

Learning Part Segmentation through Unsupervised Domain Adaptation from Synthetic Vehicles

Qing Liu^{1*}, Adam Kortylewski¹, Zhishuai Zhang¹, Zizhang Li^{2†}
Mengqi Guo^{3†}, Qihao Liu¹, Xiaoding Yuan^{4†}, Jiteng Mu^{5†}, Weichao Qiu¹, Alan Yuille¹
¹Johns Hopkins University ²Zhejiang University
³Beihang University ⁴Tongji University ⁵University of California, San Diego

Abstract

Part segmentations provide a rich and detailed part-level description of objects. However, their annotation requires an enormous amount of work, which makes it difficult to apply standard deep learning methods. In this paper, we propose the idea of learning part segmentation through unsupervised domain adaptation (UDA) from synthetic data. We first introduce UDA-Part, a comprehensive part segmentation dataset for vehicles that can serve as an adequate benchmark for UDA¹. In UDA-Part, we label parts on 3D CAD models which enables us to generate a large set of annotated synthetic images. We also annotate parts on a number of real images to provide a real test set. Secondly, to advance the adaptation of part models trained from the synthetic data to the real images, we introduce a new UDA algorithm that leverages the object’s spatial structure to guide the adaptation process. Our experimental results on two real test datasets confirm the superiority of our approach over existing works, and demonstrate the promise of learning part segmentation for general objects from synthetic data. We believe our dataset provides a rich testbed to study UDA for part segmentation and will help to significantly push forward research in this area.

1. Introduction

Part-based object representations are of key importance for many computer vision tasks such as object recognition [1, 9, 72, 98], pose estimation [13, 32, 93, 96], action detection [85], and scene understanding [65, 70, 75]. Currently, part-based approaches often represent objects as a set of sparse keypoints, because these are easy to annotate in large-scale datasets for training deep neural networks. By contrast, part segmentations provide a richer and more detailed part-level object description. Instead of recognizing specific parts sparsely on the object (e.g., keypoint or part




Figure 1. An illustration of learning part segmentation through unsupervised domain adaptation (UDA) from synthetic vehicles. Based on part annotation on 3D CAD models, we propose to use UDA to learn from large-scale labeled synthetic samples and unlabeled real images, and the goal is to make accurate part segmentation predictions on real test images.

detection), part segmentation gives a complete description of an object by assigning every pixel belonging to the object one and only one part label. This is a lot more challenging task and requires a much greater annotation effort.

Given their recent success, deep learning methods have dominated the studies of computer vision, including object segmentation [7, 10, 47]. However, these deep models usually require a large amount of annotated training data to achieve satisfying performance. Existing part segmentation datasets mostly contain only a small number of images [9, 74], or define only a small number of parts per object category [74, 97], or focus on a single object category, such as humans [18, 19, 43, 94] and faces [32, 38, 39]. These limitations inhibit effective training of standard deep segmentation networks and have largely impeded the development of computer vision models that leverage part information. By contrast, 3D CAD models are available for many different objects, and, once annotated, can be used to generate large-scale part segmentation datasets automatically.

In this work, we propose to solve part segmentation on general objects by learning from synthetic data (Figure 1) and address the problem in two steps. In the first step, we introduce UDA-Part, a comprehensive part segmentation dataset that can serve as an adequate benchmark for

*Corresponding author: qingliu.research@gmail.com

†This work is done during internship at Johns Hopkins University.

¹<https://qliu24.github.io/udapart/>

UDA. UDA-Part is composed of 21 3D CAD models from 5 vehicle categories. For each category, we define a fine-grained set of parts which are consistently annotated across all CAD models of the corresponding category. Based on these CAD models and their part annotations, we are able to render a large-scale synthetic image dataset with automatically generated part segmentation ground-truth. These synthetic data are sufficient to train deep neural networks and may also be used for model evaluation or diagnosis. To evaluate how the models trained from synthetic data perform on the real images, we also label parts on 200 real images collected from PASCAL3D+ [93] and include them as the target test set in UDA-Part.

Secondly, we introduce a new unsupervised domain adaptation (UDA) algorithm for part segmentation. UDA has been explored for image classification [30,59], keypoint detection [53,103], and semantic segmentation [31,105], where it achieves satisfactory results on real images with little annotation cost. To further advance UDA performance on part segmentation, we introduce Geometric-Matching Guided domain adaptation (GMG). GMG conducts cross-domain geometric matching based on a global transformation function between real and synthetic images. The function puts a smoothness constraint on the matching and thus adaptively preserves the spatial relations between the parts. Once an optimal match is found, GMG transfers the synthetic labels to the real images and retains the high-confidence results as pseudo-labels for a joint training process. In short, GMG makes part-relation-aware adaptation by explicitly using the object structure depicted in the synthetic samples. In our experiments, GMG outperforms other UDA baselines for part segmentation on both the UDA-Part real test images and the PascalPart [9] test set.

In summary, our main contributions are:

1. We propose to learn part segmentation for general objects through unsupervised domain adaptation (UDA) from synthetic data.
2. We introduce a new part segmentation dataset for vehicles called UDA-Part which can serve as a comprehensive benchmark for part segmentation through UDA.
3. We introduce a new UDA algorithm for part segmentation called Geometric-Matching Guided domain adaptation (GMG), which leverages the object’s spatial structure to guide the adaptation and achieves superior results.

2. Related Work

Part segmentation. Both rigid [49,69] and non-rigid objects [18,43,82,94] have been studied in part segmentation, where structure-based methods, e.g., compositional models, are widely used [49,82,84,91,92]. Architecture-wise, both fully convolutional network (FCN) [84,90] and long short-term memory (LSTM) [42,44] have been studied. Many works explore the use of auxiliary tasks (e.g.,

pose estimation) to boost part segmentation and get promising results [17,19,55,91]. Moreover, 3D information such as 3D geometric features [69] and depth [66] could also be embedded into the models to improve the performance. Beyond single-object parsing, multi-object parsing is recently proposed and studied [51,101]. While all these methods require pixel-level annotations on real images, some works start to use synthetic data to advance human parsing [29,81], but they are limited to human object. In this work, we aim to solve part segmentation on general objects by learning from synthetic vehicles.

