

# Geometry-Aware Network for Domain Adaptive Semantic Segmentation

Yinghong Liao<sup>1,2\*</sup>, Wending Zhou<sup>1,2\*</sup>, Xu Yan<sup>1,2</sup>, Zhen Li<sup>2,1†</sup>, Yizhou Yu<sup>3</sup>, Shuguang Cui<sup>2,1</sup>

<sup>1</sup> The Future Network of Intelligence Institute, The Chinese University of Hong Kong (Shenzhen)

<sup>2</sup> School of Science and Engineering, The Chinese University of Hong Kong (Shenzhen)

<sup>3</sup> Department of Computer Science, The University of Hong Kong

{yinghongliao@link., lizhen@}cuhk.edu.cn

## Abstract

Measuring and alleviating the discrepancies between the synthetic (source) and real scene (target) data is the core issue for domain adaptive semantic segmentation. Though recent works have introduced depth information in the source domain to reinforce the geometric and semantic knowledge transfer, they cannot extract the intrinsic 3D information of objects, including positions and shapes, merely based on 2D estimated depth. In this work, we propose a novel Geometry-Aware Network for Domain Adaptation (GANDA), leveraging more compact 3D geometric point cloud representations to shrink the domain gaps. In particular, we first utilize the auxiliary depth supervision from the source domain to obtain the depth prediction in the target domain to accomplish structure-texture disentanglement. Beyond depth estimation, we explicitly exploit 3D topology on the point clouds generated from RGB-D images for further coordinate-color disentanglement and pseudo-labels refinement in the target domain. Moreover, to improve the 2D classifier in the target domain, we perform domain-invariant geometric adaptation from source to target and unify the 2D semantic and 3D geometric segmentation results in two domains. Note that our GANDA is plug-and-play in any existing UDA framework. Qualitative and quantitative results demonstrate that our model outperforms state-of-the-arts on GTA5→Cityscapes and SYNTHIA→Cityscapes.

## Introduction

Semantic segmentation has obtained admirable progress owing to the success of deep learning in computer vision. However, the pixel-level annotations on large-scale image data require time-consuming and laborious manual work. To overcome this challenge, the synthetic datasets whose photo-realistic images and annotations are automatically generated via computer graphics, *e.g.*, GTA5 (Richter et al. 2016) and SYNTHIA (Ros et al. 2016), are leveraged to explore training on the synthetic images instead of real-scene ones (synthetic-to-real). Nevertheless, the model trained on synthetic data exhibits poor generalization on real-world images, *i.e.*, the domain discrepancy occurs. The discrepancy

\*These authors contributed equally.

†Corresponding author.

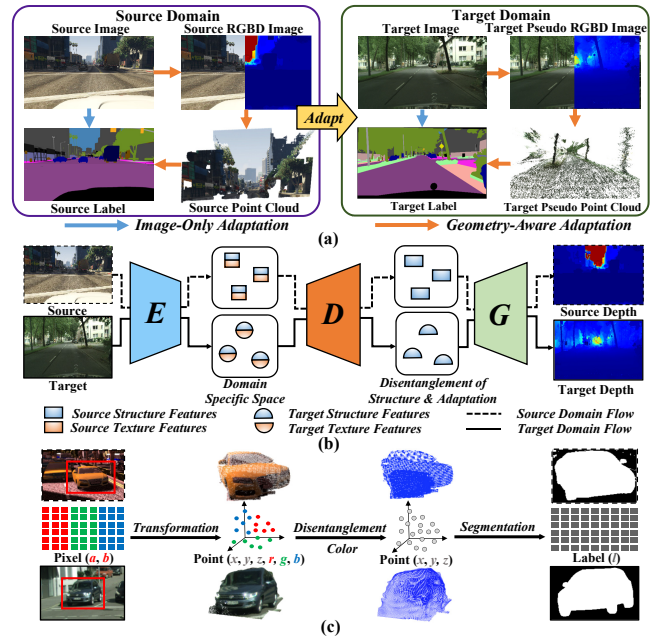


Figure 1: (a) Illustrations of the traditional image-only adaptation and our proposed geometry-aware adaptation. Geometry-aware adaptation only utilizes the depths in the source domain and leverages the domain-invariant geometric information in the point cloud transformed from RGB-D images. (b) The disentanglement of structural features in the depth-aware domain adaptation. (c) The disentanglement of spatial coordinates and RGB colors in the 3D points transformed from 2D pixels. Zoom in for better view.

any lies in the differences between the synthetic and real-scene data, *e.g.*, the contrast, lighting, object shape, and surface textures. Therefore, an effective Unsupervised Domain Adaptation (UDA) model, where the target annotations are not provided, is needed.

Instead of aligning the domain distributions that may generate sub-optimal solutions (Yang et al. 2020b; Vu et al. 2019a), the recent UDA works (Zhang et al. 2021; Hoyer, Dai, and Van Gool 2022a,b) adopt self-training to exploit the target information: first predicting pseudo labels based on

the trained parameters from the annotated data and then rectifying the pseudo labels via the target knowledge. Besides, depth-based multi-task learning (Vu et al. 2019b; Wang et al. 2021) also has demonstrated the advantages in utilizing the target image features. Although the self-training and depth-based methods have achieved significant progress in domain adaptive semantic segmentation, they still have several issues: **1)** The predicted low-confidence pseudo labels have uncertainty in the patches of the object edges and might lead to assigning incorrect labels to these edge pixels, especially for the small objects. **2)** Despite the effectiveness in extracting structure knowledge, auxiliary depth estimation cannot accurately manifest the object spatial positions and the inherent 3D geometric object shapes, leading to the unsatisfying performances in the small objects. **3)** The depth-based methods (Wang et al. 2021) obtain the depths for the target Cityscapes images via the self-supervised depth estimation and utilizes them as the ground truth to supervise the multi-task training. However, when the target dataset consists of only the images, the target depths are not available to perform a similar self-supervised depth estimation.

To address the above issues, in this work, we propose to incorporate geometric information into the prevalent image-only UDA frameworks and learn a complete 3D geometric object shape with the only source depth information, as shown in Figure 1(a). Since the depth information is accessible in the generation of synthetic images, we consider leveraging the source depths to estimate the target depths. As displayed in Figure 1(b), the generation of depths can be seen as disentanglement of the object structures from the RGB images without colors and textures (Chang et al. 2019). Thus, the depth information from the source and target domains shares more domain-invariant geometric structure knowledge and avoids influences from the domain-specific textures. Furthermore, with the 2D CNNs, the mere depth estimation is insufficient to exploit the implicit 3D geometric shape knowledge in RGB and Depth (RGB-D) images. To facilitate the adaptation of the geometric information, we convert the RGB-D images to more compact 3D scenes in point clouds. With point clouds and 3D point-based networks, the 3D geometric topology on the source and target domains can be exploited, including the explicit object positions and distinctive object shapes. As shown in Figure 1(c), the 3D coordinates and colors can be disentangled, and the 3D spatial coordinates that represent the domain-invariant object shapes can be explored independently.

