

IST-Net: Prior-free Category-level Pose Estimation with Implicit Space Transformation

Jianhui Liu¹, Yukang Chen², Xiaoqing Ye³, Xiaojuan Qi^{1†}

¹The University of Hong Kong, ²The Chinese University of Hong Kong, ³Baidu

jhliu0212@gmail.com, yukangchen@cse.cuhk.edu.hk, yexiaoqing@baidu.com, xjq@eee.hku.hk

Abstract

Category-level 6D pose estimation aims to predict the poses and sizes of unseen objects from a specific category. Thanks to prior deformation, which explicitly adapts a category-specific 3D prior (i.e., a 3D template) to a given object instance, prior-based methods attained great success and have become a major research stream. However, obtaining category-specific priors requires collecting a large amount of 3D models, which is labor-consuming and often not accessible in practice. This motivates us to investigate whether priors are necessary to make prior-based methods effective. Our empirical study shows that the 3D prior itself is not the credit to the high performance. The keypoint actually is the explicit deformation process, which aligns camera and world coordinates supervised by world-space 3D models (also called canonical space). Inspired by these observations, we introduce a simple prior-free implicit space transformation network, namely **IST-Net**, to transform camera-space features to world-space counterparts and build correspondences between them in an implicit manner without relying on 3D priors. Besides, we design camera- and world-space enhancers to enrich the features with pose-sensitive information and geometrical constraints, respectively. Albeit simple, IST-Net achieves state-of-the-art performance based-on prior-free design, with top inference speed on the REAL275 benchmark. Our code and models are available at <https://github.com/CVMI-Lab/IST-Net>.

1. Introduction

Category-level pose estimation [39] draws great attention and plays an important role in practical applications, including robotic manipulation [6, 28, 22], augmented reality [34], and scene understanding [29, 43, 19, 20, 18, 14, 46]. Unlike instance-level pose estimation [15, 17, 23, 30, 13, 42], which requires a 3D CAD model for each object instance, this task aims at exploiting category-specific in-

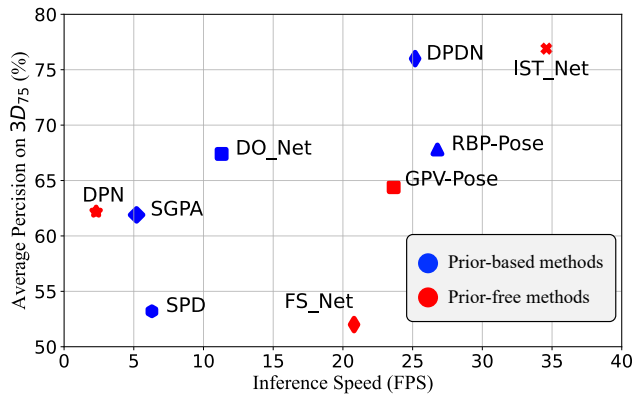


Figure 1. Comparison with competitive methods on REAL275 dataset. DPN refers to the DualPoseNet [26]. All speeds are measured on a single RTX3090Ti GPU. We use blue/red to distinguish prior-based/free methods. IST-Net achieves top performance on 3D₇₅ with the best inference speed.

formation and thus can further generalize to unseen objects within given categories.

Recently, many methods [25, 3, 35, 40, 4, 2, 5] have been proposed for category-level pose estimation, which can be categorized into two groups: prior-free methods and prior-based methods. Prior-free methods [4, 26, 39, 38] mainly focus on designing network structures to fit the training data better. These methods are relatively simple but struggle to generalize to novel objects and suffer from poor performance.

To address this issue, prior-based methods [25, 3, 35, 7, 21, 10] leverage category-specific 3D priors (templates) to guide pose estimation. They adopt a prior-driven deformation module [35] to deform the prior for synthesizing the target object in world-space. And then, they formulate the pose estimation problem as a camera- and world-space correspondence learning problem which explicitly aligns the coordinates [35]. Although considerable progress has been attained with prior-based methods [3, 45, 40], the requirements of collecting a large amount of ground-truth 3D models of target objects for obtaining the 3D prior and super-

vising training the prior deformation module hinders their practical applicability.