Datasets with part annotations for rigid objects. ImageNetPart [83] provides bounding box annotation for parts on 6 vehicle categories. PASCAL3D+ [93] includes keypoint annotation on 12 rigid object categories. CarFusion [61] and ApolloCar3D [68] contain keypoint annotation for cars in street scenes. For 3D object part recognition, PartNet [52] provides hierarchical part annotations on 3D models covering 24 object categories, most of which are indoor furniture and none is a vehicle. Yi et al. [97] label parts on 3D models selected from 16 categories in ShapeNetCore [4], while their definition of part is coarse and the average number of parts per category is less than 4. In the context of image part segmentation, PascalPart [9] has been widely studied. It includes 20 categories but provides limited number of samples and parts. MVP [46] is recently introduced to provide detailed part segmentation labels for cars in the wild. An earlier dataset ETHZ [74] provides annotations of 5 parts on 141 wheelchair images and 6 parts on 139 car images, which is not adequate for deep network training. PartImageNet [23] is a very recent large scale dataset consisting of 158 classes with up to 5 parts annotated per class, requiring several hundred hours of annotation. Our aim is to build a large-scale part segmentation dataset for generic objects efficiently using 3D computer graphics models.

Learning from synthetic data. Synthetic data generated by computer graphics techniques are effective for model diagnosis [28,100] and have boosted performance in many real-world application domains [15,22,37,53,76,81]. However, the domain shift between synthetic data and real-world data limits the improvement. To overcome this, domain adaptation is proposed [58]. Maximum Mean Discrepancy (MMD) and its kernel variants [48,71,78,79] have been studied to reduce the difference between source and target domain distributions. In the context of unsupervised domain adaptation (UDA) for semantic segmentation, self-training is one of the mainstream research directions [40,99,104,105], while there has also been an increasing interest in using style transfer [5,24,41,54,64,89] or feature alignment [25,77] to encourage domain-wise marginal distribution matching. Other methods adopt category-aware feature alignment or local contextual feature simi-




Figure 2. An overview of our UDA-Part dataset. UDA-Part is a comprehensive part segmentation dataset for five vehicle categories. Each category is composed of several 3D CAD models/prototypes. By labeling parts on these 3D CAD models, we render a large-scale synthetic image dataset with automatically generated part segmentation ground-truth (left panel). To benchmark unsupervised domain adaptation methods, we also label parts on a moderate number of real images for testing (right panel).

larity [14, 26, 50, 99]. We promote the idea of using UDA to solve part segmentation and leverage object structure to guide the knowledge transfer from synthetic to real.

Geometric matching. Geometric matching aims at finding spatial correspondences among images belonging to the same category at a fine-grained level. Both hand-engineered descriptors [3, 33, 45] and pre-trained convolutional neural network (CNN) features [20, 73, 80, 95] are explored in early works. Recent progress is made in trainable image descriptors [34, 35, 56] and trainable geometric models [6, 21, 62, 63]. However, they have only explored geometric matching for images in the same domain. Bai et al. [2] demonstrate that pre-trained CNN features can be effectively used to find sparse spatial correspondence between synthetic images and real images. Zhou et al. [102] use 3D cycle consistency to learn dense correspondence between real-to-real and real-to-synthetic pairs. However, their method does not consider the global structure nor constrain it with a geometric transformation function, thus the result is less competitive. In this work, we explore cross-domain geometric matching without supervision and integrate it into our part segmentation framework.

3. The UDA-Part Dataset

In this section, we introduce UDA-Part, a part segmentation dataset for vehicles that provides detailed annotations on both synthetic images for training and real images for testing (Figure 2). UDA-Part can serve as a comprehensive benchmark for part segmentation through UDA.

3.1. Data Generation

We build UDA-Part in five steps. Examples of the annotations on the synthetic images and the real test images are shown in Figure 2. More details (e.g., CAD models, part list, etc.) could be found in the supplementary material.

(1) Select 3D CAD models. UDA-Part is composed of 21 3D CAD models from 5 vehicle categories: *car*, *motorbike*, *aeroplane*, *bus*, and *bicycle*. Each CAD model represents a common prototype (i.e., subtype) of the category it belongs to. For example, for the category *bicycle*, the CAD models are different bicycle subtypes such as *utility*, *sports*, *road*, and *tandem*. These can effectively represent the structural variability of bicycle objects. In total, we select 21 CAD models to be included and annotated in UDA-Part.

(2) Define list of parts to be studied. We take references from existing vehicle part datasets [9, 74] and Wikipedia [86–88] to determine the part list for each category. Comparing with the part list in PascalPart [9], we make more fine-grained level definitions. For example, in PascalPart, the category *car* has a part “back side”. We instead distinguish “back windshield”, “tail light”, “back bumper”, and “trunk”. The fine-grained definition used here makes it possible to merge the part list and map to other coarse lists defined in existing datasets.

(3) Annotate parts on 3D CAD models. We adopt the Blender [11] plugin built by Kim et al. [36] to perform per-mesh part labeling on the 3D CAD models. The plugin allows the user to assign a label to a group of selected meshes and save the results to a JSON file. We include a quality control step to ensure each surface mesh is assigned one and




Figure 3. Histogram of number of parts per image in PascalPart [9] test split and UDA-Part real test images. The x-axis represents count of images.

only one part label. Annotating all CAD models in UDA-Part takes roughly 300 working hours.

(4) Render synthetic images with part annotations.

We use Blender [11] as our renderer to generate the synthetic images with part segmentation ground-truth. Following previous work [60, 76], we randomize the render parameters (e.g., viewpoint, lighting, object texture, etc.) to enable nuisance factor control and facilitate domain generalization. We generate 8000 synthetic images with resolution 2048×1024 for each 3D CAD model and split the training and test set with a ratio of 3 : 1, resulting in a total of 126000 images for training and 42000 for testing.

(5) **Annotate parts on real test images.** We manually label part segmentations on 200 real vehicle images (40 images per category) for testing. The images are selected from PASCAL3D+ dataset [93] to contain objects with different subtypes and evenly distributed viewpoints. We use the VGG Image Annotator (VIA) [16] to label the parts on the images, which takes about 150 working hours.

3.2. Dataset Comparison

In Table 1, we compare UDA-Part with existing datasets that provide vehicle part segmentation labels. UDA-Part contains the largest number of parts per category, and its annotation on both synthetic images and real images makes it adequate for UDA studies. In Figure 3, we compare the number of parts per image labeled on PascalPart and UDA-Part real test images. Generally, UDA-Part contains 2 to 3 times more annotations per image, making it a more challenging part segmentation benchmark. More comprehensive comparisons between UDA-Part and PascalPart can be found in the supplementary material.

