

UDA-COPE: Unsupervised Domain Adaptation for Category-level Object Pose Estimation

Taeyeop Lee

Byeong-Uk Lee

Inkyu Shin

Jaesung Choe

Ukcheol Shin

In So Kweon

Kuk-Jin Yoon

KAIST

Abstract

Learning to estimate object pose often requires ground-truth (GT) labels, such as CAD model and absolute-scale object pose, which is expensive and laborious to obtain in the real world. To tackle this problem, we propose an unsupervised domain adaptation (UDA) for category-level object pose estimation, called UDA-COPE. Inspired by recent multi-modal UDA techniques, the proposed method exploits a teacher-student self-supervised learning scheme to train a pose estimation network without using target domain pose labels. We also introduce a bidirectional filtering method between the predicted normalized object coordinate space (NOCS) map and observed point cloud, to not only make our teacher network more robust to the target domain but also to provide more reliable pseudo labels for the student network training. Extensive experimental results demonstrate the effectiveness of our proposed method both quantitatively and qualitatively. Notably, without leveraging target-domain GT labels, our proposed method achieved comparable or sometimes superior performance to existing methods that depend on the GT labels.

1. Introduction

Object pose estimation is one of the crucial tasks used in various robotics and computer vision applications for robot manipulation [8, 34, 37, 39] and augmented reality (AR) [23, 24, 28]. Using sensor data such as images or point clouds, this task aims to estimate the poses of target objects including 3D orientation, 3D location, and size information.

Previous 6D object pose estimation methods follow the instance-level pose estimation schemes [12, 13, 25, 27, 31, 34, 38] that rely on given 3D CAD model information (*e.g.*, keypoints, geometry) and the size of known objects. However, these methods typically have difficulty estimating the pose of unknown objects since they do not yet have 3D CAD models as priors.

In contrast to the instance-level scheme, category-level object pose estimation [4, 5, 20, 30, 35, 36] approaches are more efficient in that a single network can infer multiple classes at once. In particular, Wang *et al.* [35] introduced a pioneering representation called Normalized Object Coordinate Space (NOCS), to align different object instances within one category in a shared 3D orientation. By estimating per-category NOCS maps, it is able to estimate the 6D pose of unseen objects without prior 3D CAD models. Its strengths have led to the use of NOCS representation in the following studies [4, 5, 20, 30, 36].

However, current object pose estimation research mostly relies on supervised learning, which requires expensive GT labels such as 3D object CAD models and absolute object pose. These labels are not only difficult to obtain in the real world but are also unreliable due to the human-annotation. Because of this difficulty, most of the training depends on synthetic datasets [15, 29, 31] and is usually not feasible in real-world applications due to domain gaps.

To cope with the real-world data scarcity problem, we take a look at unsupervised domain adaptation (UDA) methods [14, 19, 42]. UDA approaches often consider two types of datasets, the source domain (*i.e.* synthetic dataset) and the target domain (*i.e.* real-world dataset) dataset. The main goal of the UDA methods is to successfully make deep learning networks robust to the target domain using only the GT labels of the source domain. Various techniques exist, such as pseudo label generation [14, 19], teacher and student networks with momentum updates [1, 40], adversarial learning [2, 3, 16], and etc.

In this paper, we propose an Unsupervised Domain Adaptation for Category-level Object Pose Estimation (UDA-COPE). The proposed method effectively transfers task knowledge from a synthetic domain to a real domain by exploiting a multi-modal self-supervised learning scheme using pseudo labels. Our UDA-COPE concentrates on how to make high-quality pseudo-labels that are efficiently tar-

geted for the category-level pose estimation task. To this end, we designed *bidirectional point filtering* to remove noisy and inaccurate points based on pose optimization. Extensive experiments demonstrate that our UDA-COPE and bidirectional point filtering successfully can reduce the domain gap between synthetic and real datasets. Moreover, our framework achieved better performance than the previous supervised methods [4, 30, 35, 36]. The contributions of our method are summarized as:

- We propose an RGB-D based Unsupervised Domain Adaptation for Category-level Object Pose Estimation (UDA-COPE) framework that addresses the problem of data deficiency in real-world scenarios.
- We design a teacher-student framework where high-quality pose-aware pseudo labels can be obtained via the proposed bidirectional point filtering.
- Our method shows comparable or sometimes better results than supervised pose estimation approaches.

2. Related Work

Category-level object pose estimation. This task [4–6, 20, 30, 35] deals with objects of unseen instances but known categories. The most recent category-level object pose and size methods [5, 17, 30, 35, 36] use the dense Normalized Object Coordinate Space (NOCS) representation as a basic way to estimate pose. It aligns different object instances within one category in a shared 3D orientation. Shape Prior [30] improves the quality of NOCS maps by generating a representative shape prior and a deformation NOCS map for each category. CR-Net [36] extends the shape-prior method [30] by using the cascade relation and recurrent reconstruction methods. Similarly, SGPA [5] proposes a prior adaptation method.