Accordingly, we propose a novel plug-and-play **Geometry-Aware Network for Domain Adaptation (GANDA)** for the UDA segmentation task. GANDA consists of two stages: depth-aware adversarial adaptation and geometry-aware adaptation. In the depth-aware adversarial adaptation stage, GANDA first accomplishes the depth estimation on the annotated source images, then predicts the target depths based on the learned parameters. With the obtained target pseudo depths, the target RGB-D images are generated, and the corresponding point clouds are reconstructed. In the geometry-aware adaptation stage, GANDA leverages 3D point clouds to perform 2D semantic and 3D geometric segmentation simultaneously. Specifi-

cally, we design a Geometric Adaptation Module to align the domain-invariant geometric information on 3D point clouds from source to target. The target predicted labels are enhanced with the geometric information by unifying the semantic and geometric segmentation results on two domains. Geometry-aware adaptation is especially effective in segmentation of the small-object classes. The main contributions of this paper are as follows:

- We propose GANDA, a novel plug-and-play UDA segmentation framework to leverage the complete 3D geometric representations with only source depth image. This model introduces domain-invariant geometric information based on the structure-texture and coordinate-color disentanglements.
- The Geometry-Aware Adaptation stage is introduced to align the domain-invariant geometric knowledge from source to target on the reconstructed point clouds from RGB-D images.
- The proposed GANDA outperforms the state-of-the-art methods on the classic synthetic-to-real segmentation tasks, including GTA5  $\rightarrow$  Cityscapes and SYNTHIA  $\rightarrow$  Cityscapes.

## Related Work

**Unsupervised Domain Adaptation in Segmentation.** Appropriate measurement of the domain gap and accurate disentanglement of the domain-invariant knowledge are the keys to effective Unsupervised Domain Adaptation (UDA) in the synthetic-to-real semantic segmentation task. Accordingly, two mainstream approaches come to stage: domain alignment and knowledge extraction from the target domain. Based on adversarial learning, the first commonly used methods perform domain alignment to find domain-invariant features from different levels, including image level (Chen et al. 2019b; Ma et al. 2021), feature level (Nuriel, Benaim, and Wolf 2021), and semantic level (Zhang et al. 2021). Though effective in narrowing the domain gap, these works fail to achieve better segmentation results since they focus on the information from source domain without further mining the knowledge from target domain. By comparison, the latter ones (Yang et al. 2020a; Zhang et al. 2021; Hoyer, Dai, and Van Gool 2022b) seizes the domain-specific knowledge from target domain to enhance the pseudo labels.

**Depth-based Multi-task Learning in UDA.** Some researchers (Lee et al. 2019; Vu et al. 2019b) resort to the auxiliary tasks to reinforce knowledge mining and transfer, *e.g.*, depth estimation. DADA (Vu et al. 2019b) introduces an auxiliary depth regression task to obtain the dense depth knowledge from the source domain to facilitate the adaptation. GIO-Ada (Chen et al. 2019a) proposes to utilize the auxiliary geometric information to ease the domain gap on both the input and output level. However, these methods are based on adversarial learning and their performances are limited. Thus the advantages of appending auxiliary information or tasks are not prominent. Very recently, CorDA (Wang et al. 2021) uses the depth information on both domains to supervise the training of two depth decoders. Though CorDA outperforms the previous works that

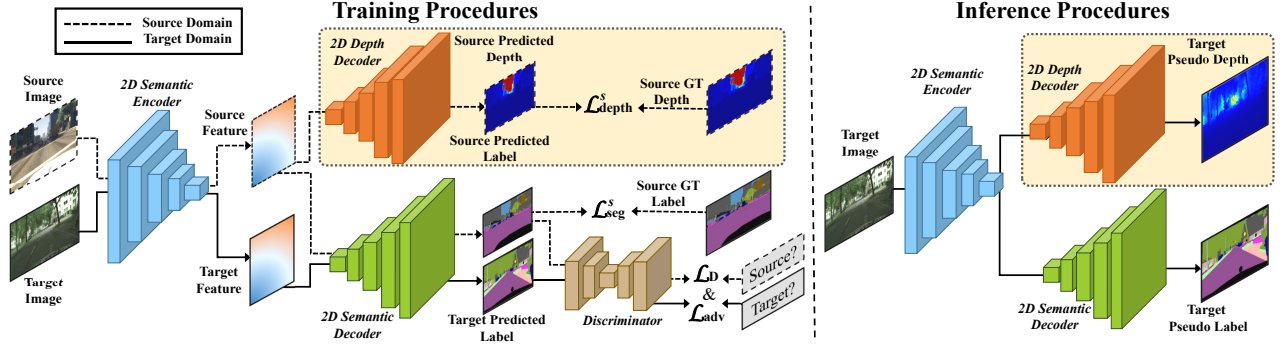


Figure 2: Depth-Aware Adversarial Adaptation. The proposed plug-and-play depth-aware adaptation modules are highlighted in rounded rectangles with dotted lines.

employ depth estimation significantly, it allows the supervision from target depth maps that might be inaccessible.

**3D Point Cloud Segmentation.** The methods of 3D point cloud segmentation are primarily divided into two categories: point-based methods (Qi et al. 2017a,b; Thomas et al. 2019; Yan et al. 2020) and voxel-based methods (Wu et al. 2015; Choy, Gwak, and Savarese 2019; Yan et al. 2021; Tang et al. 2020; Yan et al. 2022; Zhu et al. 2021). In the point-based methods, PointNet (Qi et al. 2017a) and its improved PointNet++ (Qi et al. 2017b) are the pioneers of applying permutation-invariant set functions on the raw 3D geometric points based on the MLP networks. Early voxel-based methods (Wu et al. 2015; Choy, Gwak, and Savarese 2019) convert the points to voxels and utilize 3D convolutions to process the voxels. Recently, Cylinder3D (Zhu et al. 2021) proposes an asymmetrical residual block to reduce computational costs of 3D convolutions. Since the point clouds generated by RGB-D images are single-view and have more irregular shapes than the usual point clouds, our model applies the point-based model to extract 3D geometric information from point clouds.

## Our Approach

In this section, we propose GANDA for domain adaptive semantic segmentation, which applies a two-stage self-training pipeline, including depth-aware adversarial adaptation and geometry-aware adaptation.

