

MSDA: Monocular Self-supervised Domain Adaptation for 6D Object Pose Estimation

Dingding Cai¹, Janne Heikkilä², and Esa Rahtu¹

¹ Tampere University, Finland
{dingding.cai, esa.rahtu}@tuni.fi
² University of Oulu, Finland
janne.heikkila@oulu.fi

Abstract. Acquiring labeled 6D poses from real images is an expensive and time-consuming task. Though massive amounts of synthetic RGB images are easy to obtain, the models trained on them suffer from noticeable performance degradation due to the synthetic-to-real domain gap. To mitigate this degradation, we propose a practical self-supervised domain adaptation approach that takes advantage of real RGB(-D) data without needing real pose labels. We first pre-train the model with synthetic RGB images and then utilize real RGB(-D) images to fine-tune the pre-trained model. The fine-tuning process is self-supervised by the RGB-based pose-aware consistency and the depth-guided object distance pseudo-label, which does not require the time-consuming online differentiable rendering. We build our domain adaptation method based on the recent pose estimator SC6D and evaluate it on the YCB-Video dataset. We experimentally demonstrate that our method achieves comparable performance against its fully-supervised counterpart while outperforming existing state-of-the-art approaches.

Keywords: 6D object pose estimation · self-supervised domain adaptation.

1 Introduction

Estimating the 6D object pose from a monocular RGB image is one of the fundamental tasks in computer vision. This task aims to infer the 3D rotation and 3D translation of a rigid object with respect to the camera coordinate system, which is an important application in robotic manipulation [4,28] and augmented reality [20]. Due to the success of deep learning on large-scale image classification tasks [25], learning-based 6D pose estimation methods [6,16,32] have recently achieved promising performance on benchmark datasets [12].

Most learning-based approaches [6,9,16,26,29] primarily benefit from large amounts of real training images with ground-truth 6D pose labels. Nevertheless, annotating the 6D object pose requires expert knowledge, making the pose label acquisition prohibitively expensive and time-consuming for many applications. As a result, leveraging cost-free auto-annotated synthetic data for training has

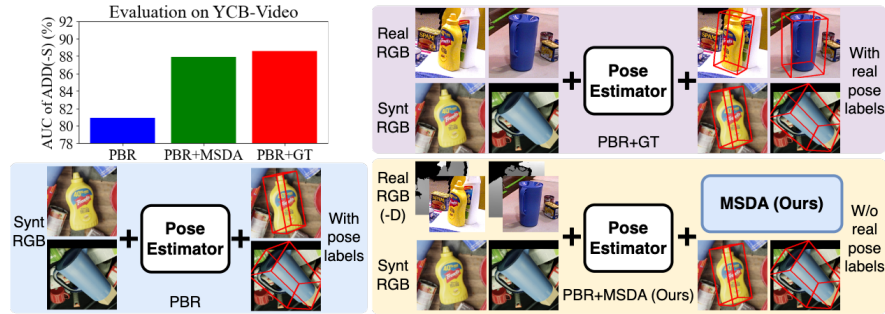


Fig. 1: Comparison of different training strategies and their corresponding performance on YCB-Video. Top left: performance comparison in terms of the AUC of ADD-(S). Bottom left: the object pose estimator is pre-trained using synthetic physically-based rendering (PBR) images. Top right: the pose estimator is fine-tuned using synthetic and real images with ground-truth real pose labels. Bottom right: the pose estimator is fine-tuned using synthetic and real images based on the proposed MSDA approach without needing real pose labels.

become indispensable in many recent 6D object pose estimation works [1,16,27]. Yet, these models trained using synthetic data suffer noticeable performance degradation due to the synthetic-to-real domain gap.

To bridge the domain gap, one seminal work called Self6D [31] was proposed by Wang *et al.* for self-supervised monocular 6D object pose estimation by leveraging synthetic RGB images for pre-training and then real unlabeled RGB-D images for fine-tuning. Later, Wang *et al.* presented an extended version called Self6D++ [30] to better handle object occlusion. Both methods employ the differentiable rendering technique [3] to establish the supervision for fine-tuning. Nevertheless, online differentiable rendering introduces an additional burden of computation during training.

Similarly, this work aims to bridge the synthetic-to-real domain gap for monocular 6D object pose estimation without using real 6D pose labels. To this end, we propose a monocular self-supervised domain adaptation method based on the recent symmetry-agnostic and correspondence-free 6D object pose estimator SC6D [2]. To supervise the domain adaptation, we first introduce a pose-aware data augmentation based on monocular real RGB images for learning the object pose-aware consistency. Next, we present an effective depth-guided pseudo-label generation strategy to supervise the object distance estimation task. Unlike the state-of-the-art method Self6D++, our method requires neither the object symmetry priors nor the time-consuming online differentiable rendering during training.

We experimentally evaluate our method on the YCB-Video dataset [33] and demonstrate that it outperforms not only the top-tier self-supervised approach Self6D++ but also some state-of-the-art fully-supervised methods using real pose labels. In addition, the experimental results show that our method achieves

competitive performance with its fully-supervised counterpart (see Fig. 1), indicating our approach is capable of nearly closing the synthetic-to-real domain gap without using real pose labels.

To summarize, our contributions are as follows. 1). We present a pose-aware consistency learning method based on only monocular RGB images for self-supervised domain adaptation. 2). We propose a practical depth-guided pseudo-label generation approach for supervising the object translation estimation along the z-axis. 3). We conduct experiments on the YCB-Video benchmark and demonstrate that our approach outperforms the state-of-the-art self-supervised method Self6D++ and some current fully-supervised methods.