4. Geometric-Matching Guided Adaptation

To facilitate studies targeting real-world data applications, we explore unsupervised domain adaptation (UDA) algorithms which enable the models trained on the synthetic data of UDA-Part to perform well on the real test images. In this section, we introduce our proposed Geometric-Matching Guided domain adaptation (GMG) approach. The key steps of GMG are illustrated in Figure 4 and will be discussed in detail in the following.

4.1. Preliminaries for Part Segmentation and UDA

We start with the preliminaries for part segmentation and unsupervised domain adaptation (UDA). In our work and

	UDA-Part (Ours)	PascalPart [9]	ETHZ [74]	SNPart [97]
3D Models	✓	✗	✗	✓
Syn. Img.	✓	✗	✗	✗
Real Img.	✓	✓	✓	✗
Avg. #Parts per Cat.	24	9	6	4

Table 1. Comparison of UDA-Part with existing datasets that have part segmentation labels on vehicles.

all baseline experiments, we assume the category label is known, such that the part segmentation models are trained and tested for each object category separately. This assumption is reasonable since many off-the-shelf classifiers are freely available for the objects included in UDA-Part.

Similar to semantic segmentation, a part segmentation model \mathbf{M} can be formulated as a mapping function from the image domain to the output label domain: $\mathbf{M} : I \rightarrow Y$, which predicts a pixel-wise category label $Y \in \{1, \dots, C\}^{H \times W}$, where H and W denote the image size, C is the total number of part categories.

In UDA, the data are usually collected from two domains: source domain \mathcal{S} (i.e., synthetic data) and target domain \mathcal{T} (i.e., real data). During training, we have access to the labeled training samples (I^s, Y^s) from $\mathcal{S}_{\text{train}}$ and unlabeled training samples I^t from $\mathcal{T}_{\text{train}}$. The goal is to train a model \mathbf{M} that can predict accurately on the test samples in $\mathcal{T}_{\text{test}}$. Most existing approaches start with training a *Source-Only* model \mathbf{M}^S on the labeled source data. For segmentation models with softmax output, the cross-entropy loss is widely used during the optimization process:

$$\mathcal{L}_{ce}(I^s, Y^s) = - \sum_{i=1}^H \sum_{j=1}^W \sum_{c=1}^C y_{(i,j),c}^s \log p_{(i,j)}(c|I^s; \mathbf{w}),$$

where (i, j) are the pixel coordinates in I^s , c is the category index, $y_{(i,j),c}^s \in \{0, 1\}$ is entry in the one-hot vector of the ground-truth label, i.e., $\forall (i, j), \sum_c y_{(i,j),c}^s = 1$, and $p_{(i,j)}(c|I^s; \mathbf{w})$ is the predicted category probability based on the model parameters \mathbf{w} .

Generally, \mathbf{M}^S has limited generalization capability and does not perform well on target samples due to domain discrepancy between the real and synthetic data. One common approach in UDA is using \mathbf{M}^S to generate pseudo-labels \hat{Y}^t on I^t , which enables a joint training on $\mathcal{S}_{\text{train}}$ and $\mathcal{T}_{\text{train}}$ based on the following loss function:

$$\mathcal{L}(\mathcal{S}_{\text{train}}, \mathcal{T}_{\text{train}}) = \sum_{\mathcal{S}_{\text{train}}} \mathcal{L}_{ce}(I^s, Y^s) + \lambda \sum_{\mathcal{T}_{\text{train}}} \mathcal{L}_{ce}(I^t, \hat{Y}^t),$$

where λ balances the loss between source and target domain. The joint training process encourages the learning of domain-invariant features and shared decision boundaries.

Different UDA methods have been proposed to find reliable pseudo-labels \hat{Y}^t [53, 99, 104], select a relevant subset of $\mathcal{S}_{\text{train}}$ [40], or add regularization terms and adversarial losses [41, 105] for the joint training process. However,

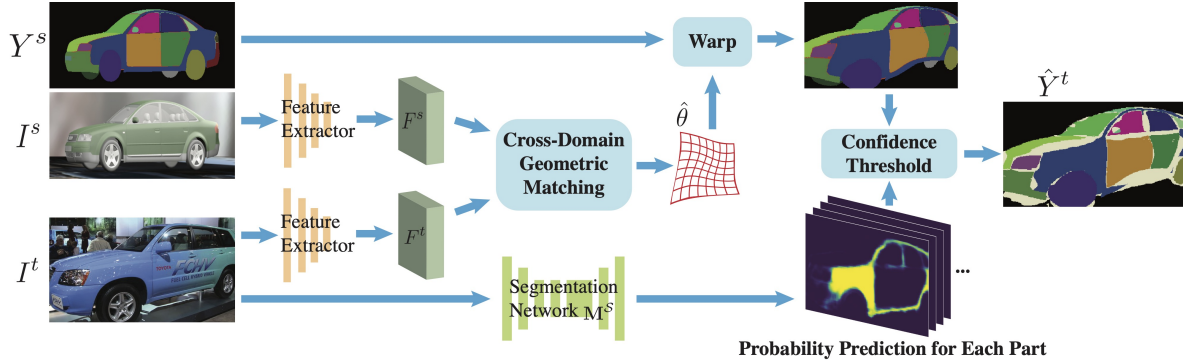


Figure 4. An illustration of Geometric-Matching Guided (GMG) domain adaptation. A pair of synthetic and real images are passed through a feature extractor to get their feature maps, and cross-domain geometric matching is used to estimate a 2-D transformation based on the feature similarities. The transformation is then applied to the segmentation ground-truth of the synthetic image to make it match the parts in the real image. Finally, a confidence threshold is used to filter out unreliable warping results and make high-quality pseudo-labels for the joint training step in unsupervised domain adaptation.

since most of these methods are designed for semantic segmentation, none of them takes advantage of the object’s spatial structure to predict parts. On the contrary, our proposed Geometric-Matching Guided domain adaptation (GMG) explicitly uses the structural relation between parts to generate pseudo-labels for joint training.

4.2. Cross-Domain Geometric Matching

In GMG, we use cross-domain geometric matching to find an optimal global transformation that can be used to transfer segmentation labels from synthetic to real. Cross-domain geometric matching has been previously explored for few-shot learning and correspondence learning [2, 102], but either with sparse alignment or strong supervision. Here we use it to facilitate the domain adaptation for part segmentation at dense pixel-level and require no labels on the real images in the whole training process.