### Preliminaries

In the UDA segmentation task, we denote the source dataset containing  $n^s$  images as  $\mathbf{X}^s = \{\mathbf{x}^s \in \mathbb{R}^{3 \times H \times W}\}_{i=1}^{n^s}$  and the corresponding semantic labels as  $\mathbf{Y}^s = \{\mathbf{y}^s \in \{0, 1\}^{n_c \times H \times W}\}_{i=1}^{n^s}$ , where each image  $\mathbf{x}^s$  and label  $\mathbf{y}^s$  share the same height of  $H$  and the width of  $W$ , and  $n_c$  represents the number of classes. Accordingly, the target dataset has  $n^t$  images,  $\mathbf{X}^t = \{\mathbf{x}^t\}_{j=1}^{n^t}$ . Since the only supervision from the source domain is limited, self-training (Pan et al. 2020) is exploited to handle the unlabeled dataset in the semi-supervised learning and UDA tasks. After training multiple steps on the source domain, the model is able to produce the coarse pseudo labels on the target domain,

$\hat{\mathbf{Y}}^t = \{\hat{\mathbf{y}}^t\}_{j=1}^{n^t} = \{\arg \max \mathbf{p}^t\}_{j=1}^{n^t}$ , where  $\mathbf{p}^t$  is the softmax probability of classes for the target image  $\mathbf{x}^t$ . The target pseudo label  $\hat{\mathbf{y}}^t$  for  $\mathbf{x}^t$  is obtained from the class that has the highest soft confidence scores. With the source label  $\mathbf{y}^s$  and the target pseudo label  $\hat{\mathbf{y}}^t$ , the Cross-Entropy (CE) loss in self-training segmentation is calculated as:

$$\mathcal{L}_{\text{seg}} = - \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \mathbf{y}_{ik}^s \log \mathbf{p}_{ik}^s - \sum_{j=1}^{HW} \sum_{k=1}^{n_c} \hat{\mathbf{y}}_{jk}^t \log \mathbf{p}_{jk}^t. \quad (1)$$

### Depth-Aware Adversarial Adaptation

In this work, other than the image-only adaptation, we take the accessible depth maps in the source domain into account. Specifically, as shown in Fig 1(b), our GANDA first learns to disentangle the structural information from single source RGB images  $\mathbf{X}^s$  under the supervision of source depths  $\mathbf{D}^s = \{\mathbf{d}^s \in \mathbb{R}^{H \times W}\}_{i=1}^{n^s}$ . Then it applies the learned depth knowledge to generate pseudo depth maps  $\hat{\mathbf{D}}^t = \{\hat{\mathbf{d}}^t \in \mathbb{R}^{H \times W}\}_{j=1}^{n^t}$  on the raw target images  $\mathbf{X}^t$ . In the training part of the depth-aware adversarial adaptation, as displayed in Figure 2, we incorporate the depth-aware knowledge transfer into the classic adversarial UDA.

The depth-aware adversarial adaptation adopts adversarial learning that 2D semantic encoder is optimized to “fool” the discriminator with the help of source depth estimation:

$$\mathcal{L}_{\text{D}} = - \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \mathbf{b}_{ik}^s \log P(\mathbf{D}_{ik} = 0) - \sum_{j=1}^{HW} \sum_{k=1}^{n_c} \mathbf{b}_{jk}^t \log P(\mathbf{D}_{jk} = 1), \quad (2)$$

$$\mathcal{L}_{\text{adv}} = - \sum_{j=1}^{HW} \sum_{k=1}^{n_c} \mathbf{b}_{jk}^t \log P(\mathbf{D}_{jk} = 0) \quad (3)$$

where  $\mathbf{b}_k^s$  and  $\mathbf{b}_k^t$  are the binary indicator (source or target) of class  $k$ ,  $P(\mathbf{D}_k = 0)$  and  $P(\mathbf{D}_k = 1)$  represent the predicted probability of class  $k$  from the source or target domain, respectively.

The 2D depth decoder is trained simultaneously using the depth loss  $\mathcal{L}_{\text{depth}}^s$  between the predicted depth map  $\tilde{\mathbf{d}}^s$  and the source depth map  $\mathbf{d}^s$ :

$$\mathcal{L}_{\text{depth}}^s = \mathcal{L}_{\text{SILog}}^s + \lambda \mathcal{L}_{\text{grad}}^s, \quad (4)$$

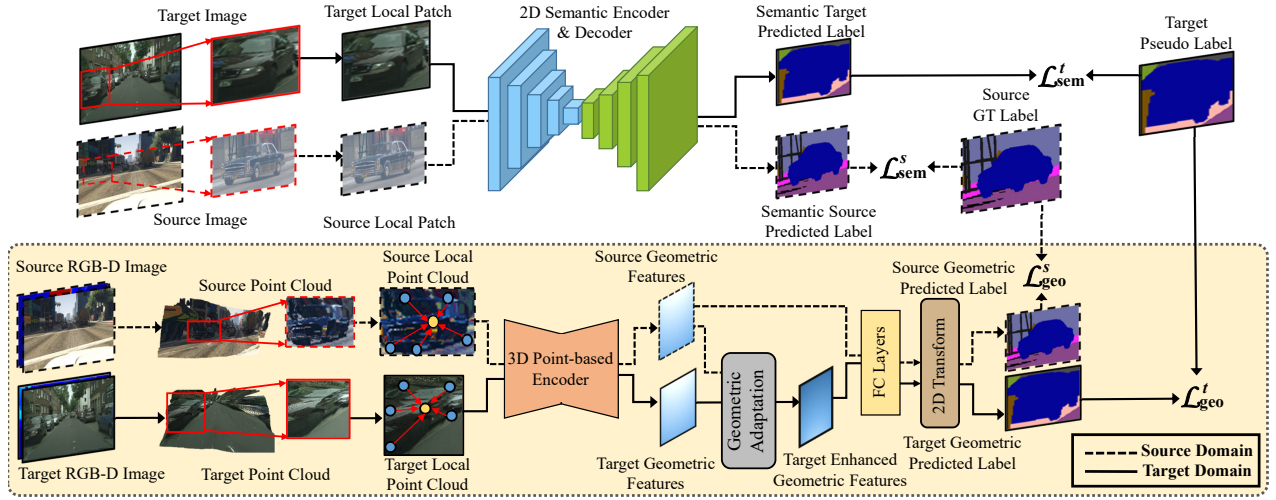


Figure 3: Geometry-Aware Adaptive Training Scheme. The proposed plug-and-play geometry-Aware adaptation pipeline is highlighted in the dotted rounded rectangle.

where  $\lambda$  is the loss weight.  $\mathcal{L}_{\text{SILog}}^s$  is the Scale-Invariant Logarithmic loss for depth estimation (Yuan et al. 2022):

$$\Delta \mathbf{d}_i^s = \log \tilde{\mathbf{d}}_i^s - \log \mathbf{d}_i^s,$$

$$\mathcal{L}_{\text{SILog}}^s = \psi \sqrt{\frac{1}{K} \sum_{i=1}^{HW} \Delta \mathbf{d}_i^s{}^2 - \frac{\gamma}{K} \left( \sum_{i=1}^{HW} \Delta \mathbf{d}_i^s \right)^2}, \quad (5)$$

where  $\psi$  is the scale constant set as 10 and  $\gamma$  denotes the variance minimizing factor set as 0.85.  $\mathcal{L}_{\text{grad}}$  is the smoothness loss with the image gradients from two directions,  $g_x$  and  $g_y$ . We train the whole stage using the following total loss with the weight  $\mu$ :

$$\mathcal{L}_{\text{depth-aware}} = \mathcal{L}_{\text{seg}}^s + \mathcal{L}_{\text{depth}}^s + \mu \mathcal{L}_{\text{adv}}. \quad (6)$$

In the inference part in Figure 2, the target pseudo label  $\hat{\mathbf{y}}^t$  and pseudo depth map  $\hat{\mathbf{d}}^t$  are generated for the following self-training based geometry-aware adaptation.