2 Related Work

In this section, we mainly focus on recent learning-based 6D object pose estimation approaches that are most related to this work.

Correspondence-based methods. This line of work aims to first learn the 2D-3D correspondences, between the 2D image pixel coordinates and the 3D surface coordinates of the object CAD model, from the monocular RGB images and then recover the 6D object poses by solving the Perspective-n-Point (PnP) [17] or PnP-RANSAC [7] algorithms. BB8 [24] is one of the most representative correspondence-based works and aims to detect the 2D projected pixel locations of the 3D object bounding box corners (keypoints) to build sparse 2D-3D correspondences. Different from BB8, PVNet [23] selects the pre-defined 3D keypoints from the surface of the 3D object CAD model and then localizes the 2D pixel coordinates of these 3D keypoints in RGB images based on the pixel-wise voting schema. Apart from building these sparse keypoint-based correspondences, learning the dense pixel-wise correspondences [22,18,34,11,8] has become mainstream in this field. Pix2Pose [22] directly regresses the dense pixel-wise 2D-3D correspondences from RGB images. Instead, DPOD [34] employs a classification-based framework to estimate the 3D coordinates for the 2D object pixels. To better train the object 3D coordinates classification task, ZebraPose [26] presents an efficient 3D surface coordinate representation using hierarchical binary encoding codes and achieves state-of-the-art performance on the YCB-Video dataset [33]. However, most of these methods require prior information about object symmetries to explicitly handle the visual ambiguities. In contrast, EPOS [11] and SurfEmb [8] aim to learn the 2D-3D correspondence distribution over the object surface and can handle visual ambiguity implicitly without needing object symmetry priors.

Regression-based methods. Instead of recovering the object pose from the intermediate 2D-3D correspondences, the regression-based methods estimate the pose parameters directly from the monocular RGB images in an end-to-end fashion. PoseCNN [33] disentangles the 6D pose representation into two components, *i.e.*, the 3D rotation and 3D translation, and separately estimates each component from a monocular RGB image. GDR-Net [32] proposes a geometry-guided

regression network for predicting the 6D pose parameters directly from the estimated 2D-3D correspondences. Based on GDR-Net, SO-Pose [6] further boosts the performance by imposing additional geometry information regarding the object self-occlusion during training. Meanwhile, CosyPose [16] and RePose [15] are capable of predicting the relative pose parameters based on an initial pose estimate to refine the initial pose. Still, these approaches are sensitive to visual ambiguities and require additional effort to deal with object symmetries during training, while SC6D [2] introduces a symmetry-agnostic 6D object pose estimation framework and can effectively deal with visual ambiguities. Therefore, we adapt SC6D as our pose estimator to demonstrate the proposed self-supervised domain adaptation approach.

Self-supervised methods. This type of work aims at learning deep models by exploiting massive unlabeled real-world data with related proxy tasks. To the best of our knowledge, Wang *et al.* [31] presented the first self-supervised monocular 6D object pose estimation work called Self6D. They first train a pose regression network fully supervised with synthetic data, then fine-tune the network using unlabeled real RGB-D images. Self6D adapts an advanced differentiable neural rendering technique DIR-B [3] to establish the supervision during fine-tuning. The follow-up work Self6D++ [30] is developed to account for the occlusion-aware object pose estimation by leveraging both visible and amodal object masks. Meanwhile, Self6D++ employs the much stronger GDR-Net [32] as the pose estimator, thus outperforming Self6D by a large margin.

3 Method

Given an RGB image with the known object, the task is to estimate the 6D object pose $[R|t]$, *i.e.*, the 3D rotation $R \in SO(3)$ and the 3D translation $t \in \mathbb{R}^3$, with respect to the camera coordinate system. We propose a monocular self-supervised domain adaptation (MSDA) approach for the pose estimation task lacking real pose labels. As illustrated in Fig. 1, we first pre-train the pose estimator fully supervised with synthetic RGB images generated with the physically-based rendering (PBR) technique [5] and then fine-tune the pose estimator using a mixture of synthetic and real-world data without needing real pose labels. Our method can be easily applied as a plugin on the two-stage based object pose estimators (*i.e.*, first performing the object detection and then estimating the 6D object pose parameters), such as GDR-Net [32] and SC6D [2]. In this work, we implement our MSDA approach based on the recent pose estimator SC6D due to its top-performing performance.

3.1 Framework

Architecture. We build our method upon the recent symmetry-agnostic pose estimator SC6D [2] that directly predicts the object pose parameters from an object-centric RGB image. We slightly modify the original SC6D to output the

same size segmentation mask as the input and denote it as SC6D⁺⁺, see Fig. 2. SC6D⁺⁺ consists of two major components, *i.e.*, a pose prediction network Θ and a 3D rotation decoder Ω , and estimates the object 6D pose from the object-centric RGB image, as shown in Fig.2. The pose prediction network Θ takes a known object RGB image $I \in \mathbb{R}^{S \times S \times 3}$ as input and outputs the object distance $\delta_z \in \mathbb{R}$, 2D projection center offset $\delta_{xy} \in \mathbb{R}^2$, segmentation mask $M \in \mathbb{R}^{S \times S}$, and RGB embedding vector $I_{emb} \in \mathbb{R}^C$. Meanwhile, the 3D rotation decoder outputs the 3D object rotation matrix $R \in \mathbb{R}^{3 \times 3}$ based on the RGB embedding vector I_{emb} .