This motivates us to investigate the mechanism that makes prior-based methods effective. We experiment with the shape deformation module which is used to deform the given shape prior to the desired instance (Fig. 3) by replacing the shape priors with random noise and fixed shape prior from another category (see Fig. 2). We observe that the deformation module can adapt any inputs (noise or fixed prior) into a target world-space object (see Fig. 2 (b)). Besides, the model performance remains high regardless of the 3D priors (see Fig. 2 (a)). The above suggests: the **shape prior itself is not necessary** for the high performance of prior-based methods, but the deformation module that learns to synthesize world-space target objects and explicitly builds the correspondence between camera and world-space is the key as the performance degrades dramatically without prior deformation. This promotes us to investigate new ways to build camera-to-world correspondence without requiring 3D priors and models.

In this paper, we propose a simple yet effective prior-free model, named **Implicit Space Transformation Network (IST-Net)**, which implicitly sets up feature correspondence between camera-space and world-space without requiring 3D priors or ground-truth 3D models of target objects. Specifically, given the camera-space features, the network transforms them into world-space features which together with the camera-space features are further used for estimating camera poses. For learning the transformation, we propose a world-space enhancer that distills standard world-space features to supervise the transformed features. Note that the standard world-space features are obtained by transforming the input target object into its world-space with the ground-truth pose and feeding them into a feature extractor. Besides, given camera-space inputs, the backbone network’s feature extraction capabilities are boosted by introducing an auxiliary pose estimation loss, namely the camera-space enhancer. Notably, both enhancers are training only, which brings considerable performance improvements without introducing computational overhead.

Our main contributions are summarized as follows:

- We investigate prior-based methods and find that shape priors are not necessary for obtaining high performance while building the camera and world-space correspondence with prior deformation is a key factor.
- We propose a simple yet effective **Implicit Space Transformation Network (IST-Net)** that implicitly builds the correspondence between camera- and world-space on feature-level without requiring 3D models or priors.
- We introduce two different space enhancers to facili-

tate learning the transformation and enhance their representation capability for pose estimation.

- We conduct a series of experiments on REAL275 [39] and Wild6D [10] datasets to demonstrate the effectiveness of the proposed method. Notably, IST-Net is currently the only prior-free method that achieves state-of-the-art performance on the REAL275 benchmark and attains notable gains over the prior-based method in terms of efficiency and accuracy (see Fig. 1).

2. Related Works

2.1. Prior-free Methods

Prior-free methods focus on designing architectures for predicting the object pose in a concise manner. Sahin *et al.* [33] propose an “Intrinsic Structure Adaptor” to adapt the distribution shifts arising from shape discrepancies. Wang *et al.* [39] first introduce a new category-level benchmark by normalizing all object instances into a shared canonical representation named *Normalized Object Coordinate Space-(NOCS)* and try to recover the angle of view in *NOCS* for pose estimation. Chen *et al.* [2] introduce a learned canonical shape space to handle intra-class variation. Chen *et al.* [5] attempt to synthesize image matches upon neural rendering in order to verify the probability of each possible pose candidate for pose estimation. Wang *et al.* [37] propose 6D-PACK which learns to compactly represent an object by a handful of 3D key points based on the motion information and compute the pose by tracking. In pursuit of more efficient and direct pose estimation, a few methods [4, 7, 26] work on designing the network in an end-to-end manner. Chen *et al.* [4] decouple the rotation into two mutually orthogonal vectors to fully decode the orientation information which allows the network to naturally handle the circle symmetry object. Di *et al.* [7] embody the geometric insights with bounding box projection to enhance the learning of category-level pose-sensitive features. Lin *et al.* [26] introduce DualPoseNet which is composed of two parallel pose decoders on top of a shared pose encoder. The two decoders work in an implicit and explicit manner with the restriction of the predicted pose consistency.