Specifically, cross-domain geometric matching aims to find spatial correspondence between a pair of synthetic and real images (I^s, I^t). To achieve this, we optimize a global transformation function \mathbf{W}_θ that matches the two images I^s and I^t based on their feature similarities. \mathbf{W}_θ puts smoothness constraint on the matching, so it is able to preserve the spatial relations between parts. After the matching, the transformation is applied to transfer the synthetic label Y^s to the real image I^t as a pseudo-label \hat{Y}^t . In the following, we first assume that the input image pair is given and the objects in both images belong to the same prototype and have a similar viewpoint. We then discuss how to search such input pairs in an unsupervised manner. Note that cross-domain geometric matching is only used to generate pseudo-labels for the joint-training. At test time, neither a paired input nor a geometric matching is necessary in our framework.

Given (I^s, I^t) with similar appearances, we first use CNN convolutional layers to extract their feature maps. The output F^s and F^t are $h \times w \times d$ tensors, which can be interpreted as $h \times w$ grids of d -dimensional local features

$f^{(i,j)} \in \mathbb{R}^d$. The similarity between two feature vectors in F^s and F^t can be measured by cosine similarity:

$$\phi(f^{(i,j)}, f^{(k,l)}) = \frac{f^{(i,j)} \cdot f^{(k,l)}}{\|f^{(i,j)}\|_2 \|f^{(k,l)}\|_2}.$$

To emphasize the similarities are measured for features sampled from different domains, we use (i, j) to denote spatial coordinates in I^s , and (k, l) for spatial coordinates in I^t .

We then define a 2D geometric transformation function $\mathbf{W}_\theta : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ so the spatial correspondence between I^t and I^s could be found by $(k', l') = \mathbf{W}_\theta(k, l)$, where θ denotes the transformation parameters, (k', l') are the corresponding coordinates of (k, l) in I^s .

The quality of a geometric transformation can be measured by the sum of feature similarities at corresponding coordinates:

$$\Phi_\theta(F^s, F^t) = \sum_{(k,l)} \phi(f_{\mathbf{W}_\theta(k,l)}^s, f_{(k,l)}^t),$$

and our goal is to find the best parameters $\hat{\theta}$ such that:

$$\hat{\theta} = \arg \max_{\theta} \Phi_\theta(F^s, F^t).$$

In practice, we follow Rocco et al. [63] and use a spatial transformer layer [27] to implement the warping, which makes Φ differentiable w.r.t θ . Note that we just optimize the transformation parameters based on the feature similarities. The CNN backbone is fixed in this step.

Unsupervised selection of input pairs. For each real training image $I^t \in \mathcal{T}_{\text{train}}$, we perform a grid search over the viewpoint and prototype in the synthetic images and select the best one based on $\Phi_{\hat{\theta}}$. More specifically, we first build a pool of prototypical synthetic images by selecting samples from each prototype with 24 diverse viewpoints (i.e., azimuth angles sampled from $\{0, 30, 60, \dots, 330\}$ and elevation angles sampled from $\{5, 20\}$). Then, we perform geometric matching for I^t and each I^s in this pool. The synthetic image that achieves the highest $\Phi_{\hat{\theta}}$ is selected, and

its label is warped to infer the pseudo-label for I^t .

4.3. Confidence Threshold of Pseudo-Labels

Given $\hat{\theta}$, we can warp Y^s to I^t and use it as pseudo-supervision for joint training. However, the warping results can contain errors due to the variability of the object shapes and 3D poses, or due to a sub-optimal estimation of the transformation parameters. To correct such errors, in GMG, we use the confidence of the prediction provided by the *Source-Only* model \mathbf{M}^S . Specifically, for spatial coordinates (k, l) in I^t , if the corresponding coordinates $\mathbf{W}_{\hat{\theta}}(k, l)$ in I^s have ground-truth label \bar{c} , we use the predicted probability of \bar{c} by \mathbf{M}^S at (k, l) as a confidence score:

$$z_{kl} = p_{(k,l)}(\bar{c}|I^t; \mathbf{w}_S),$$

and threshold z_{kl} with γ to obtain the final pseudo-label:

$$\hat{y}_{(k,l)}^t = \begin{cases} y_{\mathbf{W}_{\hat{\theta}}(k,l)}^s, & \text{if } z_{kl} > \gamma \\ \mathbf{0}, & \text{otherwise,} \end{cases}$$

where $\mathbf{0}$ denotes vector with 0 entries everywhere, and is ignored by the cross-entropy loss. Consequently, GMG is able to select high-confidence warping results as pseudo-labels for joint training.

5. Experiments

In the experiments, different UDA methods are trained using the synthetic data in the source dataset and unlabeled real training images in the target dataset, and are then evaluated on the real test images in the target dataset.

Datasets and evaluation metric. In all experiments, we use the synthetic samples of the UDA-Part dataset as the source data. UDA-Part provides 30000 / 24000 / 24000 / 24000 / 24000 training samples for the object categories car / motorbike / aeroplane / bus / bicycle. We evaluate all methods on two target datasets. In the first set of experiments, **PASCAL3D+** [93] is used as the target dataset, which contains 2763 / 624 / 986 / 548 / 661 unlabeled training images for car / motorbike / aeroplane / bus / bicycle. After training, the domain adaptation models are evaluated on 200 (40 per category) real images in the test split, which are selected and annotated as the real test set of UDA-Part. In the second set of experiments, **PascalPart** [9] is used as the target dataset. We pre-process PascalPart data by cropping images to contain single object which leads to 538 / 261 / 266 / 221 / 252 training images, and 520 / 255 / 280 / 229 / 263) test images. Note that the segmentation labels for the PascalPart training images are not used. We train with the dense part labels defined in UDA-Part first and then merge the predictions to PascalPart label space during testing. All models are trained and tested on samples belonging to each vehicle category separately. **Mean Intersection over Union (mIoU)** is used as the metric for the part segmentation task, where IoU is first computed for each part and then averaged over all parts belonging to the corresponding category.

Baseline methods. For comparison purposes, we adapt

several popular UDA methods from related tasks to object part segmentation. BDL [41], CRST [105], and CAG [99] are all methods proposed for semantic segmentation but follow different strategies: BDL uses cycleGAN to reduce pixel-level domain discrepancy and encourage marginal feature alignment; CRST performs self-training with smoothness regularization; while CAG applies adversarial training as initialization and explores category-aware feature alignment during self-training. In addition, we test CCSSL [53], a self-training-based method designed for keypoint detection. CCSSL uses consistency constraints to select reliable pseudo-labels and applies strong data augmentation to improve the model’s generalization capability. The code for all baselines was adapted from the original public repositories and is modified to use the same backbone and input size.