Scale-invariant 3D translation estimation. The outputs δ_{xy} and δ_z of the pose prediction network Θ are grounded on the object-centric RGB image I , and this image is obtained using the object bounding box centered at (x_{bbox}, y_{bbox}) with the size of (w_{bbox}, h_{bbox}) . Thus, the 3D object translation is recovered by $t = [t_x, t_y, t_z]^T = r \cdot \delta_z \cdot K^{-1}[o_x, o_y, 1]^T$, where $r = S / \max(w_{bbox}, h_{bbox})$ is the rescaling factor, $[o_x, o_y]^T = S \cdot (\delta_{xy} + 0.5)$ is the projected 2D object center location, and K is the corresponding (virtual) camera intrinsic regarding the rescaled object-centric RGB image I with the fixed scale S .

Symmetry-agnostic 3D rotation estimation. The 3D rotation decoder Ω aims to learn the 3D object rotation distribution based on the visual representation to implicitly handle the symmetry ambiguities. To this end, an SO(3) encoder Ψ is utilized to learn a latent 3D rotation representation $R_{emb}^i \in \mathbb{R}^C$ for each 3D rotation R^i uniformly sampled from the SO(3) space, *i.e.* $R_{emb}^i = \Psi(R^i)$. The 3D object rotation distribution is approximated by

$$p(R^i|I) \approx \frac{\exp(I_{emb} \cdot R_{emb}^i)}{\sum_{j=1}^{N_r} \exp(I_{emb} \cdot R_{emb}^j)}, \forall i = 1, 2, \dots, N_r, \quad (1)$$

where N_r is the number of the 3D rotation samples, and $I_{emb} \in \mathbb{R}^C$ denotes the representation vector learned from the object-centric RGB image I . By doing so, the object symmetry ambiguities can be implicitly handled, for example, by mapping all symmetric 3D rotations to similar latent representations correlated with the object observation. The 3D rotation R with the highest probability $p(R|I)$ is thus selected as the prediction.

Supervised learning with synthetic data. The pose estimator is pre-trained using synthetic data in a fully supervised manner, and the training loss is summarized as

$$\mathcal{L}^{syn} = \lambda_{xy}^{syn} \mathcal{L}_{xy}^{syn} + \lambda_z^{syn} \mathcal{L}_z^{syn} + \lambda_R^{syn} \mathcal{L}_R^{syn} + \lambda_M^{syn} \mathcal{L}_M^{syn}, \quad (2)$$

where $\mathcal{L}_{\{xy,z,R,M\}}^{syn}$ represent the training objectives for supervising the estimation tasks of the 2D object center offset δ_{xy} , the object distance δ_z , the 3D object rotation R , and the object segmentation mask M , and $\lambda_{\{xy,z,R,M\}}^{syn}$ denote the balancing parameters for the corresponding loss terms. We kindly refer the reader to [2] for more implementation details.

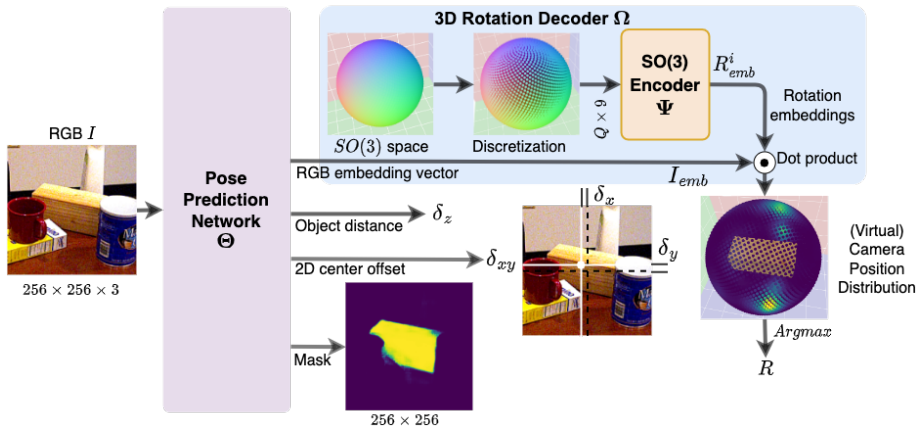


Fig. 2: Overview of SC6D⁺⁺. The pose prediction network Θ takes a known object RGB image I as input that predicts the 2D object center offset δ_{xy} , the object distance δ_z , and the image embedding vector I_{emb} . Subsequently, the 3D rotation decoder takes as input the embedding vector I_{emb} and outputs the 3D object orientation R with respect to the object itself. To visualize the estimated 3D rotation probability distribution (see the right part), we highlight the possible positions of the (virtual) camera on a sphere centered upon the object. The brighter color indicates higher probabilities.

Initial 6D pose estimate. Like Self6D⁺⁺ [30], our method requires an initial 6D pose estimate from the real RGB image to establish the supervision. Thus, we leverage the pre-trained pose estimator to provide the initial 6D pose $[R|t]$ utilized for our self-supervised domain adaptation.

3.2 Self-supervised Pose-Aware Consistency Learning

We first introduce our self-supervised pose-aware consistency learning based on monocular RGB images. Given a real RGB image $\mathcal{I}_{real} \in \mathbb{R}^{H \times W \times 3}$ with an object bounding box $B_{real} = (x_{bbox}, y_{bbox}, w_{bbox}, h_{bbox})$, we generate a pair of object-centric RGB samples with the known relative pose and enforce the pose estimator to learn the pose-aware consistency (see Fig.3).