2.2. Prior-based Methods

Since the severe intra-class variation, the generalization of the prior-free models is greatly suppressed. To alleviate this issue, some literature [10, 35, 25, 44, 41] turn to focus on prior-based methods. Tian *et al.* [35] present a general solution. They set up shape prior for each category upon an autoencoder and then use these priors as the standard template to reconstruct the canonical model for each instance. Chen *et al.* [2] use a variational autoencoder (VAE) [16] for reconstructing standard object shape, followed by a fully sparse convolution network for pose regression. Wang *et*

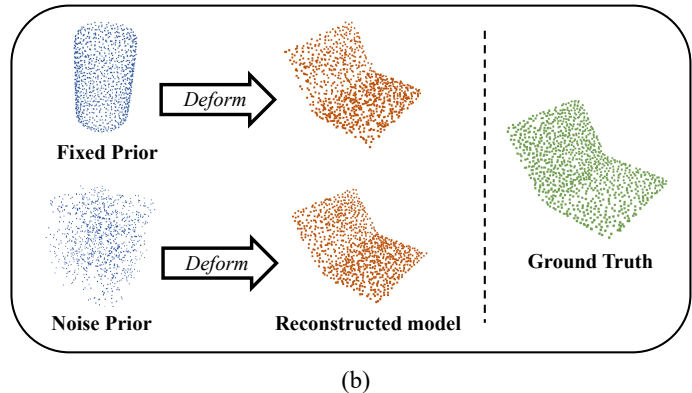
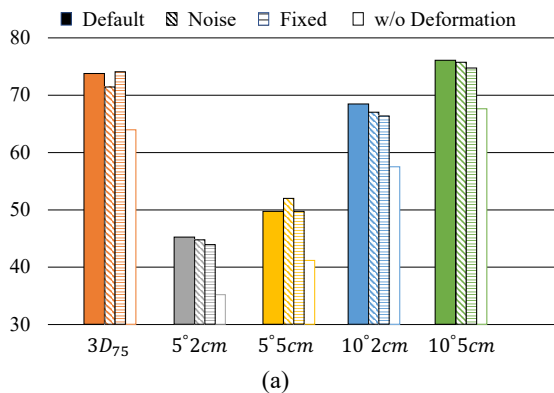


Figure 2. Illustration of the empirical experiments on the shape prior. Subfigure (a) shows the performance impact of using different forms of prior. Subfigure (b) presents the visualization of the prior deformation. Even if the class-independent and noise-generated prior are used, high-quality reconstruction results can still be obtained.

al. [40] propose a cascaded relation network to capture the underlying relations of multi-source inputs. Kai *et al.* [3] utilize a transformer network [8] to model the global structure similarity between prior and target object, based on which the object semantic information is injected into the prior feature to dynamically adapt the category-level prior to each particular object. Fan *et al.* [9] adopt a shape prior guided reconstruction network and a discriminator network to learn high-quality canonical representations. Zhang *et al.* [45] use the shape priors as the indicator to predict pose and zero-mean residual vectors which encapsulate the spatial cues of the pose and enable geometry-guided consistency terms. Zhang *et al.* [44] learn dense correspondences between input images and the canonical shape prior via surface embedding. Lin *et al.* [25] establish deep correspondence in the feature space between shape prior and canonical model, which yields a surprising performance boost.

3. Analysis of Shape Priors

3.1. Preliminary

To overcome intra-class variation, the prior deformation, as a practical module, has been widely adopted by recent works [9, 3, 25]. The vanilla version of prior deformation can be divided into two parts: 1) generating shape priors and 2) leveraging shape priors to develop prior deformation techniques.

For the former one, the common solution is to train an autoencoder with **various object models** sampled from ShapeNet [1], then acquire the category-level shape embedding by averaging the latent vectors output by the encoder. These shape embeddings will be fed into the decoder to get the shape priors. It is worth mentioning that this process needs to rely on a large number of 3D models to obtain a general prior of a category.

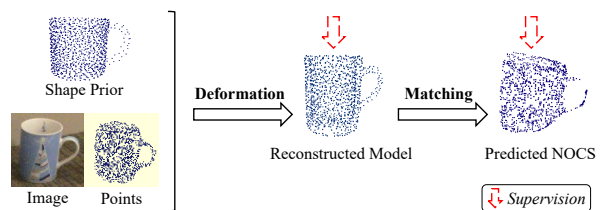


Figure 3. Illustration of the vanilla version of prior deformation. NOCS is equivalent to the coordinate in the world-space.

For the latter one, we use Fig. 3 to illustrate the process. Given the shape prior, image patch, and observed points, the network first learns a deformation field that deforms the shape prior to the desired object instance, which is **supervised by a ground-truth 3D model**. Furthermore, the network outputs a matching matrix that indicates the point-to-point correspondences between the observed target points and the reconstructed models. These correspondences transform the models to the viewpoint in the world coordinate system. With the information from camera-space (depth images) and world-space (matched priors), pose parameters can be easily solved via Umeyama algorithm [36] or pose regression by neural networks.

