Category-Level 6D Object Pose and Size Estimation using Self-Supervised Deep Prior Deformation Networks

Jiehong Lin^{1,2}, Zewei Wei¹, Changxing Ding¹, and Kui Jia^{1,3}*

¹ South China University of Technology ² DexForce Co. Ltd. ³ Peng Cheng Laboratory {lin.jiehong,eeweizewei}@mail.scut.edu.cn; {chxding,kuijia}@scut.edu.cn

Abstract. It is difficult to precisely annotate object instances and their semantics in 3D space, and as such, synthetic data are extensively used for these tasks, e.g., category-level 6D object pose and size estimation. However, the easy annotations in synthetic domains bring the downside effect of synthetic-to-real (Sim2Real) domain gap. In this work, we aim to address this issue in the task setting of Sim2Real, unsupervised domain adaptation for category-level 6D object pose and size estimation. We propose a method that is built upon a novel Deep Prior Deformation Network, shortened as DPDN. DPDN learns to deform features of categorical shape priors to match those of object observations, and is thus able to establish deep correspondence in the feature space for direct regression of object poses and sizes. To reduce the Sim2Real domain gap, we formulate a novel self-supervised objective upon DPDN via consistency learning; more specifically, we apply two rigid transformations to each object observation in parallel, and feed them into DPDN respectively to yield dual sets of predictions; on top of the parallel learning, an inter-consistency term is employed to keep cross consistency between dual predictions for improving the sensitivity of DPDN to pose changes, while individual intra-consistency ones are used to enforce selfadaptation within each learning itself. We train DPDN on both training sets of the synthetic CAMERA25 and real-world REAL275 datasets; our results outperform the existing methods on REAL275 test set under both the unsupervised and supervised settings. Ablation studies also verify the efficacy of our designs. Our code is released publicly at https://github.com/JiehongLin/Self-DPDN.

Keywords: 6D Pose Estimation, Self-Supervised Learning

1 Introduction

The task of category-level 6D object pose and size estimation, formally introduced in [24], is to estimate the rotations, translations, and sizes of unseen object instances of certain categories in cluttered RGB-D scenes. It plays a crucial

^{*} Corresponding author



Fig. 1. An illustration of our proposed self-supervised Deep Prior Deformation Network (DPDN). DPDN deforms categorical shape priors in the feature space to pair with object observations, and establishes deep correspondence for direct estimates of object poses and sizes; upon DPDN, a novel self-supervised objective is designed to reduce synthetic-to-real domain gap via consistency learning. Specifically, we apply two rigid transformations to the point set \mathcal{P}_o of an object observation, and feed them into DPDN in parallel to make dual sets of predictions; on top of the parallel learning, an interconsistency ones within each learning to form the self-supervision. For simplicity, we omit the image input of object observation. Notations are explained in Sec. 3.

role in many real-world applications, such as robotic grasping[27,17], augmented reality[2], and autonomous driving [12,26,5,6].

For this task, existing methods can be roughly categorized into two groups, i.e., those based on direct regression and those based on dense correspondence learning. Methods of the former group [4,15,14] are conceptually simple, but struggle in learning pose-sensitive features such that direct predictions can be made in the full SE(3) space; dense correspondence learning [24,20,25,3,11] makes the task easier by first regressing point-wise coordinates in the canonical space to align with points of observations, and then obtaining object poses and sizes via solving of Umeyama algorithm [21]. Recent works [20,25,3,11] of the second group exploit strong categorical priors (*e.g.*, mean shapes of object categories) for improving the qualities of canonical point sets, and constantly achieve impressive results; however, their surrogate objectives for the learning of canonical coordinates are one step away from the true ones for estimating object poses and sizes, making their learning suboptimal to the end task.

The considered learning task is further challenged by the lack of real-world RGB-D data with careful object pose and size annotations in 3D space. As such, synthetic data are usually simulated and rendered whose annotations can be freely obtained on the fly [24,7]. However, the easy annotations in synthetic domains bring a downside effect of synthetic-to-real (Sim2Real) domain gap; learning with synthetic data with no consideration of Sim2Real domain adaptation would inevitably result in poor generalization in the real-world domain. This naturally falls in the realm of Sim2Real, unsupervised domain adaptation (UDA) [1,16,19,30,28,29,23,11].

In this work, we consider the task setting of Sim2Real UDA for category-level 6D object pose and size estimation. We propose a new method of self-supervised Deep Prior Deformation Network; Fig. 1 gives an illustration. Following dense correspondence learning, we first present a novel Deep Prior Deformation Network, shortened as **DPDN**, which implements a deep version of shape prior deformation in the feature space, and is thus able to establish deep correspondence for direct regression of poses and sizes with high precision. For a cluttered RGB-D scene, we employ a 2D instance segmentation network (e.q, Mask RCNN [8])to segment the objects of interest out, and feed them into our proposed DPDN for pose and size estimation. As shown in Fig. 2, the architecture of DPDN consists of three main modules, including a Triplet Feature Extractor, a Deep Prior Deformer, and a Pose and Size Estimator. For an object observation, the Triplet Feature Extractor learns point-wise features from its image crop, point set, and categorical shape prior, respectively; then Deep Prior Deformer deforms the prior in feature space by learning a feature deformation field and a correspondence matrix, and thus builds deep correspondence from the observation to its canonical version; finally, Pose and Size Estimator is used to make reliable predictions directly from those built deep correspondence.

On top of DPDN, we formulate a self-supervised objective that combines an inter-consistency term with two intra-consistency ones for UDA. More specifically, as shown in Fig. 1, we apply two rigid transformations to an input point set of object observation, and feed them into our DPDN in parallel for making dual sets of predictions. Upon the above parallel learning, the inter-consistency term enforces cross consistency between dual predictions w.r.t. two transformations for improving the sensitivity of DPDN to pose changes, and within each learning, the individual intra-consistency term is employed to enforce self-adaptation between the correspondence and the predictions. We train DPDN on both training sets of the synthetic CAMERA25 and real-world REAL275 datasets [24]; our results outperform the existing methods on REAL275 test set under both unsupervised and supervised settings. We also conduct ablation studies that confirm the advantages of our designs. Our contributions can be summarized as follows:

- We propose a *Deep Prior Deformation Network*, termed as DPDN, for the task of category-level 6D object pose and size estimation. DPDN deforms categorical shape priors to pair with object observations in the feature space, and is thus able to establish deep correspondence for direct regression of object poses and sizes.
- Given that the considered task largely uses synthetic training data, we formulate a novel self-supervised objective upon DPDN to reduce the synthetic-toreal domain gap. The objective is built upon enforcing consistencies between parallel learning w.r.t. two rigid transformations, and has the effects of both improving the sensitivity of DPDN to pose changes, and making predictions more reliable.
- We conduct thorough ablation studies to confirm the efficacy of our designs. Notably, our method outperforms existing ones on the benchmark dataset of real-world REAL275 under both the unsupervised and supervised settings.