Specifically, we first obtain an anchor sample I^{anc} using an enlarged square bounding box B^{anc} centered at $c_{bbox}^{anc} = (x_{bbox}, y_{bbox})$ with the scale of $s_{bbox}^{anc} = f_{anc} \cdot \max(w_{bbox}, h_{bbox})$, where f_{anc} is a random scaling factor. Then, we transform the anchor bounding box B^{anc} using a random rescaling factor Δs , a random offset ΔP_{xy} , and a random in-plane rotation $\Delta R_z \in \text{SO}(2)$ to obtain an augmented variant I^{aug} for the anchor sample dynamically. The pose relation-

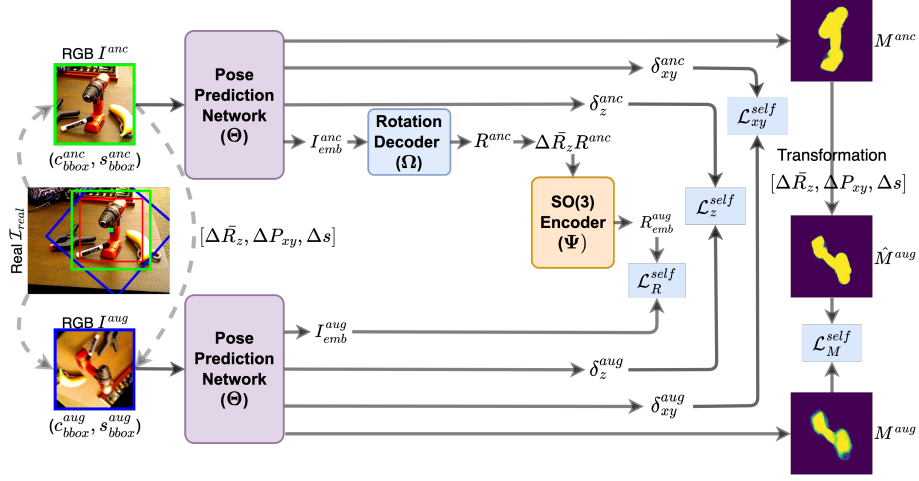


Fig. 3: Illustration of the pose-aware consistency learning for self-supervised domain adaptation. Given a real RGB image \mathcal{I}_{real} and an object bounding box (red box), we first obtain an anchor RGB sample I^{anc} using an expanded box (green box) and an augmented RGB sample I^{aug} using a transformed box (blue box) with a known relative pose $\Delta = [\Delta \bar{R}_z, \Delta P_{xy}, \Delta s]$. Then the pose prediction network Θ separately predicts the outputs of the paired samples (I^{anc}, I^{aug}), *i.e.*, the object center projection offsets ($\delta_{xy}^{anc}, \delta_{xy}^{aug}$), the object distances ($\delta_z^{anc}, \delta_z^{aug}$), the RGB embedding vectors ($I_{emb}^{anc}, I_{emb}^{aug}$), and the segmentation masks (M^{anc}, M^{aug}). The supervision signals are established by enforcing the pose-aware consistencies between the anchor sample and its augmented one. The network weights in all modules are shared during training.

ship between the training pair (I^{aug}, I^{anc}) can be formulated as

$$\begin{cases} R^{aug} &= \Delta \bar{R}_z R^{anc} \\ \delta_z^{aug} &= \Delta s \cdot \delta_z^{anc} \\ \delta_{xy}^{aug} &= \Delta R_z (\delta_{xy}^{anc} - \Delta P_{xy} / s_{bbox}^{anc}) / \Delta s \end{cases}, \quad (3)$$

where $[\delta_{xy}^{anc}, \delta_z^{anc}, R^{anc}]$ and $[\delta_{xy}^{aug}, \delta_z^{aug}, R^{aug}]$ are the predictions of the model for the anchor sample I^{anc} and its augmented one I^{aug} , respectively. $\Delta \bar{R}_z \in \mathbb{R}^{3 \times 3}$ denotes the 3×3 matrix representation of ΔR_z .

By this means, we construct online pairwise training samples (I^{aug}, I^{anc}) with the known relative transformation $\Delta = [\Delta s, \Delta P_{xy}, \Delta R_z]$ based on monocular real RGB images without using pose labels. Subsequently, we leverage the L1 loss to measure the translation consistency, *i.e.*,

$$\begin{cases} \mathcal{L}_z^{self} &= \|\delta_z^{aug} - \Delta s \cdot \delta_z^{anc}\|_1 \\ \mathcal{L}_{xy}^{self} &= \|\delta_{xy}^{aug} - \Delta R_z (\delta_{xy}^{anc} - \Delta P_{xy} / s_{bbox}^{anc}) / \Delta s\|_1 \end{cases}. \quad (4)$$

Regarding the rotation consistency, we enforce the network to minimize the negative log-likelihood $p(R^{aug}|I^{aug})$,

$$\begin{aligned}\mathcal{L}_R^{self} &= -\log(p(R^{aug}|I^{aug})) \\ &= -\log \frac{\exp(I_{emb}^{aug} \cdot R_{emb}^{aug}/\tau)}{\sum_i^{N_r} \exp(I_{emb}^{aug} \cdot R_{emb}^i/\tau)},\end{aligned}\quad (5)$$

where $I_{emb}^{aug} \in \mathbb{R}^C$ is the image embedding vector of the augmented sample I^{aug} , $R_{emb}^{aug} = \Psi(R^{aug})$ is the latent representation vector of the derived 3D rotation R^{aug} , and $\tau = 0.1$ is the temperature parameter.