3.2. Is Shape Prior Necessary?

We conduct extensive experiments to verify whether the shape prior is necessary to address the intra-class variation problem. Specifically, we choose a competitive candidate from prior-based methods, DPDN [25]. (For other methods please refer to the Appendix). We set up the following settings:

- **Case-1:** *Official baseline.*
- **Case-2:** *All the categories share the same prior (can) in replace of the class-specific priors.*

- **Case-3:** Using random noise restricted to the unit cube instead of standard shape priors.
- **Case-4:** Removing the prior deformation from original framework.

From Fig. 2, we can conclude that *category-specific priors are not necessary* as the model can learn to deform a shape to match the target from even random noise. The reason stems from the explicit supervision of ground-truth 3D models during training, which enables the model to learn to deform any given prior (e.g., random noise) to reconstruct a target object. Nevertheless, the prior deformation module contributes to the prior-based model as the performance drops dramatically without this module (see Fig. 2 (a) w/o deformation). The above suggests that the key to the success of *prior-based methods is the deformation module that aligns objects in the camera and world-space and facilitates building correspondences, but not the prior itself.*

Hence, we explore new ways to transform camera space inputs to world-space and implicitly builds the correspondence between them without relying on priors or 3D ground-truth models in training. We present a prior-free implicit feature transformation network (**IST-Net**) with details unfolded in Sec. 4.2. On one hand, our method gets rid of the dependence on a large number of 3D models required by prior-based methods. On the other hand, we make the model aware of the information in world-space with a simple design which further allows us to develop an efficient prior-free pose estimator without sacrificing model performance.

4. Method

Before delving into the details, we first show some mathematical notations in the category-level 6D pose estimation task for clarity. $P_o \in \mathbb{R}^{N_o \times 3}$ and $I_o \in \mathbb{R}^{h \times w \times 3}$ refer to the observed point cloud and corresponding RGB image, where (h, w) denotes the size of the image and N_o is the number of object points. Given these inputs, the objective is to estimate the pose of the input instance, including rotation $R \in SO(3)$, translation $t \in \mathbb{R}^3$ and size $s \in \mathbb{R}^3$.

4.1. Overview

In this section, we will introduce the detailed architecture of the proposed IST-Net, as shown in Fig. 4. Given camera-space RGB image I_o and point cloud P_o as inputs, a CNN backbone network, and PointNet++ are used to extract image feature F_{I_o} and point feature F_{P_o} and F_e respectively. Then, the learned features are fed into the implicit space transformation module (Sec. 4.2) which performs implicit camera-space to world-space feature $F_{\tilde{Q}_o}$ transformation. We also provide a detailed analysis of our implicit design and compare it with its explicit counterparts in Sec. 4.2.

Further, given point cloud P_o , point-wise aligned world-space feature ($F_{\tilde{Q}_o}$) and camera-space feature (F_{I_o} and F_{P_o}), the pose estimator directly regresses the poses $\{R, t, s\}$. We adopt the pose estimator in DPDN [25]. Details are included in the Appendix. Finally, we introduce two auxiliary modules, namely camera-space enhancer and world-space enhancer (Sec. 4.3 and Sec. 4.4), to boost feature representations and facilitate learning the implicit transformation network.

4.2. Implicit Space Transformation

Since we have shown in Sec. 3.2, shape prior itself is not necessary, the important factor is how to transform camera-space inputs into world-space, align them and build their correspondences. To address these issues, we propose an implicit space transformation module, which transforms camera-space features to world-space in an implicit manner without resorting to ground-truth 3D models during training.