Despite notable improvements, these studies rely on fully supervised learning, and as such, they require a large amount of manually labeled data, such as the 6D pose of objects, 3D CAD models, and NOCS maps. It is also time-consuming and expensive to create accurate GT information in real-world scenarios. To the best of our knowledge, CPS++ [22] is the only existing approach that addresses this data dependency problem. It aims to predict 3D shape using an RGB image to optimize pose and utilizes an unsupervised learning scheme by computing consistency between the observed depth map and the rendered depth. The rendered depth is obtained by projecting an estimated 3D shape with the predicted pose. However, 3D shape reconstruction from a single image is challenging and makes their unsupervised guidance unreliable. Recently, Li *et al.* [18] suggested a self-supervised method that leverages SE(3) equivalent representation for category-level pose estimation. However,




Figure 1. Network for multi-modal NOCS map estimation.

they do not fully consider pose parameters at the category-level that consist of 3D orientation, 3d location, and size information.

Multi-modal UDA. xMUDA [14] is a pioneer approach for unsupervised domain adaptation (UDA) in 2D/3D semantic segmentation for multi-modal scenarios. A few methods have considered an extra modality (*i.e.*, depth) during training time and leveraged such privileged information to boost adaptation performance [16, 33]. Reza *et al.* [21] proposed a multi-modal UDA in instance-level object pose estimation, but it only considered relative pose on a 2D image level. However, multi-modal UDA approaches have not yet been explored in the category-level object pose estimation task.

3. Method

Given an RGB image I , point cloud P , and segmentation labels S , our architecture aims to regress the 6D pose and size $s \in \mathbb{R}^3$ of objects. The 6D pose is defined as the rigid transformation of $[R|t]$: rotation $R \in SO(3)$, and translation $t \in \mathbb{R}^3$. Following previous studies [4, 5, 20, 30, 36], the segmentation labels S are used to crop the RGB images and point clouds. We leverage the NOCS representation to align different object instances within one category in a shared orientation. The categorical object pose $[R|t]$ and size s are estimated by Umeyama algorithm [32] with RANSAC [9], which optimizes $[R|t]$ and s by minimizing the distances between point cloud P and an estimated NOCS map N .

We first illustrate our network architecture (Sec. 3.1). Then, we introduce training methods of supervised learning using synthetic dataset (Sec. 3.2) and unsupervised domain adaptation using real-world dataset (Sec. 3.3).

3.1. Network Architecture

Recent category-level object pose estimation methods [4, 20, 30] take an RGB-D input to extract the 2D/3D features. We designed separate 2D/3D branches to extract features from both modalities. We use PSPNet [41] with ResNet34 [11] for 2D feature extraction and the Mink16UNet34 [7] for 3D feature extraction. At this time, the 2D feature is extracted by sampling features that are validly matched with point cloud P from the feature vol-

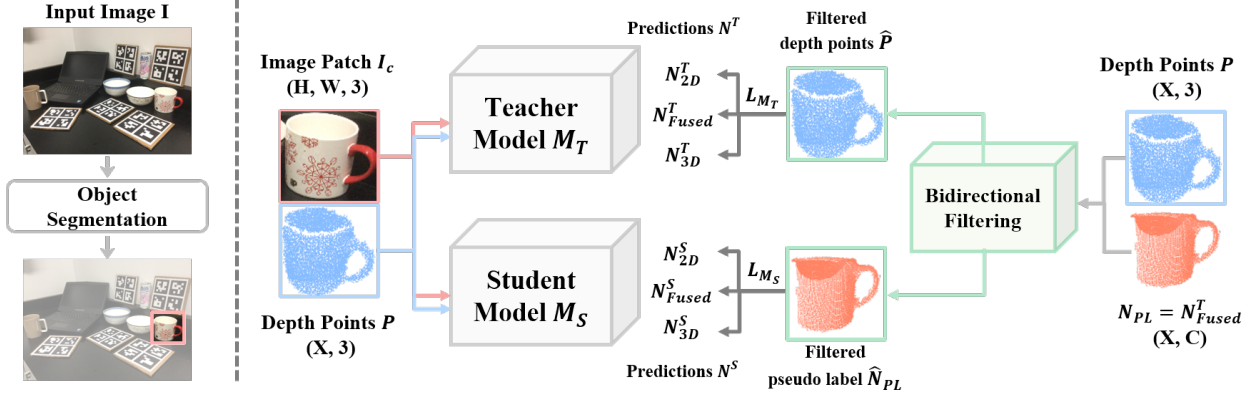


Figure 2. **Overview of unsupervised domain adaptation for category-level object pose estimation (UDA-COPE).** UDA-COPE utilizes pseudo label based teacher/student training scheme. Our proposed bidirectional point filtering method removes the noisy pseudo labels and gives reliable guidance to the student network. At the same time, filtered depth points gives additional self-supervision to the teacher network so that it can be robust to the domain gap between the synthetic and real dataset.

ume. Finally, we have a fused branch that combines each feature from both branches. Every branch estimates a NOCS map (N) with a separate NOCS header, which consists of three multi-layer perceptrons (MLP) layers. Our multi-modal NOCS map prediction network is illustrated in Fig. 1. We designate the NOCS map estimation of each branch as N_{2D} , N_{3D} , N_{Fused} , according to respective feature property.