In addition, we also leverage the object segmentation consistency loss \mathcal{L}_M^{self} to supervise the domain adaptation process. To do this, we first transform the 2D pixel coordinate grid G^{anc} of the anchor RGB sample I^{anc} based on the transformation Δ and then sample the segmentation mask \hat{M}^{aug} from the predicted anchor mask M^{anc} according to the transformed G^{anc} (see the right side in Fig.3). Therefore, the self-supervised pose-aware consistency loss is summarized as

$$\mathcal{L}^{self} = \lambda_{xy}^{self} \mathcal{L}_{xy}^{self} + \lambda_z^{self} \mathcal{L}_z^{self} + \lambda_R^{self} \mathcal{L}_R^{self} + \lambda_M^{self} \mathcal{L}_M^{self} \quad (6)$$

where $\lambda_{\{xy,z,R,M\}}^{self}$ denote the weight hyper-parameters for the corresponding loss terms.

3.3 Depth-guided Object Distance Learning

When depth images \mathcal{D}_{real} are available, we propose a depth-guided pseudo-label generation method and utilize the generated *pseudo-label* \bar{t}_z to supervise the object translation estimation along the z-axis (t_z). However, it is challenging to estimate the object distance (*i.e.*, from the camera origin to the object centroid) from the depth image directly because these depth values represent the measurements from the camera plane to the surface of the 3D object. Moreover, this task becomes more challenging when the object is occluded and only partially visible in the depth image.

To tackle these issues, we propose estimating a depth offset δ_d between the two closest points, one from the observed real depth image and the other from the 3D object point cloud transformed using the initial pose $[R|t]$, to compensate for the initial z-axis translation estimate t_z . To this end, we first obtain the object-centric real depth image $D \in \mathbb{R}^{S \times S}$ using the anchor bounding box B^{anc} , then the 3D object points $P_s \in \mathbb{R}^{N_s \times 3}$ are uniformly sampled from the surface of the 3D object model, where N_s is the number of the 3D points. Next, we transform the sampled 3D object points P_s using the estimated pose $[R|t]$ and the intrinsic K , *i.e.*,

$$D_s^i [U_s^i, V_s^i, 1]^T = K(RP_s^i + t), \quad \forall i = 1, \dots, N_s \quad (7)$$

where D_s^i and (U_s^i, V_s^i) denote the depth value and the projected 2D pixel coordinate of the 3D object point P_s^i after the transformation, respectively.




Fig. 4: Illustration of the depth-guided object z-axis translation pseudo-label generation. Given the object-centric RGB and depth images, we first utilize the pre-trained pose estimator to predict the object mask M and an initial pose $[R|t]$ from the RGB image. Then we transform the 3D object point cloud using the estimated initial pose and the corresponding camera intrinsic to obtain the projected depth image D_s . A depth set \bar{D}_s is sampled from the projected depth D_s based on the predicted mask M , and another depth set \bar{D}_r is sampled from the input depth image D based on both M and D_s . On the right side, the depth offset δ_d between two point clouds is considered as the distance between two virtual parallel planes perpendicular to the camera’s optical axis. One plane intersects with the closest point in the partial point cloud (blue) lifted from the observed depth image \bar{D}_r . The other plane intersects with the closest point in the complete 3D object point cloud (red) transformed by the initial pose estimate $[R|t]$.

Based on the predicted object mask M and the projected coordinate (U_s^i, V_s^i) , we select the corresponding depth values from the set $\{D_s^i\}_{i=1}^{N_s}$ and the object-centric real depth image D . This process can be formulated as

$$\begin{aligned} \bar{D}_s^i &= \begin{cases} D_s^i, & \text{if } \{U_s^i, V_s^i\} \in M \text{ and } M_{U_s^i, V_s^i} \geq \rho, \\ \epsilon, & \text{otherwise,} \end{cases} \\ \bar{D}_r^i &= \begin{cases} D_{U_s^i, V_s^i}, & \text{if } \{U_s^i, V_s^i\} \in M \text{ and } M_{U_s^i, V_s^i} \geq \rho, \\ \epsilon, & \text{otherwise,} \end{cases} \end{aligned} \quad (8)$$

where \bar{D}_s^i and \bar{D}_r^i are the sampled depth values from the transformed 3D object point cloud and the observed depth image D (see the left side in Fig. 4), respectively. $\rho = 0.9$ is the confidence threshold, and $\epsilon \gg \max(D)$ denotes the invalid (penalty) depth value for the background.

In this way, we obtain two depth sets, *i.e.*, $\bar{D}_s = \{\bar{D}_s^i\}_{i=1}^{N_s}$ and $\bar{D}_r = \{\bar{D}_r^i\}_{i=1}^{N_s}$, the former from the transformed 3D object point cloud P_s and the latter from

the observed depth image D . The depth offset δ_d between these two sets can be computed to generate the *pseudo-label* \bar{t}_z for the object z-axis translation, *i.e.*, $\bar{t}_z = \delta_d + t_z$. However, it is non-trivial to compute this offset δ_d because \bar{D}_s is derived from the entire 3D object point cloud while \bar{D}_r is sampled from its partial observation. To approximate δ_d , we compute the depth offset between the closest point observed in \bar{D}_r and the one from \bar{D}_s , *i.e.*, $\delta_d = \min(\bar{D}_r) - \min(\bar{D}_s)$. In other words, δ_d can be regarded as the distance between two planes perpendicular to the camera optical axis, where one plane intersects with the closest point in the observed object depth image, and the other intersects with the closest point in the transformed object point cloud, see the right in Fig. 4.

