Scale Aggregation for Scene Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [20] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu. Dual Attention Network for Scene Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [21] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu. Dual Attention Network for Scene Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [22] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-Adversarial Training of Neural Networks. *The Journal of Machine Learning Research (JMLR)*, 2016.
- [23] J. Hoffman, E. Tzeng, T. Park, J. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *International Conference on Machine Learning (ICML)*, 2018.
- [24] J. Hoffman, D. Wang, F. Yu, and T. Darrell. FCNs in the Wild: Pixel-Level Adversarial and Constraint-based Adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- [25] J. Hu, L. Shen, and G. Sun. Squeeze-and-Excitation Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [26] X. Jin, C. Lan, W. Zeng, Z. Zhang, and Z. Chen. CaseNet: Content-Adaptive Scale Interaction Networks for Scene Parsing. *arXiv preprint arXiv:1904.08170*, 2019.
- [27] M. Kim and H. Byun. Learning Texture Invariant Representation for Domain Adaptation of Semantic Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [28] D. P Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [29] N. Kussul, M. Lavreniuk, S. Skakun, and A. Shelestov. Deep Learning Classification of Land Cover and Crop Types Us-

- ing Remote Sensing Data. *IEEE Geoscience and Remote Sensing Letters*, 2017.
- [30] H. Liang and S. Newsam. Estimating the Spatial Resolution of Overhead Imagery Using Convolutional Neural Networks. In *International Conference on Image Processing (ICIP)*, 2019.
  - [31] D. Lin, Y. Ji, D. Lischinski, D. Cohen-Or, and H. Huang. Multi-Scale Context Intertwining for Semantic Segmentation. In *European Conference on Computer Vision (ECCV)*, 2018.
  - [32] Y. Liu, B. Fan, L. Wang, J. Bai, S. Xiang, and C. Pan. Semantic Labeling in Very High Resolution Images via a Self-Cascaded Convolutional Neural Network. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2018.
  - [33] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
  - [34] X. Ma, Z. Dai, Z. He, J. Ma, Y. Wang, and Y. Wang. Learning Traffic as Images: A Deep Convolutional Neural Network for Large-Scale Transportation Network Speed Prediction. *Sensors*, 2017.
  - [35] L. van der Maaten and G. Hinton. Visualizing Data Using t-SNE. *Journal of Machine Learning Research (JMLR)*, 2008.
  - [36] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez. High-Resolution Aerial Image Labeling with Convolutional Neural Networks. *IEEE Transactions on Geoscience and Remote Sensing*, 2017.
  - [37] V. Mnih and G. E. Hinton. Learning to Detect Roads in High-Resolution Aerial Images. In *European Conference on Computer Vision (ECCV)*, 2010.
  - [38] L. Mou, Y. Hua, and X. X. Zhu. A Relation-Augmented Fully Convolutional Network for Semantic Segmentation in Aerial Scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
  - [39] Z. Murez, S. Kolouri, D. Kriegman, R. Ramamoorthi, and K. Kim. Image to Image Translation for Domain Adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
  - [40] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic Differentiation in PyTorch. In *Advances in Neural Information Processing Systems (NeurIPS) Autodiff Workshop*, 2017.
  - [41] J. Sherrah. Fully Convolutional Networks for Dense Semantic Labelling of High-Resolution Aerial Imagery. *arXiv preprint arXiv:1606.02585*, 2016.
  - [42] Y. Tian, X. Deng, Y. Zhu, and S. Newsam. Cross-Time and Orientation-Invariant Overhead Image Geolocalization Using Deep Local Features. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020.
  - [43] Y. Tsai, W. Hung, S. Schuler, K. Sohn, M. Yang, and M. Chandraker. Learning to Adapt Structured Output Space for Semantic Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
  - [44] Y. Tsai, K. Sohn, S. Schuler, and M. Chandraker. Domain Adaptation for Structured Output via Discriminative Representations. In *International Conference on Computer Vision (ICCV)*, 2019.
  - [45] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial Discriminative Domain Adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
  - [46] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
  - [47] M. Volpi, G. Camps-Valls, and D. Tuia. Spectral Alignment of Multi-Temporal Cross-Sensor Images with Automated Kernel Canonical Correlation Analysis. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2015.
  - [48] T. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez. DADA: Depth-aware Domain Adaptation in Semantic Segmentation. In *International Conference on Computer Vision (ICCV)*, 2019.
  - [49] X. Wang, R. Girshick, A. Gupta, and K. He. Non-Local Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
  - [50] L. Yan, B. Fan, S. Xiang, and C. Pan. Adversarial Domain Adaptation with a Domain Similarity Discriminator for Semantic Segmentation of Urban Areas. In *International Conference on Image Processing (ICIP)*, 2018.
  - [51] J. Yuan. Learning Building Extraction in Aerial Scenes with Convolutional Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2017.
  - [52] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. Self-Attention Generative Adversarial Networks. In *International Conference on Machine Learning (ICML)*, 2019.
  - [53] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu. Image Super-Resolution Using Very Deep Residual Channel Attention Networks. In *European Conference on Computer Vision (ECCV)*, 2018.
  - [54] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid Scene Parsing Network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
  - [55] L. Zhou, C. Zhang, and M. Wu. D-LinkNet: LinkNet With Pretrained Encoder and Dilated Convolution for High Resolution Satellite Imagery Road Extraction. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018.
  - [56] Y. Zhu, K. Sapra, F. A. Reda, K. J. Shih, S. Newsam, A. Tao, and B. Catanzaro. Improving Semantic Segmentation via Video Propagation and Label Relaxation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
  - [57] Z. Zhu, M. Xu, S. Bai, T. Huang, and X. Bai. Asymmetric Non-local Neural Networks for Semantic Segmentation. In *International Conference on Computer Vision (ICCV)*, 2019.
  - [58] Y. Zou, Z. Yu, X. Liu, B. Kumar, and J. Wang. Confidence Regularized Self-Training. In *International Conference on Computer Vision (ICCV)*, 2019.