3.2. Pre-Training with Synthetic Data

Inspired by pseudo label (PL) based methods [14, 19], our method consists of a teacher and a student model. Fig. 2 shows the overview of our teacher and student model. The initial prediction of teacher model M_T becomes a pseudo label N_{PL} for a student model, and student model M_S learns from the pseudo label as a GT. Our teacher and student model have the same structure as was described in Sec. 3.1.

We first train our teacher model in a supervised manner using the labeled synthetic dataset. For the NOCS map prediction using the GT information, we utilize cross-entropy loss, as in, $H(N_{gt}, N^T)$, where the supervision is given to all predictions from three branches. Additionally, to make our teacher network more robust, we apply the 2D image and 3D points augmentation and use consistency loss L_C so that each modality can output consistent results. Total loss for the teacher network on the synthetic dataset is formulated as follows:

$$\begin{aligned} L_{M_T} &= \lambda_N H(N_{gt}, N^T) + \lambda_C L_C, \\ L_C &= H(N^T, N_{Aug}^T) \end{aligned} \quad (1)$$

where N_{Aug}^T is the NOCS map prediction from the augmented input, and λ_N and λ_C are weighting parameters. Notation for the modality of the predictions is discarded for better readability.

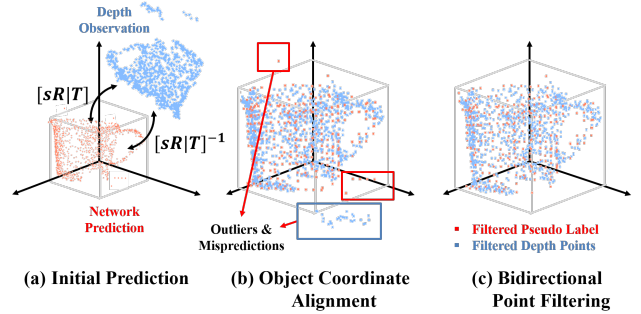


Figure 3. **Overview of bidirectional point filtering method.** Given pseudo labels and depth points (a), we estimate the pose and size using the Umeyama [32] algorithm and RANSAC [9], and align the depth points to normalized object coordinate (b). The pseudo label and aligned depth points (blue) have noisy and inaccurate points. After our bidirectional point filtering, the noisy points are removed to give more reliable supervision for both teacher and student (c).

3.3. Pose-Aware Unsupervised Domain Adaptation

After training from the synthetic dataset, the most straight-forward yet naive approach is to train the student network using the prediction of the teacher network. However, using the initial prediction from the teacher model as a pseudo label can be risky. The risk is due to the lack of robustness of the teacher model itself, or more importantly, because of the insufficient knowledge that the teacher model holds with respect to real-world scenarios, due to the domain gap between the simulated and real worlds. Techniques such as data augmentation and momentum update might help the feasibility but are still restricted. Therefore, we need additional guidance for the teacher model to estimate high-quality predictions, and more reliable pseudo labels for our student model to learn from.

3.3.1 Bidirectional Point Filtering

To solve these problems, we propose the bidirectional point filtering method which simultaneously removes the noise of the pseudo labels for the student and filters noisy depth points P for a teacher network. Fig. 3 shows an overview of the proposed bidirectional filtering method. Our bidirectional filtering method uses the P and N_{PL} as input and initially estimates the pose $[R|t]$ and size s using the Umeyama algorithm [32] with RANSAC [9]. Then it aligns the depth points P to the NOCS coordinate by applying the inverse of the estimated pose, as in multiplying the matrix $[sR|t]^{-1}$. We denote aligned depth points as P' . And then we calculate the point-wise 3D distance d between the aligned depth points P' and pseudo label N_{PL} to filter out noisy points from both sides using ρ as the threshold. Finally, we get the refined pseudo label \hat{N}_{PL} and filtered aligned depth \hat{P} . Our bidirectional point filtering can be expressed as:

$$\begin{aligned} d(n) &= \|P'(n) - N_{\text{PL}}(n)\| \quad \text{where } \forall n \in [1|P'], \\ \hat{N}_{\text{PL}} &= \{N_{\text{PL}}(n) : d(n) < \rho\}, \\ \hat{P} &= \{P'(n) : d(n) < \rho\}, \end{aligned} \quad (2)$$

Fig. 3 shows that our bidirectional filtering method removes outliers of pseudo label N_{PL} and depth P , and results in refined pseudo label \hat{N}_{PL} and filtered depth points \hat{P} .