As depicted in Fig. 4, given an RGB image I_o and corresponding object point cloud P_o , we first use the feature extractor to acquire semantic feature which sampled as $F_{I_o} \in \mathbb{R}^{N_o \times d}$ and geometry feature $F_{P_o} \in \mathbb{R}^{N_o \times d}$. Notably, before entering PointNet++, the point cloud will be pre-processed, which means subtracting the average value of all coordinates. In other words, the feature extractor can focus on the relative geometric relationship, avoiding the interference of spatial locations. However, this will make the network insensitive to location information, which is indispensable for the estimation of the translation and rotation. Hence, we utilize a Multi-Layer Perceptron (MLP) to encode the accurate position information into latent space as F_e . Then we concatenate three of them as the input of the implicit space transformation module. Sequentially, an MLP is employed to fuse these input features followed by a global average pooling layer. Then the local and global features are concatenated and fed into an MLP to get the world-space feature $F_{\tilde{Q}_o}$. This process can be formulated as follows:

$$\begin{aligned} F_L &= \text{MLP}([F_{P_o}, F_{I_o}, F_e]), \\ F_G &= G(F_L), \\ F_{\tilde{Q}_o} &= \text{MLP}([F_L, F_G]), \end{aligned} \quad (1)$$

where $[\dots]$ refers to concatenation and G denotes the global average pooling. F_L and F_G refers to the local and global feature.

It should be noted that the implicit transformation network may not certainly transform the camera space feature to world space without meaningful supervision. We thus introduce a reconstruction-based loss L_{rec} to regularize the learning. Specifically, given $F_{\tilde{Q}_o}$ as inputs, we use another MLP to predict reconstruct per-point coordinate in the

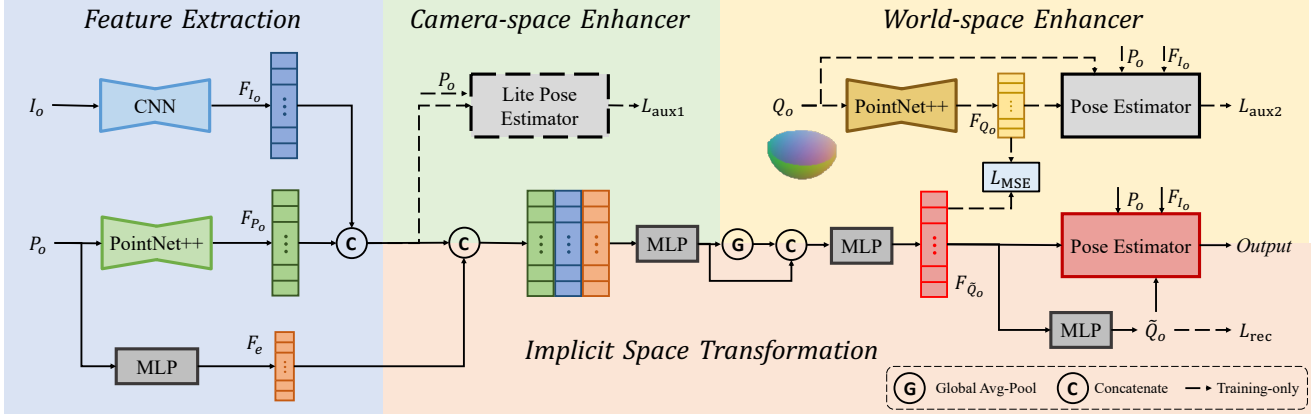


Figure 4. Illustration of our proposed Implicit Space Transformation Network IST-Net for 6D pose estimation.

world-space \tilde{Q}_o and supervise it using ground-truth world-space coordinate Q_o . Note that Q_o can be obtained by transforming input point cloud P_o using ground-truth poses without resorting to 3D models. Following [35], we adopt smooth-L1 loss as:

$$L_{rec} = L_{SL1}(\tilde{Q}_o, Q_o), Q_o = \Gamma(P_o, R, t, s), \quad (2)$$

where L_{SL1} denotes the Smooth-L1 loss [27] and Γ indicates the 3D geometric transformation operation according to pose. This supervision will encourage the transformed features to be in the world-space.

Explicit Counterpart To better showcase the merits of our implicit space transformation, we set up a comparison with its explicit counterpart. As shown in Fig. 5, we firstly use an extra pose estimator to predict a group of pose parameters and then use them to explicitly transform the camera coordinate to the world coordinate as \tilde{Q}_o . And then we utilize another PointNet++ to extract world-space features $F_{\tilde{Q}_o}$. Finally, we feed $F_{\tilde{Q}_o}$, \tilde{Q}_o , and camera space items into the estimator for pose regression. Explicit transformation is more intuitive but struggles in model redundancy. In contrast, our method is quite concise during inference, reaching a better balance between efficiency and performance. More analysis are presented in Sec. 5.5.