Implementation details. We use DeepLabv3+ [8] as the segmentation network for GMG and all baseline methods. The weights are initialized from an ImageNet [12] pre-trained model. We implement our model using Pytorch [57] on two TitanX GPUs. Synthetic training images are resized to have a long edge of 800 pixels while real training images are resized to have a short edge of 224. For the geometric matching, we use thin plate spline transformation with 25 anchor points and take the first four convolutional blocks of an ImageNet [12] pre-trained VGG16 network [67] as the feature extractor. The confidence threshold γ is set to the 60th percentile of the scores obtained from all samples in the corresponding category. We implement the pair selection using python with 8 parallel CPU threads. For a pool of 24 candidates, it takes 2.1 seconds per image and roughly 54 (195) minutes for the training set of PascalPart (PASCAL3D+). When the ground-truth viewpoint is given, the matching takes less than 0.4 second per image, and thus roughly 10 (30) minutes for PascalPart (PASCAL3D+). During joint training, we apply strong augmentations to synthetic images following [53]. The joint training takes 10000 iterations and the learning rate is fixed at $2.5e - 4$. The real loss coefficient λ is set to 1.0 for all “w/vp” experiments and 0.1 for all others. More training details could be found in the supplementary material.

5.1. Main Results

In Table 2, we report object part segmentation results on the real test images of UDA-Part. In the header row, the number of parts is denoted in parenthesis next to the category name. The *Source-Only* model \mathbf{M}^S serves as a naïve baseline. We observe that the mIoU is generally low, indicating that the domain shift between UDA-Part synthetic images and real-world images indeed disturbs the model performance and domain adaptation is necessary. By looking at the overall performance, segmenting parts on *motorbike*, *aeroplane*, and *bicycle* are more challenging since their parts are often small, with irregular shape and self-

	car (32)	motorbike (22)	aeroplane (23)	bus (33)	bicycle (18)
Source Only M^S	37.19	15.91	13.71	19.43	18.01
CRST [105]	39.35	18.39	13.38	21.03	15.31
BDL [41]	43.60	17.15	16.71	22.95	17.78
CAG [99]	48.97	17.89	16.06	25.74	19.50
CCSSL [53]	49.31	21.87	17.50	24.75	17.27
GMG (Ours)	<u>49.93</u>	<u>23.09</u>	<u>17.89</u>	<u>25.78</u>	19.07
GMG (Ours) w/ vp	53.77	23.28	17.98	26.31	<u>19.09</u>

Table 2. Part segmentation results (mIoU) on UDA-Part real test images. The number of parts is denoted in parenthesis next to the category name. “GMG” uses unsupervised grid search to find the input pairs, while “GMG w/ vp” uses the ground-truth viewpoint of the real training images to reduce matching error. **Best** and **second best** UDA results are marked accordingly.

		car (14)	motorbike (7)	aeroplane (7)	bus (14)	bicycle (8)
Fully-Supervised Learning		40.36	38.08	42.47	34.42	40.57
UDA	Source Only M^S	14.24	20.48	23.18	16.02	20.21
	CRST [105]	14.44	24.71	23.04	16.56	22.57
	BDL [41]	19.02	26.89	29.08	17.29	22.45
	CAG [99]	18.39	24.09	28.60	17.12	22.22
	CCSSL [53]	24.23	28.80	32.58	18.59	22.25
	GMG (Ours)	<u>25.61</u>	<u>29.68</u>	<u>33.50</u>	<u>19.30</u>	<u>22.91</u>
GMG (Ours) w/ vp	27.59	30.73	33.98	21.20	23.63	

Table 3. Object part segmentation results (mIoU) on PascalPart test images. The number of parts is denoted in parenthesis next to the category name. “GMG” uses unsupervised grid search to find the input pairs, while “GMG w/ vp” uses the ground-truth viewpoint of the real training images to reduce matching error. **Best** and **second best** UDA results are marked accordingly.

occlusion, which can also be observed in Figure 2.

Among the baseline methods, CRST only achieves marginal improvement and may even hurt the performance in some cases, implying that simple smoothness regularization is not working very well on this task. Moreover, since our synthetic images are generated with random texture and background, it is hard to reduce pixel-level domain discrepancy, so the results of BDL are not satisfactory either. CAG achieves more improvement on *car* and *bus*, showing category-aware feature alignment may work well on parts with regular shapes and large areas. CCSSL improves most categories, suggesting part segmentation benefits more from consistency constraints and strong augmentations.

Compared with the baselines, GMG achieves very competitive results, outperforming all other methods in most cases. Furthermore, we introduce a variant of GMG, where pseudo-label errors are reduced by using the ground-truth viewpoints of the real training images during geometric matching (the grid search step only looks for the best prototype). The results from this model are shown in the row “GMG (Ours) w/ vp”, which show further improvement.

We report object part segmentation results on PascalPart test images in Table 3. For this dataset, we include part segmentation results from a DeepLabv3+ network directly trained from PascalPart labels on real images in a fully-supervised manner. These results can be considered as upper-bound for UDA approaches. Compared with UDA-Part, the PascalPart test set seems more manageable for some categories since the parts are more coarse-grained with larger sizes. On the other hand, the images in PascalPart are overall with lower resolution and contain more

truncated/occluded objects, making the segmentation task more challenging. Therefore, the result patterns on different categories are not the same as what we observe on UDA-Part test images. Despite the variations in individual test cases, GMG still outperforms all baseline methods on all categories. Similarly, adding ground-truth viewpoints during training consistently improves GMG performance. Comparing with the fully-supervised learning results, we are aware there is still a large room for improvement, indicating a relevant future research direction.

In Figure 5, we compare object part segmentation results from GMG and the *Source-Only* model. Generally, GMG is better at recognizing the shape and the boundary of parts. GMG can also eliminate wrong predictions that violate the part relations, implying it gains more knowledge about the object structure. On the other hand, GMG is prone to misclassifying small parts, such as the *mirror* of *car* and the *pedal* of *bicycle*, whose labels are harder to be correctly transferred through geometric matching. We consider it as future work to improve GMG performance on smaller parts.

5.2. Ablation and Model Diagnosis

GMG Variants		car	mtbk	arpl
Source-Only M^S		37.19	15.91	13.71
+warped labels	w/o vp	47.84	20.65	16.25
	w/ vp	53.26	21.94	17.22
+conf. thresh.	w/o vp	49.93	23.09	17.89
	w/ vp	53.77	23.28	17.98

Table 4. Ablation of GMG components. Results on UDA-Part real test images are shown. (mtbk: motorbike; arpl: aeroplane.)