3.3.2 Self-Supervised Learning

After the bidirectional filtering, we jointly train the teacher network and student network using the filtered pseudo labels \hat{N}_{PL} and filtered aligned depth points \hat{P} . Noted that we only use the filtered points \hat{P} for the teacher training, which may be a smaller subset of an original P . We use cross-entropy loss to train a student model using clean pseudo labels \hat{N}_{PL} . The student model loss is defined as:

$$L_{M_S} = -\frac{1}{|\hat{N}_{\text{PL}}|} \sum_{n=1}^{|\hat{N}_{\text{PL}}|} H(\hat{N}_{\text{PL}}(n), N^S(n)), \quad (3)$$

where N^S is the predictions of our student network. At the same time, the teacher learns real data knowledge from observation. We use cross-entropy loss by utilizing geometric consistency between our filtered aligned depth \hat{P} and estimated pseudo labels N^T . The teacher model loss is defined as:

$$L_{M_T} = -\frac{1}{|\hat{P}|} \sum_{n=1}^{|\hat{P}|} H(\hat{P}(n), N^T(n)). \quad (4)$$

We train our teacher model with a small learning rate for stable teacher network training. For both teacher and student models, we compute the loss for all estimations, N_{2D} , N_{3D} , N_{Fused} , which shows better result than applying the loss to only N_{Fused} . We denote all estimation losses as all modality (AM) loss.

4. Experiments

Datasets. We employ two commonly used category-level pose estimation datasets, a synthetic dataset and a real dataset. The synthetic dataset (source domain) is the Context-Aware MixedEd ReAlity (CAMERA) dataset [35], generated by rendering and compositing synthetic objects into real images in a context-aware manner. It comprises 275k synthetic images for training. The dataset contains 1085 object instances selected from 6 different categories - bottle, bowl, camera, can, laptop and mug. We use the REAL dataset [35] as a real dataset (target domain). Compared to previous methods [4, 5, 30, 35], our method trains without object pose, 3D CAD, and NOCS map GT labels of the target domain. It consists of 43,000 real world images of 7 scenes for training and 2,750 real world images of 6 scenes for evaluation. The REAL evaluation set was designated REAL275. We evaluated our method on a standard REAL275 benchmark for the task of category-level object pose estimation.

Metrics. We followed the previous pose and size evaluation metric from Wang *et al.* [35], which evaluated the performance of 3D object detection and 6D pose estimation. We report the average precision at different Intersection-Over-Union (IoU) thresholds for 3D object detection. Threshold values of 25%, 50%, and 75% were used to evaluate the results. For 6D object pose evaluation, the average precision was computed considering rotation and translation errors. For example, the 10° , 10cm metric denotes the percentage of object instances where the prediction error was less than 10° and 10cm.

4.1. Implementation Details

We implemented our method with PyTorch [26]. To obtain object regions from the real-world dataset, we used the GT segmentation labels during training and the off-the-shelf object segmentation method, Mask R-CNN [10], during inference. For a detected instance, we resized the image patch to 192×192 , and randomly sampled 1024 point clouds. The 2D and 3D features had 64 dimensions and the fused feature had 128 dimension features, due to concatenation. We trained our teacher network for 50 epochs on the synthetic dataset. We used the Adam optimizer with an initial learning rate of 0.0001 and a batch size of 16, where the learning rate was decreased by a factor of 0.6, 0.3, 0.1, 0.01 at 15k, 30k, 45k, 60k iterations, respectively. We set $\lambda_N = 1.0$, $\lambda_C = 1e-6$ for the teacher network. The ρ was set to 0.05 for our proposed bidirectional filtering. On the unlabeled real dataset, both the teacher and student networks were trained for 1 epoch. The learning parameters for our student network were the same as those for the teacher training scheme on the synthetic dataset. We set the learning rate for real-world self-supervised learning for the teacher network to be $1e-7$.

Method	Input	Syn	Real		mAP (\uparrow)					
			w/ Label	w/o Label	$3D_{50}$	$3D_{75}$	$5^\circ 2cm$	$5^\circ 5cm$	$10^\circ 2cm$	$10^\circ 5cm$
CPS++ [22]	RGB	✓			72.6	-	-	25.8	-	-
Metric Scale [17]	RGB	✓			68.1	32.9	2.2	5.3	10.0	24.7
NOCS [35]	RGB	✓			36.7	3.4	-	3.4	-	20.4
SPD [30]	RGB-D	✓			71.0	43.1	11.4	12.0	33.5	37.8
NOCS [35]	RGB	✓	✓		78.0	30.1	7.2	10.0	13.8	25.2
SPD [30]	RGB	✓	✓		75.2	46.5	15.7	18.8	33.7	47.4
SPD [30]	RGB-D	✓	✓		77.4	53.5	19.5	21.6	43.5	54.0
CASS [4]	RGB-D	✓	✓		77.7	-	-	23.5	-	58.0
CR-Net [36]	RGB-D	✓	✓		79.3	55.9	27.8	34.3	47.2	60.8
DualPoseNet [20]	RGB-D	✓	✓		79.8	62.2	29.3	35.9	50.0	66.8
SGPA [5]	RGB-D	✓	✓		80.1	61.9	35.9	39.6	61.3	70.7
CPS++ [22]	RGB	✓		✓	72.8	-	-	25.2	-	-
Ours	RGB	✓		✓	82.0	59.0	24.4	27.0	49.3	54.8
Ours	D	✓		✓	79.6	57.8	21.2	29.1	48.7	65.9
Ours	RGB-D	✓		✓	82.6	62.5	30.4	34.8	56.9	66.0