### Geometry-Aware Adaptation

**From 2D RGB Pixels to 3D Geometric Points.** Given the source map  $\mathbf{d}^s$  and the target pseudo depth map  $\hat{\mathbf{d}}^t$ , we transform a depth map  $\mathbf{d}$  and the RGB image  $\mathbf{X}$  to an RGB-D image and further produce the corresponding point cloud  $\mathbf{V} = \{\mathbf{v} \in \mathbb{R}^6\}_{k=1}^{n_v}$  for each image, where  $n_v$  represents the number of the points  $\mathbf{v}$ . According to the mapping relationship between the RGB pixels and the 3D points,  $n_v = HW$ , and each point corresponds to one pixel in the original image  $\mathbf{X}$ . In 3D point cloud processing, a point  $\mathbf{v}$  consists of the 3D spatial coordinate  $(x, y, z)$  and the RGB color information  $(r, g, b)$  as illustrated in Figure 1(c),  $\mathbf{V} = \{\mathbf{v} \in \mathbb{R}^6\}_{k=1}^{n_v} = \{(x, y, z, r, g, b) \in \mathbb{R}^6\}_{k=1}^{n_v}$ . The geometric coordinate  $(x, y, z)$  is seen as the domain-invariant information of the object and disentangled from  $\mathbf{v}$ .

**Disentanglement of the 3D Geometric Coordinates and RGB Colors.** In the geometry-aware adaptation, as displayed in Figure 3, we discard the trained 2D depth decoder in the previous stage and integrate a 3D point-based

encoder that performs disentanglement on the spatial geometric features. Initially, the input point  $\mathbf{v}$  is decomposed into the geometric coordinates  $\mathbf{h} = (x, y, z)$  and the point feature  $\mathbf{v} = (x, y, z, r, g, b)$ , which isolates the spatial positional information individually and also preserves the RGB color information. The 3D point-based encoder  $\Phi$  takes the local regions of points as the inputs, including geometric coordinates  $\mathbf{H} = \{\mathbf{h}\}_{k=1}^{n_v}$  and point features  $\mathbf{V} = \{\mathbf{v}\}_{k=1}^{n_v}$ , and obtains geometry-aware features:

$$\mathbf{F}_{\text{geo}} = \Phi(\mathbf{H}, \mathbf{V}) = \{\Phi(\mathbf{h}, \mathbf{v})\}_{k=1}^{n_v} = \{\mathbf{f}_{\text{geo}} \in \mathbb{R}^m\}_{k=1}^{n_v}, \quad (7)$$

where  $m$  is the size of geometric feature  $\mathbf{f}_{\text{geo}}$ . Note that all the operations in the 3D point-based encoder are explicitly performed on the 3D spatial coordinates  $\mathbf{H}$ , which guarantees that the features  $\mathbf{F}_{\text{geo}}$  maintain the critical 3D geometric information to the largest extent.

**Geometry-Aware Adaptive Training Scheme.** Our proposed GANDA is plug-and-play, and the Geometry-Aware Adaptive Training (GAAT) scheme can be appended to the original 2D adaptive segmentation scheme in parallel. As displayed in Figure 3, the GAAT performs adaptive segmentation on the cross-modal inputs. In the 2D adaptive part, the input source and target images are randomly cropped to produce local patches  $\mathbf{x}_{\text{local}}^s$  and  $\mathbf{x}_{\text{local}}^t$  with a smaller size. Accordingly, in the 3D adaptive part, the cropping operation is performed on the same patch of the RGB-D images, then the local point clouds  $\mathbf{V}_{\text{local}}^s$  and  $\mathbf{V}_{\text{local}}^t$  are obtained after the pixel-to-point transformation. The local patches and point clouds  $\mathbf{x}_{\text{local}}$  and  $\mathbf{V}_{\text{local}}$  are then fed to the 2D semantic model (encoder and decoder) and 3D point-based encoder  $\Phi$ , respectively. The outputs of the 2D semantic model are semantic-aware labels  $\tilde{\mathbf{y}}_{\text{sem}}^s$  and  $\tilde{\mathbf{y}}_{\text{sem}}^t$ . The outputs of the 3D point-based encoder  $\Phi$  are geometry-aware features  $\mathbf{F}_{\text{geo}}^s = \Phi(\mathbf{H}^s, \mathbf{V}^s)$  and  $\mathbf{F}_{\text{geo}}^t = \Phi(\mathbf{H}^t, \mathbf{V}^t)$ . We further propose a Geometric Adaptation (GA) module in Figure 4 to transfer geometric knowledge from  $\mathbf{F}_{\text{geo}}^s$  to  $\mathbf{F}_{\text{geo}}^t$  and obtain the enhanced target feature  $\tilde{\mathbf{F}}_{\text{geo}}^t$ . The geometry-aware




Figure 4: Geometry Adaption (GA) module. It computes the geometric shape similarity of the objects, and extracts the geometric features of the objects with similar shapes.

features  $\mathbf{F}_{\text{geo}}^s$  and  $\tilde{\mathbf{F}}_{\text{geo}}^t$  are processed by two FC Layers and then transformed to 2D pixel space as  $\tilde{\mathbf{y}}_{\text{geo}}^s$  and  $\tilde{\mathbf{y}}_{\text{geo}}^t$ .

**Geometric Adaption Module.** After the structure-texture and coordinate-color disentanglements, the 3D domain-invariant geometric information is isolated and can be utilized to perform adaptation to alleviate the domain gaps. Based on this motivation, we propose a Geometric Adaption (GA) module  $\Gamma$  in Figure 4 to transfer geometric knowledge from  $\mathbf{F}_{\text{geo}}^s$  to  $\mathbf{F}_{\text{geo}}^t$  and derive the enhanced target feature  $\tilde{\mathbf{F}}_{\text{geo}}^t = \Gamma(\mathbf{F}_{\text{geo}}^s, \mathbf{F}_{\text{geo}}^t)$ . Since the source and target patch may be different objects, this module is designed as a residual form with the attention operation. The GA module first performs concatenation on the geometric features  $\mathbf{F}_{\text{geo}}^s$  and  $\mathbf{F}_{\text{geo}}^t$ . Then the concatenated features  $[\mathbf{F}_{\text{geo}}^s, \mathbf{F}_{\text{geo}}^t]$  are further performed the attention operation. The derived feature is finally added to the target geometric features  $\mathbf{F}_{\text{geo}}^t$ . While being trained on the abundant image patches, the GA module computes the geometric shape similarity of objects, where the geometric features of the objects with similar shapes are effectively extracted by the attention operation.