2 Related Work

Fully-Supervised Methods Methods of fully-supervised category-level 6D pose and size estimation could be roughly divided into two groups, *i.e.*, those based on direct regression [4,15,14] and those based on dense correspondence learning [24,20,25,3].

Direct estimates of object poses and sizes from object observations suffer from the difficulties in the learning of the full SE(3) space, and thus make demands on extraction of pose-sensitive features. FS-Net [4] builds an orientation-aware backbone with 3D graph convolutions to encode object shapes, and makes predictions with a decoupled rotation mechanism. DualPoseNet [15] encodes pose-sensitive features from object observations based on rotation-equivariant spherical convolutions, while two parallel pose decoders with different working mechanisms are stacked to impose complementary supervision. A recent work of SS-ConvNet [14] designs Sparse Steerable Convolutions (SS-Conv) to further explore SE(3)equivariant feature learning, and presents a two-stage pose estimation pipeline upon SS-Convs for iterative pose refinement.

Another group of works first learn coordinates of object observations in the canonical space to establish dense correspondence, and then obtain object poses and sizes by solving Umeyama algorithm from the correspondence in 3D space. NOCS [24], the first work for our focused task, is realized in this way by directly regressing canonical coordinates from RGB images. SPD [20] then makes the learning of canonical points easier by deforming categorical shape priors, rather than directly regressing from object observations. The follow-up works also confirm the advantages of shape priors, and make efforts on the prior deformation to further improve the qualities of canonical points, *e.g.*, via recurrent reconstruction for iterative refinement [25], or structure-guided adaptation based on transformer [3].

Unsupervised Methods Due to the time-consuming and labor-intensive annotating of real-world data in 3D space, UDA-COPE [11] presents a new setting of unsupervised domain adaptation for the focused task, and adapts a teacher-student scheme with bidirectional point filtering to this setting, which, however, heavily relies on the qualities of pseudo labels. In this paper, we exploit inter-/intra-consistency in the self-supervised objective to explore the data character-istics of real-world data and fit the data for the reduction of domain gap.

3 Self-Supervised Deep Prior Deformation Network

Given a cluttered RGB-D scene, the goal of *Category-Level 6D Object Pose* and Size Estimation is to detect object instances of interest with compact 3D bounding boxes, each of which is represented by rotation $\mathbf{R} \in SO(3)$, translation $\mathbf{t} \in \mathbb{R}^3$, and size $\mathbf{s} \in \mathbb{R}^3$ w.r.t. categorical canonical space.

A common practice to deal with this complicated task is decoupling it into two steps, including 1) object detection/instance segmentation, and 2) object pose and size estimation. For the first step, there exist quite mature techniques to accomplish it effectively, *e.g.*, employing an off-the-shelf MaskRCNN [8] to segment object instances of interest out; for the second step, however, it is still challenging to directly regress poses of unknown objects, especially for the learning in SO(3) space. To settle this problem, we propose in this paper a novel *Deep Prior Deformation Network* Φ , shortened as **DPDN**, which deforms categorical shape priors to match object observations in the feature space, and estimates object poses and sizes from the built deep correspondence directly; Fig. 2 gives an illustration. We will detail the architecture of Φ in Sec. 3.1.

Another challenge of this task is the difficulty in precisely annotating realworld data in 3D space. Although synthetic data at scale are available for the learning of deep models, their results are often less precise than those trained with annotated real-world data, due to downside effect of large domain gap. To this end, we take a mixture of labeled synthetic data and unlabeled real-world one for training, and design a novel self-supervised objective for synthetic-to-real (Sim2Real), unsupervised domain adaptation. Specifically, given a mini batch of *B* training instances $\{\mathcal{V}_i\}_{i=1}^B$, we solve the following optimization problem on top of Φ :

$$\min_{\Phi} \sum_{i=1}^{B} \frac{1}{B_1} \alpha_i \mathcal{L}_{supervised}^{\mathcal{V}_i} + \frac{1}{B_2} (1 - \alpha_i) \mathcal{L}_{self-supervised}^{\mathcal{V}_i}, \tag{1}$$

with $B_1 = \sum_{i=1}^{B} \alpha_i$ and $B_2 = \sum_{i=1}^{B} 1 - \alpha_i$. $\{\alpha_i\}_{i=1}^{B}$ is a binary mask; $\alpha_i = 1$ if the observation of \mathcal{V}_i is fully annotated and $\alpha_i = 0$ otherwise. In Sec. 3.2, we will give a detailed illustration on the self-supervised objective $\mathcal{L}_{self-supervised}$, which learns inter-consistency and intra-consistency upon DPDN, while the illustration on the supervised objective $\mathcal{L}_{supervised}$ is included in Sec. 3.3.

3.1 Deep Prior Deformation Network

For an object instance \mathcal{V} belonging to a category c of interest, we represent its RGB-D observation in the scene as $(\mathcal{I}_o, \mathcal{P}_o)$, where $\mathcal{I}_o \in \mathbb{R}^{H \times W \times 3}$ denotes the RGB segment compactly containing the instance with a spatial size of $H \times W$, and $\mathcal{P}_o \in \mathbb{R}^{N \times 3}$ denotes the masked point set with N object surface points. Direct regression of object pose and size from $(\mathcal{I}_o, \mathcal{P}_o)$ struggles in the learning of SO(3) space without object CAD model [4,15,14].

Alternatively, a recent group of works [20,25,3,11] achieve impressive results by taking advantages of strong categorical priors to establish dense point-wise correspondence, from which object pose and size could be obtained by solving of Umeyama algorithm [21]. Specifically, assuming $\mathcal{Q}_c \in \mathbb{R}^{M \times 3}$ is a sampled point set from the shape prior of c with M points, a point-wise deformation field $\mathbf{D} \in \mathbb{R}^{M \times 3}$ and a correspondence matrix $\mathbf{A} \in \mathbb{R}^{N \times M}$ are learned from a triplet of $(\mathcal{I}_o, \mathcal{P}_o, \mathcal{Q}_c)$. \mathbf{D} contains point-wise deviations with respect to \mathcal{Q}_c , deforming \mathcal{Q}_c to $\mathcal{Q}_v \in \mathbb{R}^{M \times 3}$, which represents a complete shape of \mathcal{V} in the canonical space. \mathbf{A} models relationships between points in \mathcal{Q}_v and \mathcal{P}_o , serving