Table 1. **Quantitative comparison with state-of-the-art methods on the REAL275 dataset.** Empty entries either could not be evaluated or were not reported in the original paper.

Method	Syn	Real		mAP (\uparrow)						
		w/o Label	$3D_{25}$	$3D_{50}$	$3D_{75}$	$5^\circ 2cm$	$5^\circ 5cm$	$10^\circ 2cm$	$10^\circ 5cm$	$10^\circ 10cm$
CPS++ (RGB)	✓		84.5	72.6	-	-	25.8	-	-	55.4
CPS++ (RGB)	✓	✓	84.6 (+0.1)	72.8 (+0.2)	-	-	25.2 (-0.6)	-	-	58.6 (+3.2)
Ours (RGB)	✓		83.3	79.9	49.7	15.4	18.3	37.6	46.7	48.9
Ours (RGB)	✓	✓	83.8 (+0.5)	82.0 (+2.1)	59.0 (+9.3)	24.4 (+9.0)	27.0 (+8.7)	49.3 (+11.7)	54.8 (+8.1)	56.9 (+8.0)
Ours (RGB-D)	✓	✓	84.0 (+0.7)	82.6 (+2.7)	62.5 (+12.8)	30.4 (+15.0)	34.8 (+16.5)	56.9 (+19.3)	66.0 (+19.3)	68.3 (+19.4)

Table 2. **Quantitative comparison of unsupervised pose estimation approaches on the REAL275 dataset.** Empty entries are either not able to be evaluated or not reported in the original paper. Performance margins are calculated compared to the synthetic-only results.

4.2. Comparison with State-of-the-art

We compared our methods with state-of-the-art methods that were trained on different datasets and labels: 1) labeled synthetic dataset, 2) labeled synthetic and real datasets, 3) labeled synthetic and unlabeled real datasets. All methods were evaluated on the REAL275 dataset. Note that only the approaches with the ability to perform multi-class category-level pose estimation using a single network were considered. RGB, Depth and RGB-D denotes the modality of the network input, and most of the RGB based approaches utilize depth information in the pose optimization or refinement process.

Supervised Pose Estimation methods. Table 1 summarizes the results of the state-of-the-art category-level object pose estimation methods. Obviously, supervised training with the real data annotation significantly improved the overall performance, as are revealed by comparing the results of NOCS [35] and SPD [30] on different training dataset conditions. However, our unsupervised method showed results superior to NOCS [35], SPD [30], CASS [4], and CR-Net [36]. Compared to two of the strongest previous approaches, SGPA [5] and DualPoseNet [20], ours still showed comparable performance. This indicates that our proposed filtered pseudo label based UDA-COPE is robust

when estimating object pose in unseen real-world instances.

Unsupervised Pose Estimation methods. Table 2 summarizes the results of CPS++ and our method on source only, and source with unlabeled target training conditions. CPS++ [22] provides self-supervision by computing the consistency between the observed depth map and the rendered depth. The rendered depth is obtained by projecting an estimated 3D shape with the predicted pose. The results from row 1 and row 2 in Table 2 show that for CPS++, using unlabeled real data marginally improved performance, and sometimes even worsened it, as in $5^\circ, 5cm$ metric. We believe that their self-supervision is unreliable because of ambiguous 3D shape reconstruction using only a single-view RGB image.

Comparing row 3 and row 4, it can be seen that our proposed method shows improved results for every metrics, with some metrics showing notable margins such as an 8.7 mAP (48%) increase in $5^\circ, 5cm$. Also, in the last row, our RGB-D result had better performance than the single modality based outputs. Therefore, we claim that our proposed algorithm is more effective by utilizing a pseudo label based learning scheme with modality and pose-aware self-supervision. The effectiveness of each components will be ablated thoroughly in the following sections.