**Loss Functions.** The semantic and geometric softmax probabilities of classes,  $\mathbf{p}_{\text{sem}}^s$ ,  $\mathbf{p}_{\text{sem}}^t$ ,  $\mathbf{p}_{\text{geo}}^s$ , and  $\mathbf{p}_{\text{geo}}^t$ , are computed training losses with the source ground truth label  $\mathbf{y}_{\text{local}}^s$  and target pseudo labels  $\hat{\mathbf{y}}_{\text{local}}^t$ , respectively. To increase the contribution of the long-tail classes, we derive the weights for each category  $c$  based on its corresponding pixel numbers. We first compute the frequency  $\omega_c^s$  of the pixel that belongs to the class  $c$  on the source domain, then we obtain the weights for each class and assign larger values to the long-tail classes  $\mathbf{w}_c^s = (\frac{\sum_{k=1}^{n_c} \omega_k^s}{\omega_c^s})^{\frac{1}{3}}$ ,  $\omega_c^s = \frac{1}{n_s HW} \sum_{i=1}^{n_s} \sum_{j=1}^{HW} \mathbf{m}_{cij}^s$ , where  $\mathbf{m}_{cij}^s$  is the mask at the pixel  $j$  that belongs to class  $c$ ,  $\mathbf{m}_{cij}^s = 1$ . We add the class weights  $\mathbf{w}^s$  to Eq. 1, and further derive the semantic and geometric segmentation loss:

$$\begin{aligned} \mathcal{L}_{\text{seg}}^s &= \mathcal{L}_{\text{sem}}^s + \mathcal{L}_{\text{geo}}^s \\ &= - \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \mathbf{w}_k^s \mathbf{y}_{\text{local}}^s \log \mathbf{p}_{\text{sem}}^s - \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \mathbf{w}_k^s \mathbf{y}_{\text{local}}^s \log \mathbf{p}_{\text{geo}}^s, \\ \mathcal{L}_{\text{seg}}^t &= \mathcal{L}_{\text{sem}}^t + \mathcal{L}_{\text{geo}}^t \\ &= - \sum_{j=1}^{HW} \sum_{k=1}^{n_c} \hat{\mathbf{y}}_{\text{local}}^t \log \mathbf{p}_{\text{sem}}^t - \sum_{j=1}^{HW} \sum_{k=1}^{n_c} \hat{\mathbf{y}}_{\text{local}}^t \log \mathbf{p}_{\text{geo}}^t. \end{aligned} \quad (8)$$

Furthermore, to minimize the semantic and geometric perceptual differences on the source and target domains, we employ the Kullback–Leibler (KL) divergence to design the Semantic-Geometric Consistency (SGC) loss  $\mathcal{L}_{\text{SGC}}$ :

$$\begin{aligned} \mathcal{L}_{\text{SGC}}^s &= \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \text{KL}(\mathbf{p}_{\text{geo}}^s \| \mathbf{p}_{\text{sem}}^s) + \text{KL}(\mathbf{p}_{\text{sem}}^s \| \mathbf{p}_{\text{geo}}^s), \\ \mathcal{L}_{\text{SGC}}^t &= \sum_{i=1}^{HW} \sum_{k=1}^{n_c} \text{KL}(\mathbf{p}_{\text{geo}}^t \| \mathbf{p}_{\text{sem}}^t) + \text{KL}(\mathbf{p}_{\text{sem}}^t \| \mathbf{p}_{\text{geo}}^t). \end{aligned} \quad (9)$$

Thus, the total loss of the geometry-aware adaptation stage is:

$$\mathcal{L}_{\text{geometry-aware}} = \mathcal{L}_{\text{seg}}^s + \mathcal{L}_{\text{seg}}^t + \alpha \mathcal{L}_{\text{SGC}}^s + \beta \mathcal{L}_{\text{SGC}}^t, \quad (10)$$

where  $\alpha$  and  $\beta$  control the weights of the SGC loss on each domain.

**Inference.** The Geometry-Aware Adaptation stage incorporates the complementary geometric information from 3D point cloud to strengthen final segmentation predictions. In inference, we maintain the 2D semantic-aware predictions  $\mathbf{p}_{\text{sem}}^t$  as the final prediction and discard the 3D adaptive part. Thus, the 3D geometry-aware predictions  $\mathbf{p}_{\text{geo}}^t$  is only applied to optimize  $\mathbf{p}_{\text{sem}}^t$  in the Geometry-Aware Adaptation.

## Experiments

Our proposed GANDA is evaluated on the benchmarks SYNTHIA→Cityscapes and GTA5→Cityscapes.

### Implementation

**Datasets.** In our experiments, we adopt the GTA5 and SYNTHIA datasets as the source synthetic datasets, respectively, and exploit the Cityscapes dataset as the target real scene dataset. The GTA5 dataset (Richter et al. 2016) has 24,966 synthetic photo-realistic scenes with a resolution of  $1914 \times 1052$ . SYNTHIA (Ros et al. 2016) is a large-scale synthetic urban image dataset, containing 9,400 synthetic urban images with the resolution of  $1280 \times 760$  in model training. The Cityscapes dataset (Cordts et al. 2016) is a realistic dataset of the  $2048 \times 1024$  street scenes from 50 cities, where 2,975 images for training and 500 ones for validation. **Setups.** Two synthetic datasets both have pixel-level segmentation annotations, and these labels are utilized to provide supervision on the source data. Besides, the GTA5 dataset shares 19 semantic categories, and the SYNTHIA dataset shares 16 common classes with the Cityscapes dataset, respectively. We compute the Intersection-over-Union (IoU) on each class and derive the mean IoU (mIoU) for the evaluation metric.

**Implementation Details.** Our model is implemented in PyTorch. Following the previous works (Zhang et al. 2021; Hoyer, Dai, and Van Gool 2022b), we utilize the DeepLabv2 (Chen et al. 2018) and SegFormer (Xie et al. 2021) as the segmentation architectures. On GTA5→Cityscapes, since the GTA5 dataset has no depth maps, we obtain the depths via training the NeWCRFs (Yuan et al. 2022) on the GTA5-based dataset DeepMVS (Huang et al. 2018). On SYNTHIA→Cityscapes, we utilize the provided depth maps from Synthia-Rand-Cityscapes. We conduct the experiments on 4 Tesla V100 GPUs.