We quantitatively evaluate the improvement introduced




Figure 5. Visualizations of GMG results. We show examples of part segmentation predictions on test images from UDA-Part dataset (column 1-5) and PascalPart dataset [9] (column 6,7). Row 1-4 are real test images, part segmentation ground-truths, *Source-Only* model predictions, and GMG predictions for each case respectively.

by different components of GMG in Table 4. Compared with the *Source-Only* model, a joint training based on directly warped synthetic labels leads to a performance gain by leveraging the pseudo-supervision on real samples. However, this method is more sensitive to misalignment caused by wrong geometric matching. Consequently, we observe a larger performance gap between the models that use viewpoint supervision (*w/ vp*) and those without (*w/o vp*). Applying confidence threshold improves the results in all cases, especially when ground-truth viewpoints are not available for real training samples. In summary, we can observe the effectiveness of both components in our model.

In input pair selection, we observe that thresholding ϕ can improve viewpoint estimation accuracy, but it leads to reduced number of pseudo-labels for joint training and eventually hurts the segmentation results for categories with less samples, e.g. motorbike. Searching in more azimuth bins indeed improves the viewpoint accuracy but has little positive effect on the final part segmentation mIoU, while searching in more elevation bins makes little difference in both metrics. We also observe GMG is generally not very sensitive to γ . When we change γ from 60th percentile of the scores to 50th (70th) percentile, the mIoU results for car on PASCAL3D+ changes from 53.77 to 53.65 (53.59). More details could be found in the supplementary results.

In Figure 6, we visualize GMG pseudo-label quality and the synthetic samples selected by grid search for the corresponding real images. In the first three examples, the search process successfully finds synthetic images with reasonable prototype and viewpoint. Consequently, geometric matching is performed on image pairs with similar appearances, and generates high-quality pseudo-labels for joint training. In the fourth row, we show a failure case where a wrong (i.e., opposite) viewpoint is selected. Nevertheless, most of the incorrect labels in this case are filtered out by confidence threshold and will not disturb the joint training in this case.




Figure 6. Quality of GMG pseudo-labels. Column 1-4 show the selected source label, the selected source image, the target image, and the final pseudo-label, respectively. Light yellow color indicates uncertain labels.

6. Conclusions

We introduce the idea of learning part segmentation for general objects from synthetic data through unsupervised domain adaptation (UDA). We first introduce UDA-Part, a comprehensive dataset for 5 vehicle categories designed to benchmark part segmentation through UDA. Extending UDA-Part to more object categories is one of our future goal. Secondly, we propose a new UDA algorithm called Geometric-Matching Guided domain adaptation (GMG) which leverages the object’s spatial structure to facilitate the adaptation process. In our experiments, GMG outperforms previous UDA methods on two real test image datasets, demonstrating the advantage of using structural information in UDA for part segmentation. On the other hand, GMG requires a grid-search process to find input pairs and is prone to misclassifying smaller parts, which could be improved in future works. In conclusion, our work provides a new solution for part segmentation on general objects with low cost and will motivate more research in this area.

Acknowledgement This work was supported by NSF BCS-1827427, NIH R01 EY029700, and ONR N00014-21-1-2812. We thank Zihao Xiao for proofreading.

References

- [1] Hossein Azizpour and Ivan Laptev. Object detection using strongly-supervised deformable part models. In *Eur. Conf. Comput. Vis.*, pages 836–849, 2012. [1](#)
- [2] Yutong Bai, Qing Liu, Lingxi Xie, Weichao Qiu, Yan Zheng, and Alan L Yuille. Semantic part detection via matching: Learning to generalize to novel viewpoints from limited training data. In *Int. Conf. Comput. Vis.*, pages 7535–7545, 2019. [3](#), [5](#)
- [3] Alexander C Berg, Tamara L Berg, and Jitendra Malik. Shape matching and object recognition using low distortion correspondences. In *IEEE Conf. Comput. Vis. Pattern Recog.*, volume 1, pages 26–33, 2005. [3](#)
- [4] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. [2](#)
- [5] Wei-Lun Chang, Hui-Po Wang, Wen-Hsiao Peng, and Wei-Chen Chiu. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1900–1909, 2019. [2](#)
- [6] Jianchun Chen, Lingjing Wang, Xiang Li, and Yi Fang. Arbicon-net: Arbitrary continuous geometric transformation networks for image registration. In *Adv. Neural Inform. Process. Syst.*, pages 3415–3425, 2019. [3](#)
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. [1](#)
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Eur. Conf. Comput. Vis.*, 2018. [6](#)
- [9] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1971–1978, 2014. [1](#), [2](#), [3](#), [4](#), [6](#), [8](#)
- [10] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12475–12485, 2020. [1](#)
- [11] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. [3](#), [4](#)
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009. [6](#)
- [13] Jian Dong, Qiang Chen, Xiaohui Shen, Jianchao Yang, and Shuicheng Yan. Towards unified human parsing and pose estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 843–850, 2014. [1](#)
- [14] Jiahua Dong, Yang Cong, Gan Sun, Yuyang Liu, and Xiaowei Xu. Cscsl: Critical semantic-consistent learning for unsupervised domain adaptation. In *Eur. Conf. Comput. Vis.*, pages 745–762, 2020. [3](#)
- [15] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Int. Conf. Comput. Vis.*, pages 2758–2766, 2015. [2](#)
- [16] Abhishek Dutta and Andrew Zisserman. The VIA annotation software for images, audio and video. In *ACMMM*, 2019. [4](#)
- [17] Hao-Shu Fang, Guansong Lu, Xiaolin Fang, Jianwen Xie, Yu-Wing Tai, and Cewu Lu. Weakly and semi supervised human body part parsing via pose-guided knowledge transfer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 70–78, 2018. [2](#)
- [18] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *Eur. Conf. Comput. Vis.*, pages 770–785, 2018. [1](#), [2](#)
- [19] Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 932–940, 2017. [1](#), [2](#)
- [20] Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow: Semantic correspondences from object proposals. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1711–1725, 2017. [3](#)
- [21] Kai Han, Rafael S Rezende, Bumsub Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. Sc-net: Learning semantic correspondence. In *Int. Conf. Comput. Vis.*, pages 1831–1840, 2017. [3](#)
- [22] Ankur Handa, Viorica Patraucean, Vijay Badrinarayanan, Simon Stent, and Roberto Cipolla. Understanding real world indoor scenes with synthetic data. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4077–4085, 2016. [2](#)
- [23] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, and Alan Yuille. Partimagenet: A large, high-quality dataset of parts. *arXiv preprint arXiv:2112.00933*, 2021. [2](#)
- [24] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *Int. Conf. Mach. Learn.*, pages 1989–1998, 2018. [2](#)
- [25] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. [2](#)
- [26] Jiaying Huang, Shijian Lu, Dayan Guan, and Xiaobing Zhang. Contextual-relation consistent domain adaptation for semantic segmentation. In *Eur. Conf. Comput. Vis.*, pages 705–722, 2020. [3](#)