Method	RGB				Depth				RGB-D			
	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$
Lower Bound	79.9	49.7	15.4	18.3	76.7	52.7	14.9	22.7	80.7	60.9	23.0	27.9
PL	78.7	51.0	11.6	13.5	76.2	51.0	15.4	23.0	80.3	58.6	23.8	28.2
PL + MU	79.0	52.8	11.9	13.6	76.9	51.3	14.7	22.7	80.0	58.3	23.3	27.8
PL + AM	80.2	56.2	21.2	24.8	77.1	57.9	18.2	24.7	81.6	60.8	24.6	29.2
PL + xMUDA	78.4	55.3	21.9	25.5	76.5	56.6	18.3	26.0	80.7	60.2	25.7	30.3
PL-F + AM	81.5	59.0	23.3	26.1	77.3	58.4	20.0	27.3	81.3	62.0	28.1	32.9
PL-F + AM + TSL	82.0 (+2.1)	59.0 (+9.3)	24.4 (+9.0)	27.0 (+8.7)	79.6 (+2.9)	57.8 (+5.1)	21.2 (+6.3)	29.1 (+6.4)	82.6 (+1.9)	62.5 (+1.6)	30.4 (+7.4)	34.8 (+6.9)
Upper Bound	82.7	66.7	29.3	32.8	79.9	64.9	23.7	29.6	82.9	70.4	31.8	35.8

Table 3. **Ablation studies on UDA components.** Lower Bound: trained with labeled source only, Upper Bound: trained with both labeled source and target, PL: Pseudo Label, MU: Momentum Update, AM: All Modality loss, PL-F: Pseudo Label Filtering, TSL: Teacher Self-supervised Learning (4). Performance margins are calculated compared to the Lower Bound.

4.3. Ablation Study

Our ablation studies were conducted using the predictions from all branches: RGB, Depth, and RGB-D. The lower bound and upper bound in Table 3 are the results from using a single network described in Sec. 3.1. The lower bound was trained source only, while the upper bound utilized both source and target data with their labels. Table 3 summarizes the results of the ablation studies.

Naive Teacher/Student UDA for Pose Estimation. To verify the effectiveness of our proposed pose-aware UDA methods, we first applied a very naive version of teacher-student learning scheme with the momentum update. When applying the momentum update, we updated our teacher network every 100 iterations with a γ of 0.999.

However, it can be seen that neither methods were not very helpful. The possible reasons are: 1) directly using the initial predictions of our teacher network as a pseudo label was unreliable due to the domain gap and noisy prediction, 2) providing supervision only to the fused branch may restrict effective training of the 2D and 3D branches, which disrupt the fused branch after all, and 3) the momentum update using the gradient information retrieved from the unreliable pseudo label and thereby unreliable student network training were insufficient for the teacher network to learn real world adaptation.

Modality-Aware UDA methods. We applied all modality (AM) loss which computes cross-entropy loss between a pseudo label and all outputs from the three branches. Comparing (PL) and (PL + AM) in Table 3, shows that the performances of the RGB-based prediction and Depth-based prediction improved by large margins. Accordingly, we achieved improved accuracy from the RGB-D-based fused branch as well.

To compare our methods with the previous multi-modal unsupervised domain adaptation method, we also applied xMUDA [14] which constraints the 2D feature and 3D feature for consistency. xMUDA consistency was given during the student network training. xMUDA showed a similar

amount of performance boost as our AM loss. This is because both AM and xMUDA help our network to output consistent predictions between N_{2D} , N_{3D} , and N_{Fused} . We chose AM since it can provide more straight-forward supervision to each branch.

Pose-Aware UDA methods. We compared our two major components, which were specifically designed for pose estimation. As shown in Table 3, pseudo label filtering (PL-F + AM) resulted in significant improvements compared with naive pseudo label (PL + AM), which indicates that providing more reliable and confident pseudo labels to our student network is important. More detailed comparisons of various approaches of selecting confident pseudo label are explained in the next session. Using AM and PL-F, our student network was more robustly trained on real world data, and now holds meaningful knowledge which can be passed back to the teacher network. Utilizing teacher self-supervised learning (TSL) at this stage results in a notable increase in performance, compared to how basic momentum update technique performed from (PL) to (PL + MU).

Pseudo Label Filtering. To show the effectiveness of our pseudo label filtering based on the proposed bidirectional point filtering, we experimented with well-known pseudo label filtering techniques. The candidates were, Top k (conf), Top k (conf, class-wise), Entropy, SoftMax Max, SoftMax Avg., and ArgMax Match. Top k filterings use the top $k\%$ predictions of the pseudo label N_{PL} based on the softmax values, where class-wise does the sorting/filtering within all instances in each class. Entropy filtering uses the bottom $k\%$ of the predictions regarding its entropy. For SoftMax Max, SoftMax Avg., and ArgMax Match, we considered all three outputs from the teacher network, N_{2D}^T , N_{3D}^T , N_{Fused}^T . Among the three predictions, the filtered pseudo label was generated by selecting or compositing values that were either the maximum of the softmax, the average of the softmax, or having indices match of the argmax operation. We set k to be 50, regarding the filtering approximate ratio of bidirectional point filtering.