Fig. 2. An illustration of our Deep Prior Deformation Network (DPDN). For an object observation \mathcal{V} , we take its image crop \mathcal{I}_o , point set \mathcal{P}_o , and shape prior \mathcal{Q}_c of the same category as inputs of a **Triplet Feature Extractor**, for the learning of their point-wise features $\mathcal{F}_{\mathcal{I}_o}$, $\mathcal{F}_{\mathcal{P}_o}$, and $\mathcal{F}_{\mathcal{Q}_c}$, respectively; then a **Deep Prior Deformer** is employed to learn a feature deformation field \mathcal{F}_D and a correspondence matrix \boldsymbol{A} to deform $\mathcal{F}_{\mathcal{Q}_c}$, yielding a feature map $\mathcal{F}_{\mathcal{Q}_o}$ in the canonical space to pair with $\mathcal{F}_{\mathcal{P}_o}$; a **Pose and Size Estimator** makes final predictions $(\boldsymbol{R}, \boldsymbol{t}, \boldsymbol{s})$ directly from the built deep correspondence between $\mathcal{F}_{\mathcal{P}_o}$ and $\mathcal{F}_{\mathcal{Q}_o}$.

as a sampler from \mathcal{Q}_v to generate a partial point set $\mathcal{Q}_o \in \mathbb{R}^{N \times 3}$, paired with \mathcal{P}_o , as follows:

$$Q_o = \mathbf{A} \times Q_v = \mathbf{A} \times (Q_c + \mathbf{D}).$$
⁽²⁾

Finally, solving of Umeyama algorithm to align \mathcal{Q}_o with \mathcal{P}_o gives out the target pose and size.

However, surrogate objectives for the learning of A and D are a step away from the true ones for estimates of pose and size; for example, small deviations of A or D may lead to large changes in the pose space. Thereby, we present a Deep Prior Deformation Network (DPDN), which implements a deep version of (2) as follows:

$$\mathcal{F}_{\mathcal{Q}_{o}} = \mathbf{A} \times \mathcal{F}_{\mathcal{Q}_{v}} = \mathbf{A} \times (\mathcal{F}_{\mathcal{Q}_{c}} + \mathcal{F}_{\mathbf{D}}), \tag{3}$$

where $\mathcal{F}_{\mathcal{Q}_c}$, $\mathcal{F}_{\mathcal{Q}_v}$, and $\mathcal{F}_{\mathcal{Q}_o}$ denote point-wise features of \mathcal{Q}_c , \mathcal{Q}_v , and \mathcal{Q}_o , respectively, and \mathcal{F}_D is a feature deformation field *w.r.t.* $\mathcal{F}_{\mathcal{Q}_c}$. The deep version (3) deforms \mathcal{Q}_c in the feature space, such that features of \mathcal{Q}_o and \mathcal{P}_o are paired to establish deep correspondence, from which object pose and size $(\mathbf{R}, \mathbf{t}, \mathbf{s})$ could be predicted via a subsequent network. Direct regression from deep correspondence thus alleviates the difficulties encountered by that from object observation. We note that upon the correspondence and the predictions, a self-supervised signal of intra-consistency could also be built for unlabeled data (see Sec. 3.2).

As depicted in Fig. 2, the architecture of DPDN consists of three main modules, including **Triplet Feature Extractor**, **Deep Prior Deformer**, and **Pose and Size Estimator**. We will give detailed illustrations shortly.

Triplet Feature Extractor Given the inputs of object observation $(\mathcal{I}_o, \mathcal{P}_o)$ and categorical shape prior \mathcal{Q}_c , we firstly extract their point-wise features $(\mathcal{F}_{\mathcal{I}_o}, \mathcal{F}_{\mathcal{P}_o}) \in$

 $(\mathbb{R}^{N \times d}, \mathbb{R}^{N \times d})$ and $\mathcal{F}_{\mathcal{Q}_c} \in \mathbb{R}^{M \times d}$, where *d* denotes the number of feature channels. Following [22], we firstly employ a PSP network [31] with ResNet-18 [9] to learn pixel-wise appearance features of \mathcal{I}_o , and then select those corresponding to \mathcal{P}_o out to form $\mathcal{F}_{\mathcal{I}_o}$. For both \mathcal{P}_o and \mathcal{Q}_c , two networks of PointNet++ [18] decorated with 4 set abstract levels are individually applied to extract their point-wise geometric features $\mathcal{F}_{\mathcal{P}_o}$ and $\mathcal{F}_{\mathcal{Q}_c}$.

Deep Prior Deformer After obtaining $\mathcal{F}_{\mathcal{I}_o}$, $\mathcal{F}_{\mathcal{P}_o}$ and $\mathcal{F}_{\mathcal{Q}_c}$, the goal of Deep Prior Deformer is to learn the feature deformation field $\mathcal{F}_{D} \in \mathbb{R}^{M \times d}$ and the correspondence matrix $A \in \mathbb{R}^{N \times M}$ in (3), and then implement (3) in the feature space to establish deep correspondence.

Specifically, as shown in Fig. 2, we obtain global feature vectors $f_{\mathcal{I}_o} \in \mathbb{R}^d$ and $f_{\mathcal{P}_o} \in \mathbb{R}^d$ of \mathcal{I}_o and \mathcal{P}_o , respectively, by averaging their point-wise features $\mathcal{F}_{\mathcal{I}_o}$ and $\mathcal{F}_{\mathcal{P}_o}$; then each point feature of $\mathcal{F}_{\mathcal{Q}_c}$ is fused with $f_{\mathcal{I}_o}$ and $f_{\mathcal{P}_o}$, and fed into a subnetwork of Multi-Layer Perceptron (MLP) to learn its deformation. Collectively, the whole feature deformation field \mathcal{F}_D could be learned as follows:

$$\mathcal{F}_{D} = \mathsf{MLP}([\mathcal{F}_{\mathcal{Q}_{c}}, \mathsf{Tile}^{M}(\boldsymbol{f}_{\mathcal{I}_{o}}), \mathsf{Tile}^{M}(\boldsymbol{f}_{\mathcal{P}_{o}})]),$$

s.t. $\boldsymbol{f}_{\mathcal{I}_{o}} = \mathsf{AvgPool}(\mathcal{F}_{\mathcal{I}_{o}}), \quad \boldsymbol{f}_{\mathcal{P}_{o}} = \mathsf{AvgPool}(\mathcal{F}_{\mathcal{P}_{o}}),$ (4)

where $[\cdot, \cdot]$ denotes concatenation along feature dimension, MLP(\cdot) denotes a trainable subnetwork of MLP, AvgPool(\cdot) denotes an average-pooling operation over surface points, and Tile^M(\cdot) denotes M copies of the feature vector.