Method	road	sideway	building	wall	fence	pole	light	sign	vege.	terrace	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
Source Only	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
Seg-Uncertainty	90.4	31.2	85.1	36.9	25.6	37.5	48.8	48.5	85.3	34.8	81.1	64.4	36.8	86.3	34.9	52.2	1.7	29.0	44.6	50.3
SAC	90.4	53.9	86.6	42.4	27.3	45.1	48.5	42.7	87.4	40.1	86.1	67.5	29.7	88.5	49.1	54.6	9.8	26.6	45.3	53.8
Coarse-to-Fine	92.5	58.3	86.5	27.4	28.8	38.1	46.7	42.5	85.4	38.4	91.8	66.4	37.0	87.8	40.7	52.4	44.6	41.7	59.0	56.1
CorDA	94.7	63.1	87.6	30.7	40.6	40.2	47.8	51.6	87.6	47.0	89.7	66.7	35.9	90.2	48.9	57.5	0.0	39.8	56.0	56.6
ProDA	87.8	56.0	79.7	46.3	44.8	45.6	53.5	53.5	88.6	45.2	82.1	70.7	39.2	88.8	45.5	59.4	1.0	48.9	56.4	57.5
DAFormer	95.7	70.2	89.4	53.5	48.1	49.6	55.8	59.4	89.9	47.9	92.5	72.2	44.7	92.3	74.5	78.2	65.1	55.9	61.8	68.3
HRDA	96.4	74.4	91.0	61.6	51.5	57.1	63.9	<b>69.3</b>	<u>91.3</u>	<u>48.4</u>	<u>94.2</u>	<u>79.0</u>	<u>52.9</u>	<u>93.9</u>	<b>84.1</b>	<u>85.7</u>	<b>75.9</b>	<u>63.9</u>	<u>67.5</u>	<u>73.8</u>
GANDA	<b>96.5</b>	<b>74.8</b>	<b>91.4</b>	<b>61.7</b>	<b>57.3</b>	<b>59.2</b>	<b>65.4</b>	<u>68.8</u>	<b>91.5</b>	<b>49.9</b>	<b>94.7</b>	<b>79.6</b>	<b>54.8</b>	<b>94.1</b>	<u>81.3</u>	<b>86.8</b>	<u>74.6</u>	<b>64.8</b>	<b>68.2</b>	<b>74.5</b>

Table 1: Quantitative results on GTA5→Cityscapes in the metrics of mIoU. mIoU is the score on 19 categories. The best score is highlighted bold and the second best one is marked underline for each column.

Method	road	sideway	building	wall*	fence*	pole*	light	sign	vege.	sky	person	rider	car	bus	motor	bike	mIoU	mIoU*
Source Only	64.3	21.3	73.1	2.4	1.1	31.4	7.0	27.7	63.1	67.6	42.2	19.9	73.1	15.3	10.5	38.9	34.9	40.3
GIO-Ada	78.3	29.2	76.9	11.4	0.3	26.5	10.8	17.2	81.7	81.9	45.8	15.4	68.0	15.9	7.5	30.4	37.3	43.0
DADA	89.2	44.8	81.4	6.8	0.3	26.2	8.6	11.1	81.8	84.0	54.7	19.3	79.7	40.7	14.0	38.8	42.6	49.8
Seg-Uncertainty	87.6	41.9	83.1	14.7	1.7	36.2	31.3	19.9	81.6	80.6	63.0	21.8	86.2	40.7	23.6	53.1	47.9	54.9
Coarse-to-Fine	75.7	30.0	81.9	11.5	2.5	35.3	18.0	32.7	86.2	90.1	65.1	33.2	83.3	36.5	35.3	54.3	48.2	55.5
SAC	<u>89.3</u>	47.2	85.5	26.5	1.3	43.0	45.5	32.0	<u>87.1</u>	89.3	63.6	25.4	86.9	35.6	30.4	53.0	51.6	59.3
CorDA	<b>93.3</b>	<b>61.6</b>	85.3	19.6	5.1	37.8	36.6	42.8	<u>84.9</u>	90.4	69.7	41.8	85.6	38.4	32.6	53.9	55.0	62.8
ProDA	87.8	45.7	84.6	37.1	0.6	44.0	54.6	37.0	<b>88.1</b>	84.4	74.2	24.3	88.2	51.1	40.5	45.6	55.5	62.0
DAFormer	84.5	40.7	88.4	41.5	<u>6.5</u>	50.0	55.0	54.6	86.0	89.8	73.2	48.2	87.2	53.2	53.9	61.7	60.9	67.4
HRDA	85.2	47.7	88.8	49.5	4.8	<u>57.2</u>	<u>65.7</u>	<b>60.9</b>	85.3	<u>92.9</u>	<u>79.4</u>	<u>52.8</u>	<u>89.0</u>	<u>64.7</u>	<b>63.9</b>	<u>64.9</u>	65.8	72.4
GANDA	89.1	<u>50.6</u>	<b>89.7</b>	<b>51.4</b>	<b>6.7</b>	<b>59.4</b>	<b>66.8</b>	<u>57.7</u>	86.7	<b>93.8</b>	<b>80.6</b>	<b>56.9</b>	<b>90.7</b>	<b>64.8</b>	<u>62.6</u>	<b>65.0</b>	<b>67.0</b>	<b>73.5</b>

Table 2: Quantitative results on SYNTHIA→Cityscapes in the metrics of mIoU. mIoU and mIoU\* are the scores on 16 categories and 13 categories. The best score is highlighted bold and the second best one is marked underline for each column.




Figure 5: Comparisons with and without the 2D and 3D segmentation on GANDA+ProDA on the SYNTHIA→Cityscapes task in the Geometry-Aware Adaptation stage in the metrics of mIoU and s-mIoU.

### Comparisons with State-Of-The-Arts

The evaluations of our GANDA are conducted via the comparisons of several state-of-the-art methods on the two scenarios: GTA5→Cityscapes in Table 1 and SYNTHIA→Cityscapes in Table 2. These methods are mainly divided into two categories. a) **Self-training Methods**: Seg-Uncertainty (Zheng and Yang 2021), SAC (Araslanov and Roth 2021), Coarse-to-Fine (Ma et al. 2021), ProDA (Zhang et al. 2021), DAFormer (Hoyer, Dai,

and Van Gool 2022a), and HRDA (Hoyer, Dai, and Van Gool 2022b). b) **Depth-based Multi-task Learning Methods**: GIO-Ada (Chen et al. 2019a), DADA (Vu et al. 2019b), and CorDA (Wang et al. 2021). On the GTA5→Cityscapes task in Table 1, our proposed GANDA outperforms of the state-of-the-art HRDA by 0.7% mIoU. On the more challenging SYNTHIA→Cityscapes task in Table 2, our proposed GANDA achieves the highest mIoU score as 67.0% and 73.5%. The performances well illustrate the capability of GANDA in shrinking the domain gaps and leveraging the geometric information to perceive the shapes and spatial positions of objects.

### Effectiveness of Geometry-Aware Adaptation

Since GANDA is a plug-and-play framework that can be applied to any existed UDA method, we perform the experiments on the multiple UDA methods with and without the Geometry-Aware Adaptation in Table 3, including self-training ProDA (Zhang et al. 2021), depth-based CorDA (Wang et al. 2021), and self-training HRDA (Hoyer, Dai, and Van Gool 2022b). The gains are evident in each method in terms of all the metrics, which demonstrates the efficacy of differentiating the small object classes.