Filter	RGB				Depth				RGB-D			
	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$
None	80.2	56.2	21.2	24.8	77.1	57.9	18.2	24.7	81.6	60.8	24.6	29.2
Top k (conf)	79.0	55.5	21.6	25.0	76.7	56.4	18.1	24.8	81.4	60.4	25.6	30.1
Top k (conf, class-wise)	78.8	55.0	20.4	23.8	76.7	56.2	18.1	24.9	80.9	59.7	25.3	29.8
Entropy	77.6	54.9	24.1	27.7	77.7	57.0	18.4	25.3	80.6	60.8	26.6	31.1
SoftMax Max	79.9	52.7	19.6	22.4	76.9	56.3	17.4	25.5	79.9	59.2	25.6	30.0
ArgMax Match	79.6	52.2	17.7	19.9	76.2	56.0	20.1	27.6	81.6	60.9	24.1	27.5
Softmax Avg.	79.1	55.9	21.8	25.3	77.1	55.9	18.6	26.1	81.3	61.9	25.5	30.0
Ours	81.5 (+1.3)	59.0 (+2.8)	23.3 (+2.1)	26.1 (+1.3)	77.3 (+0.2)	58.4 (+0.5)	20.0 (+1.8)	27.3 (+2.6)	81.3 (-0.3)	62.0 (+1.2)	28.1 (+3.5)	32.9 (+3.7)

Table 4. **Ablation study of the pseudo label filtering methods.** Performance margins were calculated as compared to the results without using any pseudo label filtering.




Figure 4. **Noisy GT label examples of the Real training dataset.** Human-annotated GT pose labels on the real dataset (top row) are sometimes more inaccurate than our predicted pseudo labels (bottom row).

Compared to other filtering methods, our filtering method achieved the best performance overall, as can be seen in Table 4. This suggests that our pose-aware bidirectional point filtering generates more reliable pseudo label. This is because while other filter methods only concentrate on the predicted logit itself, while ours consider the relationship between estimated NOCS map and observed depth points. Therefore, our approach can efficiently remove pose and depth-aware outliers.

Reliability of our Pseudo label. Fig. 4 visualizes some examples of the real training set with GT labels and our pseudo labels. 6D poses were obtained and visualized using the Umeyama algorithm [32], using GT NOCS map and our pseudo label NOCS map. The real data annotations were mainly performed automatically using aruco markers. For some of the failure cases, additional ICP or manual human annotations were needed. Therefore, frames with inaccurate labels exist, which might disrupt supervised training.

However, there were cases where our method generated more accurate annotations than the GT provided by the

dataset. For example, the GT image of the first column shows that all objects have relatively wrong poses due to the occluded aruco markers in the image, while our prediction seems much more reliable. It shows that our pseudo label is not only reliable but also sometimes more accurate than the GT, which means that our proposed approach successfully addresses both the real-world data scarcity and unreliability issues.

4.4. Qualitative Results

Fig. 5 shows qualitative results on the REAL275 dataset. We compare our results with some of the supervised methods, NOCS [35], SPD [30] and DualPoseNet [20]. Our method estimated pose and sizes more accurately than NOCS and SPD, especially on cameras and laptops. Compared to the state-of-the-art approach, DualPoseNet, ours exhibited comparable predictions, although it was not trained with the GT labels of the real dataset.




Figure 5. Qualitative comparison on the REAL275 dataset.

5. Limitations and Future Work

To the best of our knowledge, ours is the first approach that tries to solve unsupervised domain adaption for category-level 6D pose estimation using unlabeled RGB-D data. Therefore, there may exist some issues or future directions to be addressed. For example, our pose estimation depends on object classification, detection, and segmentation to produce appropriately cropped images and sampled depth points, thus sensitive to the performance of the off-the-shelf segmentation pipeline. Also, while the proposed method utilizes a single frame RGB-D image, we may utilize geometric constraints from the video input.

6. Conclusions

We propose UDA-COPE, unsupervised domain adaptation for category-level object pose estimation which ad-

resses the real-world lack-of-label problem using multi-modality (RGB-D). Specifically, we designed a bidirectional point filtering method to filter noisy pseudo labels, and observed depth points, where the filtered depth points improve the robustness of the teacher network, and the filtered pseudo label helps efficient student network training. Both provide for better domain adaptation with real-world pose estimation. Experiments showed that our proposed pipeline and pose-aware point filtering results were comparable to or sometimes better than the performance of fully supervised approaches.

Acknowledgment

This work was conducted by Center for Applied Research in Artificial Intelligence(CARAI) grant funded by Defense Acquisition Program Administration(DAPA) and Agency for Defense Development(ADD) (UD190031RD).