However, outliers could be included in the real depth set \bar{D}_r and result in significant errors due to the inaccurate mask prediction M . Instead of simply selecting the closest real depth $\bar{D}_r^o = \min(\bar{D}_r)$, we propose an adaptive selection strategy to compute \bar{D}_r^o . We first rank the depth set \bar{D}_r in ascending order, followed by a moving average operation, to obtain a smooth sorted set \hat{D}_r and then calculate the adjacent depth difference \hat{G}_r of the sorted set. Finally, the real depth value $\bar{D}_r^o = \hat{D}_r^k$ is selected if the first $k^{th} \in [1, N_s - 1]$ difference value in \hat{G}_r is smaller than the predefined threshold, *i.e.* $\hat{G}_r^k \leq \gamma$ where $\gamma = 0.001$ in our experiments, otherwise $\bar{D}_r^o = \min(\bar{D}_r)$. We utilize the generated *pseudo-label* \bar{t}_z and a truncated L1 loss to supervise the training process,

$$\mathcal{L}_{tz}^{pseudo} = \min(\|\bar{t}_z - t_z\|_1, \xi) \quad (9)$$

where $\xi = 0.1$ is an upper-bound value for the distance offset to avoid the unstable loss.

3.4 Joint Fine-tuning for Domain Adaptation

We employ both synthetic and real data to jointly fine-tune the pre-trained pose estimator. To summarize, the overall training objective is written as

$$\mathcal{L} = \lambda^{syn} \mathcal{L}^{syn} + \lambda^{self} \mathcal{L}^{self} + \lambda_{tz}^{pseudo} \mathcal{L}_{tz}^{pseudo}, \quad (10)$$

where λ^{syn} , λ^{self} , and λ_{tz}^{pseudo} represent the weight hyper-parameters of the corresponding loss terms.

4 Experiment

Dataset. The experiments are conducted on the benchmark YCB-Video [33] dataset. YCB-Video [33] contains 21 objects with textured 3D CAD models and 92 annotated real RGB-D videos collected from multiple noisy environments. More than 113,000 images from 80 videos are used as the training data, and 2949 keyframes are selected from the rest 12 videos as the testing set. The physically-based rendering (PBR) dataset [12] contains 50,000 synthetic images and is employed for training in this work.

Table 1: Ablation studies using the YCB-Video benchmark in terms of Average Recall (%) of ADD(-S) metric, and AUC of ADD-S/ADD(-S) metrics. "PBR" indicates the synthetic PBR training images. "LB" denotes the lower bound result only using PBR for training, and "UB" represents the upper bound result obtained using both PBR and real data for training with ground-truth pose labels.

	Training Data	Method	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)
A1	PBR	SC6D ⁺⁺ (LB)	56.1	87.5	80.9
A2	PBR + Real (RGB)	SC6D ⁺⁺ (UB)	79.8	92.2	88.6
B1	PBR + Real (RGB)	SC6D _{MSDA} ⁺⁺ w/o $\mathcal{L}_{tz}^{pseudo}$	59.4	88.4	82.2
B2	PBR + Real (RGB-D)	SC6D _{MSDA} ⁺⁺ w/o \mathcal{L}^{self}	77.6	91.6	87.2
B3	PBR + Real (RGB-D)	SC6D _{MSDA} ⁺⁺ (Ours)	80.7	92.1	87.9

Evaluation Metrics. We report our results in terms of the standard metrics of ADD(-S) [10] with the error tolerance threshold less than the 10% diameter of the object and the AUC (area under the curve) of ADD(-S)/ADD-S [33] with varying tolerance thresholds up to 10cm to fairly compare with other approaches. We kindly refer the reader to [10,33] for the definition details.

Implementation Details. We utilize the AdamW solver [19] with the cosine annealing learning rate to train the model on 16 Nvidia GPUs. The pose estimator is first pre-trained using synthetic data for 75 epochs with the learning rate starting from 5×10^{-4} to 1×10^{-5} and then fine-tuned using a mixture of synthetic and real data for additional five epochs with the learning rate from 1×10^{-4} to 1×10^{-5} . We set the weight hyper-parameters $\lambda_{xy}^{syn} = 10.0$, $\lambda_z^{syn} = 1.0$, $\lambda_R^{syn} = 1.0$, $\lambda_M^{syn} = 10.0$, $\lambda_{xy}^{self} = 10.0$, $\lambda_z^{self} = 10.0$, $\lambda_R^{self} = 0.1$, $\lambda_M^{self} = 10.0$, $\lambda^{syn} = 1.0$, $\lambda^{self} = 0.1$, $\lambda_{tz}^{pseudo} = 10.0$ to maintain comparable magnitudes among all loss terms. Besides, we follow the 3D rotation sampling strategy in [21,1] and uniformly sample $N_r = 5000$ rotations for training and $N_r = 480,000$ (4000 viewpoints on a sphere with 120 in-plane rotations for each viewpoint) for testing. We apply strong image augmentations on synthetic RGB images following [2] and dynamically crop the region of interest using the in-plane rotated object bounding box during training. For evaluation, we leverage the predicted object bounding boxes provided by Self6D++ [30] for a fair comparison, where the object detector is trained using purely synthetic PBR images.