- [27] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. *arXiv preprint arXiv:1506.02025*, 2015. 5
- [28] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2901–2910, 2017. 2
- [29] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7122–7131, 2018. 2
- [30] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4893–4902, 2019. 2
- [31] Guoliang Kang, Yunchao Wei, Yi Yang, Yueting Zhuang, and Alexander Hauptmann. Pixel-level cycle association: A new perspective for domain adaptive semantic segmentation. In *Adv. Neural Inform. Process. Syst.*, 2020. 2
- [32] Khalil Khan, Massimo Mauro, Pierangelo Migliorati, and Riccardo Leonardi. Head pose estimation through multi-class face segmentation. In *Int. Conf. Multimedia and Expo*, pages 175–180, 2017. 1
- [33] Jaechul Kim, Ce Liu, Fei Sha, and Kristen Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2307–2314, 2013. 3
- [34] Seungryoung Kim, Dongbo Min, Bumsub Ham, Sangryul Jeon, Stephen Lin, and Kwanghoon Sohn. Fcss: Fully convolutional self-similarity for dense semantic correspondence. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6560–6569, 2017. 3
- [35] Seungryoung Kim, Dongbo Min, Stephen Lin, and Kwanghoon Sohn. Dctm: Discrete-continuous transformation matching for semantic flow. In *Int. Conf. Comput. Vis.*, pages 4529–4538, 2017. 3
- [36] Tae Soo Kim, Bohoon Shim, Michael Peven, Weichao Qiu, Alan Yuille, and Gregory D. Hager. Learning from synthetic vehicles. *IEEE Winter Conf. on Appl. of Comput. Vis.*, 2022. In submission. 3
- [37] Adam Kortylewski, Bernhard Egger, Andreas Schneider, Thomas Gerig, Andreas Morel-Forster, and Thomas Vetter. Analyzing and reducing the damage of dataset bias to face recognition with synthetic data. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 0–0, 2019. 2
- [38] Vuong Le, Jonathan Brandt, Zhe Lin, Lubomir Bourdev, and Thomas S Huang. Interactive facial feature localization. In *Eur. Conf. Comput. Vis.*, pages 679–692, 2012. 1
- [39] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1
- [40] Guangrui Li, Guoliang Kang, Wu Liu, Yunchao Wei, and Yi Yang. Content-consistent matching for domain adaptive semantic segmentation. In *Eur. Conf. Comput. Vis.*, pages 440–456, 2020. 2, 4
- [41] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6936–6945, 2019. 2, 4, 6, 7
- [42] Xiaodan Liang, Liang Lin, Xiaohui Shen, Jiashi Feng, Shuicheng Yan, and Eric P Xing. Interpretable structure-evolving lstm. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1010–1019, 2017. 2
- [43] Xiaodan Liang, Si Liu, Xiaohui Shen, Jianchao Yang, Luqi Liu, Jian Dong, Liang Lin, and Shuicheng Yan. Deep human parsing with active template regression. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(12):2402–2414, 2015. 1, 2
- [44] Xiaodan Liang, Xiaohui Shen, Jiashi Feng, Liang Lin, and Shuicheng Yan. Semantic object parsing with graph lstm. In *Eur. Conf. Comput. Vis.*, pages 125–143, 2016. 2
- [45] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):978–994, 2010. 3
- [46] Xinchun Liu, Wu Liu, Jinkai Zheng, Chenggang Yan, and Tao Mei. Beyond the parts: Learning multi-view cross-part correlation for vehicle re-identification. In *ACMMM*, pages 907–915, 2020. 2
- [47] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440, 2015. 1
- [48] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Int. Conf. Mach. Learn.*, pages 97–105, 2015. 2
- [49] Wenhao Lu, Xiaochen Lian, and Alan Yuille. Parsing semantic parts of cars using graphical models and segment appearance consistency. In *Brit. Mach. Vis. Conf.*, 2014. 2
- [50] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2507–2516, 2019. 3
- [51] Umberto Michieli, Edoardo Borsato, Luca Rossi, and Pietro Zanuttigh. Gmnet: Graph matching network for large scale part semantic segmentation in the wild. In *Eur. Conf. Comput. Vis.*, pages 397–414, 2020. 2
- [52] Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 909–918, 2019. 2
- [53] Jiteng Mu, Weichao Qiu, Gregory D Hager, and Alan L Yuille. Learning from synthetic animals. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12386–12395, 2020. 2, 4, 6, 7
- [54] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4500–4509, 2018. 2