 \mathcal{F}_{D} is used to deform the deep prior $\mathcal{F}_{Q_{c}}$ to match \mathcal{V} in the feature space. Thereby, according to (3), we have $\mathcal{F}_{Q_{v}} = \mathcal{F}_{Q_{c}} + \mathcal{F}_{D}$, with a global feature $\mathbf{f}_{Q_{v}}$ generated by averaging $\mathcal{F}_{Q_{v}}$ over M points. Then \mathbf{A} could be learned from the fusion of $\mathcal{F}_{\mathcal{I}_{o}}$, $\mathcal{F}_{\mathcal{P}_{o}}$, and N copies of $\mathbf{f}_{Q_{v}}$, via another MLP as follows:

$$A = \mathsf{MLP}([\mathcal{F}_{\mathcal{I}_o}, \mathcal{F}_{\mathcal{P}_o}, \mathsf{Tile}^N(\boldsymbol{f}_{\mathcal{Q}_v})]),$$

s.t.
$$\boldsymbol{f}_{\mathcal{Q}_v} = \mathsf{AvgPool}(\mathcal{F}_{\mathcal{Q}_v}) = \mathsf{AvgPool}(\mathcal{F}_{\mathcal{Q}_c} + \mathcal{F}_{\boldsymbol{D}}).$$
 (5)

Compared to the common practice in [20,25,3] to learn \boldsymbol{A} by fusing $\mathcal{F}_{\mathcal{I}_o}$ and $\mathcal{F}_{\mathcal{P}_o}$ with N copies of $\boldsymbol{f}_{\mathcal{Q}_c} = \operatorname{AvgPool}(\mathcal{F}_{\mathcal{Q}_c})$, our deformed version of the deep prior $\boldsymbol{F}_{\mathcal{Q}_c}$ via adding $\mathcal{F}_{\boldsymbol{D}}$ could effectively improve the quality of \boldsymbol{A} .

We also learn \mathcal{Q}_v in (2) from $\mathcal{F}_{\mathcal{Q}_v}$ as follows:

$$\mathcal{Q}_{v} = \mathsf{MLP}(\mathcal{F}_{\mathcal{Q}_{v}}) = \mathsf{MLP}(\mathcal{F}_{\mathcal{Q}_{c}} + \mathcal{F}_{D}), \tag{6}$$

such that according to (2) and (3), we have $\mathcal{Q}_o = \mathbf{A} \times \mathcal{Q}_v$ and $\mathcal{F}_{\mathcal{Q}_o} = \mathbf{A} \times \mathcal{F}_{\mathcal{Q}_v}$, respectively. Supervisions on \mathcal{Q}_v and \mathcal{Q}_o could guide the learning of \mathcal{F}_D and \mathbf{A} .

Pose and Size Estimator Through the module of Deep Prior Deformer, we establish point-to-point correspondence for the observed \mathcal{P}_o with $\mathcal{F}_{\mathcal{P}_o}$ by learning \mathcal{Q}_o in the canonical space with $\mathcal{F}_{\mathcal{Q}_o}$. As shown in Fig. 2, for estimating object pose and size, we firstly pair the correspondence via feature concatenation and apply an MLP to lift the features as follows:

$$\mathcal{F}_{corr} = \mathsf{MLP}([\mathcal{F}_{\mathcal{I}_o}, \mathsf{MLP}(\mathcal{P}_o), \mathcal{F}_{\mathcal{P}_o}, \mathsf{MLP}(\mathcal{Q}_o), \mathcal{F}_{\mathcal{Q}_o}]).$$
(7)

We then inject global information into the point-wise correspondence in \mathcal{F}_{corr} by concatenating its averaged feature \mathbf{f}_{corr} , followed by an MLP to strengthen the correspondence; a pose-sensitive feature vector \mathbf{f}_{pose} is learned from all the correspondence information via an average-pooling operation:

$$f_{pose} = \operatorname{AvgPool}(\operatorname{MLP}([\mathcal{F}_{corr}, \operatorname{Tile}^{N}(f_{corr})])), \\ s.t. \quad f_{corr} = \operatorname{AvgPool}(\mathcal{F}_{corr}).$$
(8)

Finally, we apply three parallel MLPs to regress R, t, and s, respectively:

$$\boldsymbol{R}, \boldsymbol{t}, \boldsymbol{s} = \rho(\mathsf{MLP}(\boldsymbol{f}_{pose})), \mathsf{MLP}(\boldsymbol{f}_{pose}), \mathsf{MLP}(\boldsymbol{f}_{pose}), \tag{9}$$

where we choose a 6D representation of rotation [32] as the regression target of the first MLP, for its continuous learning space in SO(3), and $\rho(\cdot)$ represents transformation from the 6D representation to the 3×3 rotation matrix \mathbf{R} .

For the whole DPDN, we could summarize it as follows:

$$\boldsymbol{R}, \boldsymbol{t}, \boldsymbol{s}, \mathcal{Q}_{v}, \mathcal{Q}_{o} = \Phi(\mathcal{I}_{o}, \mathcal{P}_{o}, \mathcal{Q}_{c}).$$
(10)

3.2 Self-Supervised Training Objective $\mathcal{L}_{self-supervised}$

For an observed point set $\mathcal{P}_o = \{\boldsymbol{p}_o^{(j)}\}_{j=1}^N$, if we transform it with $(\Delta \boldsymbol{R}_1, \Delta \boldsymbol{t}_1, \Delta \boldsymbol{t}_1, \Delta \boldsymbol{s}_1)$ and $(\Delta \boldsymbol{R}_2, \Delta \boldsymbol{t}_2, \Delta \boldsymbol{s}_2)$, we can obtain $\mathcal{P}_{o,1} = \{\boldsymbol{p}_{o,1}^{(j)}\}_{j=1}^N = \{\frac{1}{\Delta s_1} \Delta \boldsymbol{R}_1^T(\boldsymbol{p}_o^{(j)} - \Delta \boldsymbol{t}_1)\}_{j=1}^N$ and $\mathcal{P}_{o,2} = \{\boldsymbol{p}_{o,2}^{(j)}\}_{j=1}^N = \{\frac{1}{\Delta s_2} \Delta \boldsymbol{R}_2^T(\boldsymbol{p}_o^{(j)} - \Delta \boldsymbol{t}_2)\}_{j=1}^N$, respectively. When inputting them into DPDN in parallel, we have

$$\boldsymbol{R}_{\mathcal{P}_{o,1}}, \boldsymbol{t}_{\mathcal{P}_{o,1}}, \boldsymbol{s}_{\mathcal{P}_{o,1}}, \mathcal{Q}_{v,1}, \mathcal{Q}_{o,1} = \Phi(\mathcal{I}_o, \mathcal{P}_{o,1}, \mathcal{Q}_c),$$
(11)

and
$$\boldsymbol{R}_{\mathcal{P}_{o,2}}, \boldsymbol{t}_{\mathcal{P}_{o,2}}, \boldsymbol{s}_{\mathcal{P}_{o,2}}, \mathcal{Q}_{v,2}, \mathcal{Q}_{o,2} = \Phi(\mathcal{I}_o, \mathcal{P}_{o,2}, \mathcal{Q}_c),$$
 (12)

with $Q_{v,1} = \{q_{v,1}^{(j)}\}_{j=1}^M$, $Q_{o,1} = \{q_{o,1}^{(j)}\}_{j=1}^N$, $Q_{v,2} = \{q_{v,2}^{(j)}\}_{j=1}^M$, and $Q_{o,2} = \{q_{o,2}^{(j)}\}_{j=1}^N$. There exist two solutions to $(\mathbf{R}, \mathbf{t}, \mathbf{s})$ of \mathcal{P}_o from (11) and (12), respectively;

for clarity, we use subscripts '1' and '2' for $({\pmb{R}},{\pmb{t}},{\pmb{s}})$ to distinguish them:

1) $\boldsymbol{R}_1, \boldsymbol{t}_1, \boldsymbol{s}_1 = \Delta \boldsymbol{R}_1 \boldsymbol{R}_{\mathcal{P}_{o,1}}, \Delta \boldsymbol{t}_1 + \Delta \boldsymbol{s}_1 \Delta \boldsymbol{R}_1 \boldsymbol{t}_{\mathcal{P}_{o,1}}, \Delta \boldsymbol{s}_1 \boldsymbol{s}_{\mathcal{P}_{o,1}};$

2) $\mathbf{R}_2, \mathbf{t}_2, \mathbf{s}_2 = \Delta \mathbf{R}_2 \mathbf{R}_{\mathcal{P}_{o,2}}, \Delta \mathbf{t}_2 + \Delta s_2 \Delta \mathbf{R}_2 \mathbf{t}_{\mathcal{P}_{o,2}}, \Delta s_2 \mathbf{s}_{\mathcal{P}_{o,2}}.$

Upon the above parallel learning of (11) and (12), we design a novel self-supervised objective for the unlabeled real-world data to reduce the Sim2Real domain gap. Specifically, it combines an inter-consistency term \mathcal{L}_{inter} with two intraconsistency ones ($\mathcal{L}_{intra,1}, \mathcal{L}_{intra,2}$) as follows:

$$\mathcal{L}_{self-supervised} = \lambda_1 \mathcal{L}_{inter} + \lambda_2 (\mathcal{L}_{intra,1} + \mathcal{L}_{intra,2}), \tag{13}$$

where λ_1 and λ_2 are superparameters to balance the loss terms. \mathcal{L}_{inter} enforces consistency across the parallel learning from \mathcal{P}_o with different transformations, making the learning aware of pose changes to improve the precision of predictions, while $\mathcal{L}_{intra,1}$ and $\mathcal{L}_{intra,2}$ enforce self-adaptation between correspondence and predictions within each learning, respectively, in order to realize more reliable predictions inferred from the correspondence. **Inter-Consistency Term** We construct the inter-consistency loss based on the following two facts: 1) two solutions to the pose and size of \mathcal{P}_o from those of $\mathcal{P}_{o,1}$ and $\mathcal{P}_{o,2}$ are required to be consistent; 2) as representations of a same object \mathcal{V} in the canonical space, $\mathcal{Q}_{v,1}$ and $\mathcal{Q}_{v,2}$ should be invariant to any pose transformations, and thus keep consistent to each other, as well as $\mathcal{Q}_{o,1}$ and $\mathcal{Q}_{o,2}$. Therefore, with two input transformations, the inter-consistency loss \mathcal{L}_{inter} could be formulated as follows:

$$\mathcal{L}_{inter} = \mathcal{D}_{pose}(\mathbf{R}_1, \mathbf{t}_1, \mathbf{s}_1, \mathbf{R}_2, \mathbf{t}_2, \mathbf{s}_2) + \beta_1 \mathcal{D}_{cham}(\mathcal{Q}_{v,1}, \mathcal{Q}_{v,2}) + \beta_2 \mathcal{D}_{L2}(\mathcal{Q}_{o,1}, \mathcal{Q}_{o,2}),$$
(14)

where λ_1 and λ_2 are balanced parameters, and

$$\begin{split} \mathcal{D}_{pose}(\boldsymbol{R}_{1},\boldsymbol{t}_{1},\boldsymbol{s}_{1},\boldsymbol{R}_{2},\boldsymbol{t}_{2},\boldsymbol{s}_{2}) &= ||\boldsymbol{R}_{1}-\boldsymbol{R}_{2}||_{2} + ||\boldsymbol{t}_{1}-\boldsymbol{t}_{2}||_{2} + ||\boldsymbol{s}_{1}-\boldsymbol{s}_{2}||_{2},\\ \mathcal{D}_{cham}(\mathcal{Q}_{v,1},\mathcal{Q}_{v,2}) &= \frac{1}{2M} (\sum_{j=1}^{M} \min_{\boldsymbol{q}_{v,2}} ||\boldsymbol{q}_{v,1}^{(j)} - \boldsymbol{q}_{v,2}||_{2} + \sum_{j=1}^{M} \min_{\boldsymbol{q}_{v,1}} ||\boldsymbol{q}_{v,1} - \boldsymbol{q}_{v,2}^{(j)}||_{2}),\\ \mathcal{D}_{L2}(\mathcal{Q}_{o,1},\mathcal{Q}_{o,2}) &= \frac{1}{N} \sum_{j=1}^{N} ||\boldsymbol{q}_{o,1}^{(j)} - \boldsymbol{q}_{o,2}^{(j)}||_{2}. \end{split}$$

Chamfer distance \mathcal{D}_{cham} is used to restrain the distance of two complete point sets $\mathcal{Q}_{v,1}$ and $\mathcal{Q}_{v,2}$, while for the partial $\mathcal{Q}_{o,1}$ and $\mathcal{Q}_{o,2}$, we use a more strict metric of L2 distance \mathcal{D}_{L2} for point-to-point constraints, since their points should be ordered to correspond with those of \mathcal{P}_o .

Intra-Consistency Terms For an observation $(\mathcal{I}_o, \mathcal{P}_o)$, DPDN learns deep correspondence between $\mathcal{P}_o = \{p_o^{(j)}\}_{j=1}^N$ and $\mathcal{Q}_o = \{q_o^{(j)}\}_{j=1}^N$ to predict their relative pose and size $(\mathbf{R}, \mathbf{t}, \mathbf{s})$; ideally, for $\forall j = 1, \dots, N, q_o^{(j)} = \frac{1}{||\mathbf{s}||_2} \mathbf{R}^T (\mathbf{p}_o^{(j)} - \mathbf{t})$. Accordingly, the predictions $\mathcal{Q}_{o,1}$ and $\mathcal{Q}_{o,2}$ in (11) and (12) should be restrained to be consistent with $\mathcal{Q}'_{o,1} = \{\frac{1}{||\mathbf{s}_1||_2} \mathbf{R}_1^T (\mathbf{p}_o^{(j)} - \mathbf{t}_1)\}_{j=1}^N$ and $\mathcal{Q}'_{o,2} = \{\frac{1}{||\mathbf{s}_2||_2} \mathbf{R}_2^T (\mathbf{p}_o^{(j)} - \mathbf{t}_2)\}_{j=1}^N$, respectively, which gives the formulations of two intra-consistency terms based on Smooth-L1 distance as follows:

$$\mathcal{L}_{intra,1} = \mathcal{D}_{SL1}(\mathcal{Q}_{o,1}, \mathcal{Q}'_{o,1}), \quad \mathcal{L}_{intra,2} = \mathcal{D}_{SL1}(\mathcal{Q}_{o,2}, \mathcal{Q}'_{o,2}), \tag{15}$$

where

$$\mathcal{D}_{SL1}(\mathcal{Q}_1, \mathcal{Q}_2) = \frac{1}{N} \sum_{j=1}^N \sum_{k=1}^3 \begin{cases} 5(q_1^{(jk)} - q_2^{(jk)})^2, & \text{if}|q_1^{(jk)} - q_2^{(jk)}| \le 0.1\\ |q_1^{(jk)} - q_2^{(jk)}| - 0.05, & \text{otherwise} \end{cases}$$

with $Q_1 = \{(q_1^{(j1)}, q_1^{(j2)}, q_1^{(j3)})\}_{j=1}^N$ and $Q_2 = \{(q_2^{(j1)}, q_2^{(j2)}, q_2^{(j3)})\}_{j=1}^N$.