4.1 Ablation Study

We first conduct ablation studies using the proposed self-supervised domain adaptation approach based on monocular RGB(-D) images. The experimental results are reported in Table 1 in terms of the average recall of ADD(-S) and the AUC of ADD-S/ADD(-S).

Table 2: Comparison with state-of-the-art methods in terms of Average Recall (%) of ADD(-S), and AUC of ADD-S/ADD(-S). "Synt + Real GT" denotes using both synthetic and real data with ground-truth pose labels for fully-supervised training. "Synt + Real Self" denotes using synthetic and real data for self-supervised training without needing real pose labels. "UB" indicates the upper bound results obtained using real pose labels. The best results are highlighted in **bold**, and the second-best results are underlined. "-" indicates unavailable results.

Training	Method	ADD(-S)	AUC of ADD-S	AUC of ADD(-S)
Supervision: Synt GT + Real GT	PoseCNN [33]	21.3	75.9	61.3
	PVNet [23]	-	-	73.4
	Seg-Driven [14]	39.0	-	-
	SingleStage [13]	53.9	-	-
	CosyPose [16]	-	89.8	84.5
	SO-Pose [6]	56.8	90.9	83.9
	GDR-Net [32]	60.1	91.6	84.4
	RePose [15]	62.1	88.5	82.0
	ZebraPose [26]	<u>80.5</u>	90.1	85.3
	Self6D++(UB)[30]	-	90.7	82.6
	SC6D [2]	77.1	90.8	86.4
	SC6D ⁺⁺ (UB)	79.8	92.2	88.6
Supervision: Synt GT + Real Self	Self6D++ [30]	-	91.1	80.0
	SC6D ⁺⁺ _{MSDA} (Ours)	80.7	<u>92.1</u>	<u>87.9</u>

Compared to the baseline (A1) fully supervised with synthetic data only (A1), using all the proposed self-supervised training losses (B3) can significantly boost the performance from 56.1% to 80.7% in terms of ADD(-S) and from 80.9% to 87.9% in terms of AUC of ADD(-S). Specifically, SC6D⁺⁺_{MSDA} *w/o* $\mathcal{L}_{tz}^{pseudo}$ improves the ADD(-S) recall by over 3% using only the pose-aware consistency loss based on monocular real RGB images (B1). In addition, SC6D⁺⁺_{MSDA} *w/o* \mathcal{L}^{self} can achieve 77.6% ADD(-S) recall using the loss only based on the object depth-guided z-axis translation *pseudo-labels* (B2).

Overall, our SC6D⁺⁺_{MSDA} achieves optimal performance using both the object RGB-based pose-aware consistency and the object depth-guided *pseudo-label* distance losses (B3). Noteworthy, SC6D⁺⁺_{MSDA} consistently improves over its baseline SC6D⁺⁺(LB) and approaches towards its fully-supervised counterpart SC6D⁺⁺(UB), even slightly outperforming its fully-supervised counterpart regarding the ADD(-S) metric (80.7% *vs.* 79.8%), which indicates SC6D⁺⁺_{MSDA} can nearly bridge the synthetic-to-real domain gap.

4.2 Comparison with State-of-the-art

In this part, we compare SC6D⁺⁺_{MSDA} against current state-of-the-art 6D object pose estimation approaches on the YCB-Video dataset. The comparison results

are presented in Table 2 in terms of the standard metrics, ADD(-S) and AUC of ADD-S/ADD(-S). In general, $\text{SC6D}_{\text{MSDA}}^{++}$ achieves state-of-the-art performance regarding the ADD(-S) metric and outperforms the top-tier self-supervised approach Self6D++ [30] in terms of the AUC of ADD-S/ADD(-S).

On the one hand, our self-supervised $\text{SC6D}_{\text{MSDA}}^{++}$ achieves superior performance over the state-of-the-art fully-supervised pose estimators ZebraPose [26] and SC6D [2]. Concretely, ZebraPose [26] is a top-performing correspondence-based pose estimator achieving 80.5% ADD(-S) recall using real pose labels for training, in comparison, $\text{SC6D}_{\text{MSDA}}^{++}$ can reach 80.7% without needing real pose labels. Meanwhile, compared with the original symmetry-agnostic SC6D [2], $\text{SC6D}_{\text{MSDA}}^{++}$ outperforms it by 3.6% and 1.6% regarding the metrics of ADD(-S) and AUC of ADD(-S) despite the lack of real pose labels, correspondingly. On the other hand, $\text{SC6D}_{\text{MSDA}}^{++}$ also outperforms the current state-of-the-art self-supervised pose estimator, Self6D++ [30]. In particular, without using real pose labels, Self6D++ achieves 91.1%/80.0% AUC of ADD-S/ADD(-S) behind 92.1%/87.9% obtained by $\text{SC6D}_{\text{MSDA}}^{++}$.

5 Conclusion

This work introduces a monocular self-supervised domain adaptation approach for 6D object pose estimation. By leveraging real-world RGB(-D) images, our method is capable of bridging the synthetic-to-real domain gap and achieving superior performance on the YCB-Video benchmark without using real pose labels. We accomplish this by learning the object pose-aware consistency based on real-world monocular RGB images and compensating for the inaccurate object translation estimation along the z-axis using the observed depth images. Our approach does not rely on complex differentiable rendering techniques to establish supervision, simplifying the self-supervised domain adaptation process for the 6D object pose estimation task.