Method	road	sideway	building	wall*	fence*	pole*	light	sign	vege.	sky	person	rider	car	bus	motor	bike	mIoU	mIoU*
<i>Self-training Method</i>				<i>Architecture: DeepLabv2</i>						<i>Backbone: ResNet-101</i>								
ProDA	87.0	44.0	83.2	26.9	1.0	42.0	45.8	34.2	86.7	81.3	68.4	22.1	87.7	50.0	31.4	38.6	51.9	58.5
GANDA+ProDA	87.6	44.2	84.0	31.0	1.0	42.0	50.5	38.5	87.1	81.2	68.5	27.2	88.7	50.2	44.4	46.6	54.5	61.4
Gain	<b>↑0.6</b>	<b>↑0.2</b>	<b>↑0.8</b>	<b>↑4.1</b>	<b>↑0.0</b>	<b>↑0.0</b>	<b>↑4.7</b>	<b>↑4.3</b>	<b>↑0.4</b>	↓0.1	<b>↑0.1</b>	<b>↑5.1</b>	<b>↑1.0</b>	<b>↑0.2</b>	<b>↑13.0</b>	<b>↑8.0</b>	<b>↑2.6</b>	<b>↑2.9</b>
<i>Self-training Method with Distillation</i>				<i>Architecture: DeepLabv2</i>						<i>Backbone: ResNet-101</i>								
ProDA*	87.8	45.7	84.6	37.1	0.6	44.0	54.6	37.0	88.1	84.4	74.2	24.3	88.2	51.1	40.5	45.6	55.5	62.0
GANDA+ProDA*	87.6	44.8	85.6	36.5	0.0	45.8	50.3	40.0	88.2	86.4	71.1	30.1	90.8	57.4	40.6	52.9	56.8	63.5
Gain	↓0.2	↓0.9	<b>↑1.0</b>	↓0.6	↓0.6	<b>↑1.8</b>	↓4.3	↓3.0	<b>↑0.1</b>	<b>↑2.0</b>	↓3.1	<b>↑5.8</b>	<b>↑2.6</b>	<b>↑6.3</b>	<b>↑0.1</b>	<b>↑7.3</b>	<b>↑1.3</b>	<b>↑1.5</b>
<i>Depth-based Multi-task Learning Method</i>				<i>Architecture: DeepLabv2</i>						<i>Backbone: ResNet-101</i>								
CorDA	93.3	61.6	85.3	19.6	5.1	37.8	36.6	42.8	84.9	90.4	69.7	41.8	85.6	38.4	32.6	53.9	55.0	62.8
GANDA+CorDA	87.1	45.8	86.1	28.9	4.8	37.1	40.6	45.0	87.0	87.9	69.1	39.8	89.9	59.8	33.8	57.2	56.3	63.8
Gain	↓6.2	↓15.8	<b>↑0.8</b>	<b>↑9.3</b>	↓0.3	↓0.7	<b>↑4.0</b>	<b>↑2.2</b>	<b>↑2.1</b>	↓2.5	↓0.6	↓2.0	<b>↑4.3</b>	<b>↑21.4</b>	<b>↑1.2</b>	<b>↑3.3</b>	<b>↑1.3</b>	<b>↑1.0</b>
<i>Self-training Method</i>				<i>Architecture: SegFormer</i>						<i>Backbone: MiT-B5</i>								
HRDA	85.2	47.7	88.8	49.5	4.8	57.2	65.7	60.9	85.3	92.9	79.4	52.8	89.0	64.7	63.9	64.9	65.8	72.4
GANDA+HRDA	89.1	50.6	89.7	51.4	6.7	59.4	66.8	57.7	86.7	93.8	80.6	56.9	90.7	64.8	62.6	65.0	67.0	73.5
Gain	<b>↑3.9</b>	<b>↑2.9</b>	<b>↑0.9</b>	<b>↑1.9</b>	<b>↑1.9</b>	<b>↑2.2</b>	<b>↑1.1</b>	↓3.2	<b>↑1.4</b>	<b>↑1.1</b>	<b>↑1.2</b>	<b>↑4.1</b>	<b>↑1.7</b>	<b>↑0.1</b>	↓1.3	<b>↑0.1</b>	<b>↑1.2</b>	<b>↑1.1</b>

Table 3: Effectiveness of the Geometry-Aware Adaptation in multiple UDA methods on SYNTHIA  $\rightarrow$  Cityscapes in the metrics of mIoU and mIoU\*. mIoU and mIoU\* are the scores on 16 categories and 13 categories, respectively. The symbol “\*” indicates the distillation operation is applied.

Stage & Components					mIoU	Gain
Source Only					34.9	+0.0
AA	Image Only				40.9	+6.0
	Depth-Aware				41.2	+6.3
GAA	ST	GA	WCE	SGC	mIoU	Gain
	✓			✓	51.9	+17.0
	✓			✓	52.3	+17.4
	✓	✓			52.9	+18.0
	✓		✓		53.1	+18.2
✓	✓	✓	✓	54.5	+19.6	

Table 4: Ablation studies of each component on the GANDA+ProDA on the SYNTHIA  $\rightarrow$  Cityscapes task.

## Ablation Study

The proposed components of our GANDA are evaluated based on their contributions to the model performances in Table 4, including Depth-Aware Adversarial Adaptation, Geometry Adaptation module (GA), Weighted CE loss (WCE), and Semantic-Geometric Consistency loss (SGC). At the stage of Adversarial Adaptation (AA), the introduction of depth estimation can gain the performances of 41.2% mIoU. At the stage of Geometry-Aware Adaptation (GAA), the model using only the Self-Training (ST) techniques has a gain of 17.0% mIoU compared to the source only model. After adding these parts, the model performances achieve the gains from 17.4% to 19.6%. Besides, we perform the comparisons with and without the 2D and 3D segmentation on GANDA+ProDA on the SYNTHIA  $\rightarrow$  Cityscapes task in Figure 5, which proves that 2D semantic segmentation is necessary and the 3D geometric part is complementary.

Model	Training		Inference	
	GFLOPs	Params (M)	GFLOPs	Params (M)
HRDA	146.9	85.7		
GANDA+HRDA	150.4	86.7	146.9	85.7

Table 5: Computation comparisons in HRDA with and without Geometry-Aware Adaptation in the training and inference phases with the metrics of GFLOPs and Parameters (Params) (M) on the  $512 \times 512$  input image and the extra 40000 input points in training.

## Computation Costs in Training and Inference

Table 5 presents the computation costs in training and inference in HRDA with and without Geometry-Aware Adaptation in terms of GFLOPs and parameters. Inclusion of the GANDA does not lead to the large computation costs, since it is only applied in training and discarded in inference.

## Conclusions

In this paper, we address the task of synthetic-to-real domain adaptive semantic segmentation by proposing a two-stage geometry-aware training framework. Beyond the previous depth-based UDA methods and based only on source depth information, GANDA gains a deep insight into the 3D geometric space and extracts the domain-invariant information after the structure-texture and coordinate-color disentanglements. Besides, without any structure or model dependency, GANDA is plug-and-play and boosts the performances of other existing UDA methods, especially effective in classifying small objects. Comprehensive experimental analysis demonstrate the effectiveness of GANDA that outperforms other methods on two prevalent synthetic-to-real tasks.