3.3 Supervised Training Objective $\mathcal{L}_{supervised}$

Given a triplet of inputs $(\mathcal{I}_o, \mathcal{P}_o, \mathcal{Q}_c)$ along with the annotated ground truths $(\hat{\mathbf{R}}, \hat{\mathbf{t}}, \hat{\mathbf{s}}, \hat{\mathcal{Q}}_v, \hat{\mathcal{Q}}_o)$, we generate dual input triplets by applying two rigid transformations to \mathcal{P}_o , as done in Sec. 3.2, and use the following supervised objective

on top of the parallel learning of (11) and (12):

$$\mathcal{L}_{supervised} = \mathcal{D}_{pose}(\mathbf{R}_{1}, \mathbf{t}_{1}, \mathbf{s}_{1}, \mathbf{R}, \dot{\mathbf{t}}, \hat{\mathbf{s}}) + \mathcal{D}_{pose}(\mathbf{R}_{2}, \mathbf{t}_{2}, \mathbf{s}_{2}, \mathbf{R}, \dot{\mathbf{t}}, \hat{\mathbf{s}}) + \gamma_{1}(\mathcal{D}_{cham}(\mathcal{Q}_{v,1}, \hat{\mathcal{Q}}_{v}) + \mathcal{D}_{cham}(\mathcal{Q}_{v,2}, \hat{\mathcal{Q}}_{v})) + \gamma_{2}(\mathcal{D}_{SL1}(\mathcal{Q}_{o,1}, \hat{\mathcal{Q}}_{o}) + \mathcal{D}_{SL1}(\mathcal{Q}_{o,2}, \hat{\mathcal{Q}}_{o})).$$
(16)

We note that this supervision also implies inter-consistency between the parallel learning defined in (14), making DPDN more sensitive to pose changes.

4 Experiments

Datasets We train DPDN on both training sets of synthetic CAMERA25 and real-world REAL275 datasets [24], and conduct evaluation on REAL275 test set. CAMERA25 is created by a context-aware mixed reality approach, which renders 1,085 synthetic object CAD models of 6 categories to real-world backgrounds, yielding a total of 300,000 RGB-D images, with 25,000 ones of 184 objects set aside for validation. REAL275 is a more challenging real-world dataset, which includes 4,300 training images of 7 scenes and 2,754 testing ones of 6 scenes. Both datasets share the same categories, yet impose large domain gap.

Implementation Details To obtain instance masks for both training and test sets of REAL275, we train a MaskRCNN [8] with a backbone of ResNet101 [9] on CAMERA25; for settings of available training mask labels, we use the same segmentation results as [20,15,14] to make fair comparisons. We employ the shape priors released by [20].

For DPDN, we resize the image crops of object observations as 192×192 , and set the point numbers of shape priors and observed point sets as M = N =1,024, respectively. In Triplet Feature Extractor, a PSP Network [31] based on ResNet-18 [9] and two networks of PointNet++ [18] are employed, sharing the same architectures as those in [10,3]. To aggregate multi-scale features, each PointNet++ is built by stacking 4 set abstract levels with multi-scale grouping. The output channels of point-wise features ($\mathcal{F}_{\mathcal{I}_o}, \mathcal{F}_{\mathcal{P}_o}, \mathcal{F}_{\mathcal{Q}_c}$) are all set as d = 128; other network specifics are also given in Fig. 2. We use ADAM to train DPDN with a total of 120,000 iterations; the data size of a mini training batch is B = 24with $B_1 : B_2 = 3 : 1$. The superparameters of $\lambda_1, \lambda_2, \beta_1, \beta_2, \gamma_1$ and γ_2 are set as 0.2, 0.02, 5.0, 1.0, 5.0 and 1.0, respectively. For each pose transformation, $\Delta \mathbf{R}$ is sampled from the whole SO(3) space by randomly generating three euler angles, while $\Delta t \in \mathbb{R}^3$ with each element $\Delta t \sim U(-0.02, 0.02)$, and $\Delta s \sim U(0.8, 1.2)$.

Evaluation Metrics Following [24], we report mean Average Precision (mAP) of Intersection over Union (IoU) for object detection, and mAP of $n^{\circ}m$ cm for 6D pose estimation. IoU_x denotes precision of predictions with IoU over a threshold of x%, and $n^{\circ}m$ cm denotes precision of those with rotation error less than n° and transformation error less than m cm.

Table 1. Quantitative comparisons of different methods for category-level 6D pose and size estimation on REAL275 [24]. 'Syn' and 'Real' denote the uses of training data of synthetic CAMERA25 and real-world REAL275 datasets, respectively. '*' denotes training with mask labels.

Method	Syn	Real	Real mAP								
		w/o Label	with Label	IoU_{50}	IoU_{75}	5°2cm	5°5cm	$10^{\circ}2$ cm	$10^{\circ}5$ cm		
Unsupervised											
NOCS [24]	\checkmark			36.7	3.4	-	3.4	-	20.4		
SPD [24]	\checkmark			71.0	43.1	11.4	12.0	33.5	37.8		
DualPoseNet [15]	\checkmark			68.4	49.5	15.9	27.1	33.1	56.8		
DPDN (Ours)	\checkmark			71.7	60.8	29.7	37.3	53.7	67.0		
Self-DPDN (Ours)	\checkmark	\checkmark		72.6	63.8	37.8	45.5	59.8	71.3		
UDA-COPE [11]	\checkmark	√*		82.6	62.5	30.4	34.8	56.9	66.0		
Self-DPDN (Ours)	\checkmark	√*		83.0	70.3	39.4	45.0	63.2	72.1		
Supervised											
NOCS [24]	\checkmark		\checkmark	78.0	30.1	7.2	10.0	13.8	25.2		
SPD [20]	\checkmark		\checkmark	77.3	53.2	19.3	21.4	43.2	54.1		
CR-Net [25]	\checkmark		\checkmark	79.3	55.9	27.8	34.3	47.2	60.8		
DualPoseNet [15]	\checkmark		\checkmark	79.8	62.2	29.3	35.9	50.0	66.8		
SAR-Net $[13]$	\checkmark		\checkmark	79.3	62.4	31.6	42.3	50.3	68.3		
SS-ConvNet [14]	\checkmark		\checkmark	79.8	65.6	36.6	43.4	52.6	63.5		
SGPA $[3]$	\checkmark		\checkmark	80.1	61.9	35.9	39.6	61.3	70.7		
DPDN (Ours)	\checkmark		\checkmark	83.4	76.0	46.0	50.7	70.4	78.4		