References

1. Cai, D., Heikkilä, J., Rahtu, E.: Ove6d: Object viewpoint encoding for depth-based 6d object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6803–6813 (2022)
2. Cai, D., Heikkilä, J., Rahtu, E.: Sc6d: Symmetry-agnostic and correspondence-free 6d object pose estimation. arXiv preprint arXiv:2208.02129 (2022)
3. Chen, W., Ling, H., Gao, J., Smith, E., Lehtinen, J., Jacobson, A., Fidler, S.: Learning to predict 3d objects with an interpolation-based differentiable renderer. Advances in Neural Information Processing Systems **32** (2019)
4. Collet, A., Martinez, M., Srinivasa, S.S.: The moped framework: Object recognition and pose estimation for manipulation. The international journal of robotics research **30**(10), 1284–1306 (2011)
5. Denninger, M., Sundermeyer, M., Winkelbauer, D., Olefir, D., Hodan, T., Zidan, Y., Elbadrawy, M., Knauer, M., Katam, H., Lodhi, A.: Blenderproc: Reducing the reality gap with photorealistic rendering. In: International Conference on Robotics: Science and Systems, RSS 2020 (2020)

6. Di, Y., Manhardt, F., Wang, G., Ji, X., Navab, N., Tombari, F.: So-pose: Exploiting self-occlusion for direct 6d pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12396–12405 (2021)
7. Fischler, M.A., Bolles, R.C.: 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)
8. Haugaard, R.L., Buch, A.G.: Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6749–6758 (2022)
9. He, Y., Huang, H., Fan, H., Chen, Q., Sun, J.: Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3003–3013 (2021)
10. Hinterstoisser, S., Lepetit, V., Ilic, S., Holzer, S., Bradski, G., Konolige, K., Navab, N.: Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In: Asian conference on computer vision. pp. 548–562. Springer (2012)
11. Hodan, T., Barath, D., Matas, J.: Epos: Estimating 6d pose of objects with symmetries. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11703–11712 (2020)
12. Hodaň, T., Sundermeyer, M., Drost, B., Labbé, Y., Brachmann, E., Michel, F., Rother, C., Matas, J.: BOP challenge 2020 on 6D object localization. European Conference on Computer Vision Workshops (ECCVW) (2020)
13. Hu, Y., Fua, P., Wang, W., Salzmann, M.: Single-stage 6d object pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2930–2939 (2020)
14. Hu, Y., Hugonot, J., Fua, P., Salzmann, M.: Segmentation-driven 6d object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3385–3394 (2019)
15. Iwase, S., Liu, X., Khirodkar, R., Yokota, R., Kitani, K.M.: Repose: Fast 6d object pose refinement via deep texture rendering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3303–3312 (2021)
16. Labbé, Y., Carpentier, J., Aubry, M., Sivic, J.: Cosypose: Consistent multi-view multi-object 6d pose estimation. In: European Conference on Computer Vision. pp. 574–591. Springer (2020)
17. Lepetit, V., Moreno-Noguer, F., Fua, P.: Epnp: An accurate o (n) solution to the pnp problem. *International journal of computer vision* **81**(2), 155 (2009)
18. Li, Z., Wang, G., Ji, X.: Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7678–7687 (2019)
19. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
20. Marchand, E., Uchiyama, H., Spindler, F.: Pose estimation for augmented reality: a hands-on survey. *IEEE transactions on visualization and computer graphics* **22**(12), 2633–2651 (2015)
21. Park, K., Mousavian, A., Xiang, Y., Fox, D.: Latentfusion: End-to-end differentiable reconstruction and rendering for unseen object pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10710–10719 (2020)

22. Park, K., Patten, T., Vincze, M.: Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In: The IEEE International Conference on Computer Vision (ICCV) (Oct 2019)
23. Peng, S., Zhou, X., Liu, Y., Lin, H., Huang, Q., Bao, H.: Pvnnet: pixel-wise voting network for 6dof object pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
24. Rad, M., Lepetit, V.: Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3828–3836 (2017)
25. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
26. Su, Y., Saleh, M., Fetzer, T., Rambach, J., Navab, N., Busam, B., Stricker, D., Tombari, F.: ZebraPose: Coarse to fine surface encoding for 6dof object pose estimation. *arXiv preprint arXiv:2203.09418* (2022)
27. Sundermeyer, M., Durner, M., Puang, E.Y., Marton, Z.C., Vaskevicius, N., Arras, K.O., Triebel, R.: Multi-path learning for object pose estimation across domains. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13916–13925 (2020)
28. Tremblay, J., To, T., Sundaralingam, B., Xiang, Y., Fox, D., Birchfield, S.: Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790* (2018)
29. Wang, C., Xu, D., Zhu, Y., Martín-Martín, R., Lu, C., Fei-Fei, L., Savarese, S.: Densefusion: 6d object pose estimation by iterative dense fusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3343–3352 (2019)
30. Wang, G., Manhardt, F., Liu, X., Ji, X., Tombari, F.: Occlusion-aware self-supervised monocular 6d object pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
31. Wang, G., Manhardt, F., Shao, J., Ji, X., Navab, N., Tombari, F.: Self6d: Self-supervised monocular 6d object pose estimation. In: European Conference on Computer Vision. pp. 108–125. Springer (2020)
32. Wang, G., Manhardt, F., Tombari, F., Ji, X.: GDR-Net: Geometry-guided direct regression network for monocular 6d object pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16611–16621 (June 2021)
33. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In: Proceedings of Robotics: Science and Systems (RSS) (2018)
34. Zakharov, S., Shugurov, I., Ilic, S.: Dpod: 6d pose object detector and refiner. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1941–1950 (2019)