References

- [1] Nikita Araslanov and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15384–15394, 2021. **1**
- [2] Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, et al. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4243–4250, 2018. **1**
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3722–3731, 2017. **1**
- [4] Dengsheng Chen, Jun Li, Zheng Wang, and Kai Xu. Learning canonical shape space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11973–11982, 2020. **1, 2, 4, 5**
- [5] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2773–2782, 2021. **1, 2, 4, 5**
- [6] Xu Chen, Zijian Dong, Jie Song, Andreas Geiger, and Otmar Hilliges. Category level object pose estimation via neural analysis-by-synthesis. In *European Conference on Computer Vision (ECCV)*, pages 139–156, 2020. **2**
- [7] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3075–3084, 2019. **2**
- [8] Guoguang Du, Kai Wang, Shiguo Lian, and Kaiyong Zhao. Vision-based robotic grasping from object localization, object pose estimation to grasp estimation for parallel grippers: a review. *Artificial Intelligence Review*, 54(3):1677–1734, 2021. **1**
- [9] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. **2, 3, 4**
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2961–2969, 2017. **4**
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. **2**
- [12] Yisheng He, Haibin Huang, Haoqiang Fan, Qifeng Chen, and Jian Sun. Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3003–3013, 2021. **1**
- [13] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11632–11641, 2020. **1**
- [14] Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Emilie Wirbel, and Patrick Pérez. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12605–12614, 2020. **1, 2, 3, 6**
- [15] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1521–1529, 2017. **1**
- [16] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–51, 2018. **1**
- [17] Taeyeop Lee, Byeong-Uk Lee, Myungchul Kim, and In So Kweon. Category-level metric scale object shape and pose estimation. *IEEE Robotics and Automation Letters (RA-L)*, 6(4):8575–8582, 2021. **2, 5**
- [18] Xiaolong Li, Yijia Weng, Li Yi, Leonidas Guibas, A Lynn Abbott, Shuran Song, and He Wang. Leveraging se(3) equivariance for self-supervised category-level object pose estimation. *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. **2**
- [19] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6936–6945, 2019. **1, 3**
- [20] Jiehong Lin, Zewei Wei, Zhihao Li, Songcen Xu, Kui Jia, and Yuanqing Li. Dualposenet: Category-level 6d object pose and size estimation using dual pose network with refined learning of pose consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3560–3569, 2021. **1, 2, 5, 7, 8**
- [21] Mohammad Reza Lohmani, Luca Robbiano, Mirco Planamente, Kiru Park, Barbara Caputo, and Markus Vincze. Unsupervised domain adaptation through inter-modal rotation for rgb-d object recognition. *IEEE Robotics and Automation Letters (RA-L)*, 5(4):6631–6638, 2020. **2**
- [22] Fabian Manhardt, Gu Wang, Benjamin Busam, Manuel Nickel, Sven Meier, Luca Minciullo, Xiangyang Ji, and Nassir Navab. Cps++: Improving class-level 6d pose and shape estimation from monocular images with self-supervised learning. *arXiv preprint arXiv:2003.05848*, 2020. **2, 5**
- [23] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 22(12):2633–2651, 2015. **1**

- [24] Eitan Marder-Eppstein. Project tango. In *ACM SIGGRAPH 2016 Real-Time Live!*, pages 25–25, 2016. [1](#)
- [25] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 7668–7677, 2019. [1](#)
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)*, 32:8026–8037, 2019. [4](#)
- [27] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4561–4570, 2019. [1](#)
- [28] Martin Runz, Maud Buffier, and Lourdes Agapito. Mask-fusion: Real-time recognition, tracking and reconstruction of multiple moving objects. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 10–20, 2018. [1](#)
- [29] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018. [1](#)
- [30] Meng Tian, Marcelo H Ang Jr, and Gim Hee Lee. Shape prior deformation for categorical 6d object pose and size estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [31] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning (CoRL)*, 2018. [1](#)
- [32] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence (TPAMI)*, 13(04):376–380, 1991. [2](#), [3](#), [4](#), [7](#)
- [33] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7364–7373, 2019. [2](#)
- [34] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3343–3352, 2019. [1](#)
- [35] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2642–2651, 2019. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [36] Jiaze Wang, Kai Chen, and Qi Dou. Category-level 6d object pose estimation via cascaded relation and recurrent reconstruction networks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. [1](#), [2](#), [5](#)
- [37] Jay M Wong, Vincent Kee, Tiffany Le, Syler Wagner, Gian-Luca Mariottini, Abraham Schneider, Lei Hamilton, Rahul Chipalkatty, Mitchell Hebert, David MS Johnson, et al. Segicp: Integrated deep semantic segmentation and pose estimation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5784–5789, 2017. [1](#)
- [38] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *Robotics: Science and Systems (RSS)*, 2018. [1](#)
- [39] Andy Zeng, Kuan-Ting Yu, Shuran Song, Daniel Suo, Ed Walker, Alberto Rodriguez, and Jianxiong Xiao. Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1386–1383, 2017. [1](#)
- [40] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12414–12424, 2021. [1](#)
- [41] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017. [2](#)
- [42] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5982–5991, 2019. [1](#)