4.1 Comparisons with Existing Methods

We compare our method with the existing ones for category-level 6D object pose and size estimation under both unsupervised and supervised settings. Quantitative results are given in Table 1, where results under supervised setting significantly benefit from the annotations of real-world data, compared to those under unsupervised one; for example, on the metric of 5°2 cm, SPD [20] improves the results from 11.4% to 19.3%, while DualPoseNet [15] improves from 15.9% to 29.3%. Therefore, the exploration of UDA for the target task in this paper is of great practical significance, due to the difficulties in precisely annotating real-world object instances in 3D space.

Unsupervised Setting Firstly, a basic version of DPDN is trained on the synthetic data and transferred to real-world domain for evaluation; under this setting, our basic DPDN outperforms the existing methods on all the evaluation metrics, as shown in Table 1. To reduce the Sim2Real domain gap, we further include the unlabeled Real275 training set via our self-supervised DPDN for UDA; results in Table 1 verify the effectiveness of our self-supervised DPDN (dubbed **Self-DPDN** in the table), which significantly improves the precision

Input		Deep Prior Pose&Size		mAP						
$\mathcal{P}_{o,1}$	$\mathcal{P}_{o,2}$	Deformer	Estimator	IoU_{50}	IoU_{75}	5°2cm	5°5cm	$10^{\circ}2$ cm	$10^{\circ}5$ cm	
\checkmark	×	×	\checkmark	79.9	65.4	26.7	35.6	47.2	63.9	
\checkmark	\checkmark	×	\checkmark	83.3	72.9	35.2	43.9	57.3	70.8	
\checkmark	×	\checkmark	\checkmark	83.4	76.2	39.6	46.1	65.5	76.7	
\checkmark	\checkmark	\checkmark	\checkmark	83.4	76.0	46.0	50.7	70.4	78.4	
\checkmark	\checkmark	\checkmark	×	59.6	45.6	27.9	33.0	50.2	63.9	

Table 2. Ablation studies on the variants of our proposed DPDN under supervised setting. Experiments are evaluated on REAL275 test set [24].

of the basic version, *e.g.*, a performance gain of 8.1% on $5^{\circ}2$ cm from 29.7% to 37.8%.

UDA-COPE [11] is the first work to introduce the unsupervised setting, which trains deep model with a teacher-student scheme to yield pseudo labels for real-world data; in the process of training, pose annotations of real-world data are not employed, yet mask labels are used for learning instance segmentation. To fairly compare with UDA-COPE, we evaluate DPDN under the same setting; results in Table 1 also show the superiority of our self-supervised DPDN over UDA-COPE, especially for the metrics of high precisions, *e.g.*, an improvement of 10.2% on 5°5 cm. The reason for the great improvement is that UDA-COPE heavily relies on the qualities of pseudo labels, while our self-supervised objective could guide the optimization moving for the direction meeting inter-/intra- consistency, to make the learning fit the characteristics of REAL275 and decrease the downside effect of the synthetic domain.

Supervised Setting We also compare our DPDN with the existing methods, including those of direct regression [15,14], and those based on dense correspondence learning[24,20,25,13,3], under supervised setting. As shown in Table 1, DPDN outperforms the existing methods on all the evaluation metrics, *e.g.*, reaching the precisions of 76.0% on IoU₇₅ and 78.4% on 10°5 cm. Compared with the representative SS-ConvNet [14], which directly regresses object poses and sizes from observations, our DPDN takes the advantages of categorical shape priors, and achieves more precise results by regressing from deep correspondence; compared with SGPA [3], the recent state-of-the-art method based on correspondence learning (*c.f.* Eq. (2)), our DPDN shares the same feature extractor, yet benefits from the direct objectives for pose and size estimation, rather than the surrogate ones, *e.g.*, for regression of **D** and **A** in Eq. (2); DPDN thus achieves more reliable predictions.

4.2 Ablation Studies and Analyses

In this section, we conduct experiments to evaluate the efficacy of both the designs in our DPDN and the self-supervision upon DPDN.

\mathcal{L}_{inter}	$\mathcal{L}_{intra,1}$	mAP							
	$\mathcal{L}_{intra,2}$	IoU ₅₀	IoU_{75}	5°2cm	5°5cm	$10^{\circ}2$ cm	$10^{\circ}5$ cm		
×	×	71.7	60.8	29.7	37.3	53.7	67.0		
\checkmark	×	72.6	63.2	36.9	43.7	58.7	68.7		
×	 ✓ 	70.9	58.6	35.8	43.6	56.6	69.3		
\checkmark	✓	72.6	63.8	37.8	45.5	59.8	71.3		

Table 3. Ablation studies of our proposed self-supervised objective upon DPDN under unsupervised setting. Experiments are evaluated on REAL275 test set [24].

Effects of the Designs in DPDN We verify the efficacy of the designs in DPDN under supervised setting, with the results of different variants of our DPDN shown in Table. 2. Firstly, we confirm the effectiveness of our Deep Prior Deformer with categorical shape priors; by removing Deep Prior Deformer, the precision of DPDN with parallel learning on $5^{\circ}2$ cm drops from 46.0% to 35.2%, indicating that learning from deep correspondence by deforming priors in feature space indeed makes the task easier than that directly from object observations.

Secondly, we show the advantages of using true objectives for direct estimates of object poses and sizes, over the surrogate ones for the learning of the canonical point set Q_o to pair with the observed \mathcal{P}_o . Specifically, we remove our Pose and Size Estimator, and make predictions by solving Umeyama algorithm to align \mathcal{P}_o and Q_o ; precisions shown in the table decline sharply on all the evaluation metrics, especially on IoU_x . We found that the results on IoU_x are also much lower than those methods based on dense correspondence learning [20,25,3], while results on $n^{\circ}m$ are comparable; the reason is that we regress the absolute coordinates of Q_c , rather than the deviations D in Eq. (2), which may introduce more outliers to affect the object size estimation.

Thirdly, we confirm the effectiveness of the parallel supervisions in (16), *e.g.*, inputting $\mathcal{P}_{o,1}$ and $\mathcal{P}_{o,2}$ of the same instance with different poses. As shown in Table 2, results of DPDN with parallel learning are improved (with or without Deep Prior Deformer), since the inter-consistency between dual predictions is implied in the parallel supervisions, making the learning aware of pose changes.