## Acknowledgments

This work was supported in part by JCYJ20220530143600001, by the Basic Research Project No. HZQB-KCZYZ-2021067 of Hetao Shenzhen HK S&T Cooperation Zone, by the National Key R&D Program of China with grant No.2018YFB1800800, by SGDX20211123112401002, by Shenzhen Outstanding Talents Training Fund, by Guangdong Research Project No. 2017ZT07X152 and No. 2019CX01X104, by the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001), by the NSFC 61931024&8192 2046, by NSFC-Youth 62106154, by zelixir biotechnology company Fund, by Tencent Open Fund, and by ITSO at CUHKSZ.

## References

- Araslanov, N.; and Roth, S. 2021. Self-supervised Augmentation Consistency for Adapting Semantic Segmentation. In *CVPR*, 15384–15394.
- Chang, W.-L.; Wang, H.-P.; Peng, W.-H.; and Chiu, W.-C. 2019. All About Structure: Adapting Structural Information Across Domains for Boosting Semantic Segmentation. In *CVPR*, 1900–1909.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2018. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE TPAMI*, 40(4): 834–848.
- Chen, Y.; Li, W.; Chen, X.; and Gool, L. V. 2019a. Learning Semantic Segmentation From Synthetic Data: A Geometrically Guided Input-Output Adaptation Approach. In *CVPR*, 1841–1850.
- Chen, Y.-C.; Lin, Y.-Y.; Yang, M.-H.; and Huang, J.-B. 2019b. CrDoCo: Pixel-Level Domain Transfer With Cross-Domain Consistency. In *CVPR*, 1791–1800.
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In *CVPR*, 3075–3084.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *CVPR*, 8515–8525.
- Hoyer, L.; Dai, D.; and Van Gool, L. 2022a. DAFormer: Improving Network Architectures and Training Strategies for Domain-Adaptive Semantic Segmentation. In *CVPR*, 9924–9935.
- Hoyer, L.; Dai, D.; and Van Gool, L. 2022b. HRDA: Context-Aware High-Resolution Domain-Adaptive Semantic Segmentation. In *ECCV*.
- Huang, P.-H.; Matzen, K.; Kopf, J.; Ahuja, N.; and Huang, J.-B. 2018. DeepMVS: Learning Multi-View Stereopsis. In *CVPR*, 2821–2830.
- Lee, K.-H.; Ros, G.; Li, J.; and Gaidon, A. 2019. SPIGAN: Privileged Adversarial Learning from Simulation. In *ICLR*.
- Ma, H.; Lin, X.; Wu, Z.; and Yu, Y. 2021. Coarse-to-Fine Domain Adaptive Semantic Segmentation with Photometric Alignment and Category-Center Regularization. In *CVPR*, 4051–4060.
- Nuriel, O.; Benaim, S.; and Wolf, L. 2021. Permuted AdaIN: Reducing the Bias Towards Global Statistics in Image Classification. In *CVPR*, 9482–9491.
- Pan, F.; Shin, I.; Rameau, F.; Lee, S.; and Kweon, I. S. 2020. Unsupervised Intra-domain Adaptation for Semantic Segmentation through Self-supervision. In *CVPR*, 3764–3773.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *CVPR*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *NeurIPS*.
- Richter, S. R.; Vineet, V.; Roth, S.; and Koltun, V. 2016. Playing for Data: Ground Truth from Computer Games. In *ECCV*, 102–118.
- Ros, G.; Sellart, L.; Materzynska, J.; Vazquez, D.; and Lopez, A. M. 2016. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *CVPR*, 3234–3243.
- Tang, H.; Liu, Z.; Zhao, S.; Lin, Y.; Lin, J.; Wang, H.; and Han, S. 2020. Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution. In *ECCV*, 685–702.
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. KPConv: Flexible and Deformable Convolution for Point Clouds. In *ICCV*.
- Vu, T.-H.; Jain, H.; Bucher, M.; Cord, M.; and Pérez, P. 2019a. ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. In *CVPR*, 2517–2526.
- Vu, T. H.; Jain, H.; Bucher, M.; Cord, M.; and Perez, P. P. 2019b. DADA: Depth-Aware Domain Adaptation in Semantic Segmentation. In *ICCV*, 7363–7372.
- Wang, Q.; Dai, D.; Hoyer, L.; Fink, O.; and Van Gool, L. 2021. Domain Adaptive Semantic Segmentation with Self-Supervised Depth Estimation. In *ICCV*, 8515–8525.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3D ShapeNets: A deep representation for volumetric shapes. In *CVPR*, 1912–1920.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In *NeurIPS*.
- Yan, X.; Gao, J.; Li, J.; Zhang, R.; Li, Z.; Huang, R.; and Cui, S. 2021. Sparse Single Sweep LiDAR Point Cloud Segmentation via Learning Contextual Shape Priors from Scene Completion. In *AAAI*, volume 35, 3101–3109.
- Yan, X.; Gao, J.; Zheng, C.; Zheng, C.; Zhang, R.; Cui, S.; and Li, Z. 2022. 2DPASS: 2D Priors Assisted Semantic Segmentation on LiDAR Point Clouds. In *ECCV*, 677–695.
- Yan, X.; Zheng, C.; Li, Z.; Wang, S.; and Cui, S. 2020. PointASNL: Robust Point Clouds Processing using Nonlocal Neural Networks with Adaptive Sampling. In *CVPR*, 5589–5598.
- Yang, J.; An, W.; Wang, S.; Zhu, X.; Yan, C.; and Huang, J. 2020a. Label-Driven Reconstruction for Domain Adaptation in Semantic Segmentation. In *ECCV*, 480–498.

Yang, J.; Xu, R.; Li, R.; Qi, X.; Shen, X.; Li, G.; and Lin, L. 2020b. An Adversarial Perturbation Oriented Domain Adaptation Approach for Semantic Segmentation. In *AAAI*, 12613–12620.

Yuan, W.; Gu, X.; Dai, Z.; Zhu, S.; and Tan, P. 2022. Neural Window Fully-Connected CRFs for Monocular Depth Estimation. In *CVPR*, 3916–3925.

Zhang, P.; Zhang, B.; Zhang, T.; Chen, D.; Wang, Y.; and Wen, F. 2021. Prototypical Pseudo Label Denoising and Target Structure Learning for Domain Adaptive Semantic Segmentation. In *CVPR*, 12414–12424.

Zheng, Z.; and Yang, Y. 2021. Rectifying Pseudo Label Learning via Uncertainty Estimation for Domain Adaptive Semantic Segmentation. *IJCV*, 129(4): 1106–1120.

Zhu, X.; Zhou, H.; Wang, T.; Hong, F.; Ma, Y.; Li, W.; Li, H.; and Lin, D. 2021. Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR Segmentation. In *CVPR*, 9939–9948.