- [55] Xuecheng Nie, Jiashi Feng, and Shuicheng Yan. Mutual learning to adapt for joint human parsing and pose estimation. In *Eur. Conf. Comput. Vis.*, pages 502–517, 2018. 2
- [56] David Novotny, Diane Larlus, and Andrea Vedaldi. Anchornet: A weakly supervised network to learn geometry-sensitive features for semantic matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5277–5286, 2017. 3
- [57] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Adv. Neural Inform. Process. Syst.*, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [58] Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE Signal Processing Magazine*, 32(3):53–69, 2015. 2
- [59] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *AAAI*, 2018. 2
- [60] Aayush Prakash, Shaad Boochoon, Mark Brophy, David Acuna, Eric Cameracci, Gavriel State, Omer Shapira, and Stan Birchfield. Structured domain randomization: Bridging the reality gap by context-aware synthetic data. In *Int. Conf. Robotics and Automation*, pages 7249–7255, 2019. 4
- [61] N Dinesh Reddy, Minh Vo, and Srinivasa G Narasimhan. Carfusion: Combining point tracking and part detection for dynamic 3d reconstruction of vehicles. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1906–1915, 2018. 2
- [62] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6148–6157, 2017. 3
- [63] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6917–6925, 2018. 3, 5
- [64] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3752–3761, 2018. 2
- [65] Yifei Shi, Pinxin Long, Kai Xu, Hui Huang, and Yueshan Xiong. Data-driven contextual modeling for 3d scene understanding. *Computers & Graphics*, 55:55–67, 2016. 1
- [66] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1297–1304, 2011. 2
- [67] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [68] Xibin Song, Peng Wang, Dingfu Zhou, Rui Zhu, Chenye Guan, Yuchao Dai, Hao Su, Hongdong Li, and Ruigang Yang. Apollocar3d: A large 3d car instance understanding benchmark for autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5452–5462, 2019. 2
- [69] Yafei Song, Xiaowu Chen, Jia Li, and Qingping Zhao. Embedding 3d geometric features for rigid object part segmentation. In *Int. Conf. Comput. Vis.*, pages 580–588, 2017. 2
- [70] Michael Stark, Jonathan Krause, Bojan Pepik, David Meger, James J Little, Bernt Schiele, and Daphne Koller. Fine-grained categorization for 3d scene understanding. *Int. J. Robotics Res.*, 30(13):1543–1552, 2011. 1
- [71] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *Eur. Conf. Comput. Vis.*, pages 443–450, 2016. 2
- [72] Jian Sun and Jean Ponce. Learning discriminative part detectors for image classification and cosegmentation. In *Int. Conf. Comput. Vis.*, pages 3400–3407, 2013. 1
- [73] Tatsunori Tanai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4246–4255, 2016. 3
- [74] Alexander Thomas, Vittorio Ferrari, Bastian Leibe, Tinne Tuytelaars, and Luc Van Gool. Using recognition to guide a robot’s attention. In *Robotics: Science and Systems*, volume 4, pages 246–253, 2008. 1, 2, 3, 4
- [75] Daniel Töpfer, Jens Spehr, Jan Effertz, and Christoph Stiller. Efficient scene understanding for intelligent vehicles using a part-based road representation. In *IEEE Conf. Intell. Trans. Syst.*, pages 65–70, 2013. 1
- [76] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 969–977, 2018. 2, 4
- [77] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7472–7481, 2018. 2
- [78] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Int. Conf. Comput. Vis.*, pages 4068–4076, 2015. 2
- [79] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. 2
- [80] Nikolai Ufer and Bjorn Ommer. Deep semantic feature matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6914–6923, 2017. 3
- [81] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 109–117, 2017. 2

- [82] Jianyu Wang and Alan L Yuille. Semantic part segmentation using compositional model combining shape and appearance. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1788–1797, 2015. [2](#)
- [83] Jianyu Wang, Zhishuai Zhang, Cihang Xie, Vittal Premachandran, and Alan Yuille. Unsupervised learning of object semantic parts from internal states of cnns by population encoding. *arXiv preprint arXiv:1511.06855*, 2015. [2](#)
- [84] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Joint object and part segmentation using deep learned potentials. In *Int. Conf. Comput. Vis.*, pages 1573–1581, 2015. [2](#)
- [85] Yang Wang, Duan Tran, Zicheng Liao, and David Forsyth. Discriminative hierarchical part-based models for human parsing and action recognition. *J. Mach. Learn. Res.*, 13(10), 2012. [1](#)
- [86] Wikipedia contributors. List of bicycle parts — Wikipedia, the free encyclopedia, 2020. [Online; accessed 9-March-2021]. [3](#)
- [87] Wikipedia contributors. List of auto parts — Wikipedia, the free encyclopedia, 2021. [Online; accessed 9-March-2021]. [3](#)
- [88] Wikipedia contributors. Motorcycle components — Wikipedia, the free encyclopedia, 2021. [Online; accessed 9-March-2021]. [3](#)
- [89] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S Davis. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Eur. Conf. Comput. Vis.*, pages 518–534, 2018. [2](#)
- [90] Fangting Xia, Peng Wang, Liang-Chieh Chen, and Alan L Yuille. Zoom better to see clearer: Human and object parsing with hierarchical auto-zoom net. In *Eur. Conf. Comput. Vis.*, pages 648–663, 2016. [2](#)
- [91] Fangting Xia, Peng Wang, Xianjie Chen, and Alan L Yuille. Joint multi-person pose estimation and semantic part segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6769–6778, 2017. [2](#)
- [92] Fangting Xia, Jun Zhu, Peng Wang, and Alan Yuille. Pose-guided human parsing by an and/or graph using pose-context features. In *AAAI*, 2016. [2](#)
- [93] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE Winter Conf. on Appl. of Comput. Vis.*, pages 75–82, 2014. [1](#), [2](#), [4](#), [6](#)
- [94] Kota Yamaguchi, M Hadi Kiapour, Luis E Ortiz, and Tamara L Berg. Parsing clothing in fashion photographs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3570–3577, 2012. [1](#), [2](#)
- [95] Fan Yang, Xin Li, Hong Cheng, Jianping Li, and Leiting Chen. Object-aware dense semantic correspondence. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2777–2785, 2017. [3](#)
- [96] Yi Yang and Deva Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1385–1392, 2011. [1](#)
- [97] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Trans. Graph.*, 35(6):1–12, 2016. [1](#), [2](#), [4](#)
- [98] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based r-cnns for fine-grained category detection. In *Eur. Conf. Comput. Vis.*, pages 834–849, 2014. [1](#)
- [99] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *Adv. Neural Inform. Process. Syst.*, pages 433–443, 2019. [2](#), [3](#), [4](#), [6](#), [7](#)
- [100] Yi Zhang, Weichao Qiu, Qi Chen, Xiaolin Hu, and Alan Yuille. Unrealstereo: Controlling hazardous factors to analyze stereo vision. In [3](#). [2](#)
- [101] Yifan Zhao, Jia Li, Yu Zhang, and Yonghong Tian. Multi-class part parsing with joint boundary-semantic awareness. In *Int. Conf. Comput. Vis.*, pages 9177–9186, 2019. [2](#)
- [102] Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, and Alexei A Efros. Learning dense correspondence via 3d-guided cycle consistency. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 117–126, 2016. [3](#), [5](#)
- [103] Xingyi Zhou, Arjun Karapur, Chuang Gan, Linjie Luo, and Qixing Huang. Unsupervised domain adaptation for 3d keypoint estimation via view consistency. In *Eur. Conf. Comput. Vis.*, pages 137–153, 2018. [2](#)
- [104] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Eur. Conf. Comput. Vis.*, pages 289–305, 2018. [2](#), [4](#)
- [105] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Int. Conf. Comput. Vis.*, pages 5982–5991, 2019. [2](#), [4](#), [6](#), [7](#)