Effects of the Self-Supervision upon DPDN We have shown the superiority of our novel self-supervised DPDN under the unsupervised setting in Sec. 4.1; here we include the evaluation on the effectiveness of each consistency term in the self-supervised objective, which is confirmed by the results shown in Table 3. Taking results on 5°, 2 cm as examples, DPDN with inter-consistency term \mathcal{L}_{inter} improves the results of the baseline from 29.7% to 36.9%, and DPDN with the intra-consistency ones $\mathcal{L}_{intra,1}$ and $\mathcal{L}_{intra,2}$ improves to 35.8%, while their combinations further refresh the results, revealing their strengths on reduction of domain gap. We also show the influence of data size of unlabeled real-world images on the precision of predictions in Fig. 3, where precisions improve along with the increasing ratios of training data.



Fig. 3. Plottings of mAP versus the ratio of unlabeled REAL275 training data under unsupervised setting. Experiments are evaluated on REAL275 test set [24].



Fig. 4. Qualitative results of DPDN on REAL275 test set [24]. 'Syn' and 'Real' denote the uses of synthetic CAMERA25 and real-world REAL275 training sets, respectively.

4.3 Visualization

We visualize in Fig. 4 the qualitative results of our proposed DPDN under different settings on REAL275 test set [24]. As shown in the figure, our self-supervised DPDN without annotations of real-world data, in general, achieves comparable results with the fully-supervised version, although there still exist some difficult examples, *e.g.*, cameras in Fig. 4 (a) and (b), due to the inaccurate masks from MaskRCNN trained on synthetic data. Under the unsupervised setting, our self-supervised DPDN also outperforms the basic version trained with only CAMERA25, by including unlabeled real-world data with self-supervision; for example, more precise poses of laptops are obtained in Fig. 4 (c) and (d).

Acknowledgements This work is supported in part by Guangdong R&D key project of China (No.: 2019B010155001), and the Program for Guangdong Introducing Innovative and Enterpreneurial Teams (No.: 2017ZT07X183).

References

- Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M.: Domainadversarial neural networks. arXiv preprint arXiv:1412.4446 (2014) 2
- 2. Azuma, R.T.: A survey of augmented reality. Presence: teleoperators & virtual environments 6(4), 355–385 (1997) 2
- Chen, K., Dou, Q.: Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2773–2782 (2021) 2, 4, 5, 7, 10, 11, 12, 13
- Chen, W., Jia, X., Chang, H.J., Duan, J., Shen, L., Leonardis, A.: Fs-net: Fast shape-based network for category-level 6d object pose estimation with decoupled rotation mechanism. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1581–1590 (2021) 2, 4, 5
- Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 1907–1915 (2017) 2
- Deng, S., Liang, Z., Sun, L., Jia, K.: Vista: Boosting 3d object detection via dual cross-view spatial attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8448–8457 (2022) 2
- 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: Sciene and Systems, RSS 2020 (2020) 2
- 8. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017) 3, 5, 10
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 7, 10
- He, Y., Sun, W., Huang, H., Liu, J., Fan, H., Sun, J.: Pvn3d: A deep pointwise 3d keypoints voting network for 6dof pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11632– 11641 (2020) 10
- Lee, T., Lee, B.U., Shin, I., Choe, J., Shin, U., Kweon, I.S., Yoon, K.J.: Udacope: Unsupervised domain adaptation for category-level object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14891–14900 (2022) 2, 4, 5, 11, 12
- Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., Kolter, J.Z., Langer, D., Pink, O., Pratt, V., et al.: Towards fully autonomous driving: Systems and algorithms. In: 2011 IEEE intelligent vehicles symposium (IV). pp. 163–168. IEEE (2011) 2
- Lin, H., Liu, Z., Cheang, C., Fu, Y., Guo, G., Xue, X.: Sar-net: Shape alignment and recovery network for category-level 6d object pose and size estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6707–6717 (2022) 11, 12
- Lin, J., Li, H., Chen, K., Lu, J., Jia, K.: Sparse steerable convolutions: An efficient learning of se (3)-equivariant features for estimation and tracking of object poses in 3d space. Advances in Neural Information Processing Systems 34 (2021) 2, 4, 5, 10, 11, 12
- 15. Lin, J., Wei, Z., Li, Z., Xu, S., Jia, K., Li, Y.: 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. pp. 3560–3569 (2021) 2, 4, 5, 10, 11, 12

- Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: International conference on machine learning. pp. 97–105. PMLR (2015) 2
- Mousavian, A., Eppner, C., Fox, D.: 6-dof graspnet: Variational grasp generation for object manipulation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2901–2910 (2019) 2
- Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. Advances in neural information processing systems **30** (2017) 7, 10
- Qin, C., You, H., Wang, L., Kuo, C.C.J., Fu, Y.: Pointdan: A multi-scale 3d domain adaption network for point cloud representation. Advances in Neural Information Processing Systems 32 (2019) 2
- Tian, M., Ang, M.H., Lee, G.H.: Shape prior deformation for categorical 6d object pose and size estimation. In: European Conference on Computer Vision. pp. 530– 546. Springer (2020) 2, 4, 5, 7, 10, 11, 12, 13
- Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis & Machine Intelligence 13(04), 376–380 (1991) 2, 5
- 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) 7
- Wang, G., Manhardt, F., Shao, J., Ji, X., Navab, N., Tombari, F.: Self6d: Selfsupervised monocular 6d object pose estimation. In: European Conference on Computer Vision. pp. 108–125. Springer (2020) 2
- Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.J.: 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. pp. 2642–2651 (2019) 1, 2, 3, 4, 10, 11, 12, 13, 14
- Wang, J., Chen, K., Dou, Q.: Category-level 6d object pose estimation via cascaded relation and recurrent reconstruction networks. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4807–4814. IEEE (2021) 2, 4, 5, 7, 11, 12, 13
- 26. Wang, Z., Jia, K.: Frustum convnet: Sliding frustums to aggregate local pointwise features for amodal 3d object detection. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1742–1749. IEEE (2019) 2
- 27. Wu, C., Chen, J., Cao, Q., Zhang, J., Tai, Y., Sun, L., Jia, K.: Grasp proposal networks: An end-to-end solution for visual learning of robotic grasps. Advances in Neural Information Processing Systems 33, 13174–13184 (2020) 2
- Zhang, Y., Deng, B., Jia, K., Zhang, L.: Label propagation with augmented anchors: A simple semi-supervised learning baseline for unsupervised domain adaptation. In: European Conference on Computer Vision. pp. 781–797. Springer (2020)
 2
- Zhang, Y., Deng, B., Tang, H., Zhang, L., Jia, K.: Unsupervised multi-class domain adaptation: Theory, algorithms, and practice. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020) 2

- Zhang, Y., Tang, H., Jia, K., Tan, M.: Domain-symmetric networks for adversarial domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5031–5040 (2019) 2
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017) 7, 10
- Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5745–5753 (2019) 8