4.3. Camera-space Enhancer

Since the transformed feature $F_{\tilde{Q}_o}$ is derived from camera-space features, the representative power of F_{P_o} and F_{I_o} becomes crucial. Therefore, we propose a camera-space enhancement strategy to strengthen the camera-space features supervised by the auxiliary pose estimation task. The camera-space enhancer can easily be an auxiliary pose estimator whose training doesn't rely on the world-space 3D model. This encourages the feature extractor to learn repre-

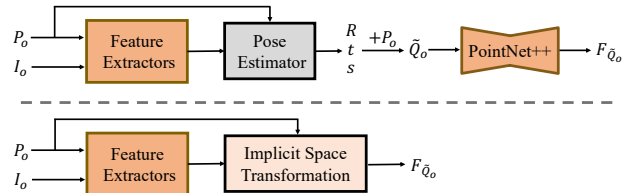


Figure 5. Comparison between implicit transformation (sub-graph below) and explicit transform (sub-graph above).

sentations that bridge camera and world-space and implicitly establish a correspondence with world-space features, which is helpful for subsequent transformations and final predictions.

4.4. World-space Enhancer

In the implicit space transformation module, we predict the world-space coordinates, which encourage the network to transform camera-space features into world-space feature $F_{\tilde{Q}_o}$ and regularize the learning indirectly. Here, we propose world-space enhancer which directly enforces $F_{\tilde{Q}_o}$ to be similar to world-space counterparts

As shown in Fig. 4, the world-space enhancer takes the world-space point cloud Q_o derived using Eq. (2) as input and produces (ground-truth) world-space features F_{Q_o} using PointNet++. F_{Q_o} is used to directly supervise $F_{\tilde{Q}_o}$ as

$$L_{feat} = L_{MSE}(F_{\tilde{Q}_o}, F_{Q_o}), \quad (3)$$

$$L_{MSE}(F_{\tilde{Q}_o}, F_{Q_o}) = \frac{1}{N_o \cdot d} \sum_{i=1}^{N_o} \sum_{j=1}^d \|f_{\tilde{Q}_o}^{ij} - f_{Q_o}^{ij}\|^2.$$

To train the feature extractor (*i.e.* PointNet++) and obtain high-quality features F_{Q_o} , we use another auxiliary pose estimator with the same architecture as the one that generates

Method	Prior	$3D_{25}$	$3D_{50}$	$3D_{75}$	$5^\circ 2cm$	$5^\circ 5cm$	$10^\circ 2cm$	$10^\circ 5cm$	$10^\circ 10cm$	Speed (FPS)
NOCS [39]	\times	84.9	80.5	30.1	7.2	10	13.8	25.2	26.7	4.8
CASS [2]	\times	84.2	77.7	-	-	23.5	-	58.0	58.3	-
FS-Net [4]	\times	-	-	-	-	28.2	-	60.8	-	20.8
DualPoseNet [26]	\times	-	79.8	62.2	29.3	35.9	66.8	-	-	2.3
GPV-Pose [7]	\times	84.2	83.0	64.4	32.0	42.9	-	73.3	74.6	23.6
SPD [35]	\checkmark	83.4	77.3	53.2	19.3	21.4	43.2	54.1	-	6.3
CR-Net [40]	\checkmark	-	79.3	55.9	27.8	34.3	47.2	60.8	-	-
SGPA [3]	\checkmark	-	80.1	61.9	35.9	39.6	61.3	70.7	-	5.2
DO-Net [24]	\checkmark	-	80.4	63.7	24.1	34.8	45.3	67.4	-	11.3
RBP-Pose [45]	\checkmark	-	-	67.8	38.2	48.1	63.1	79.2	-	26.8
DPDN [25]	\checkmark	84.3	83.4	76.0	46.0	50.7	70.4	78.4	80.4	25.2
IST-Net (Ours)	\times	84.3	82.5	76.6	47.5	53.4	72.1	80.5	82.6	34.6

Table 1. Comparisons with state-of-art methods on REAL275 dataset. We summarize the pose estimation results reported in the original papers. **Prior** refers to whether the method builds upon shape priors. ‘-’ denotes no results are reported under this metric. All the experiments for inference speed are conducted with a single RTX3090Ti GPU.

pose estimates to produce supervisory signals. Notably, the only difference from the main pose estimator (red boxes in Fig. 4) is the world-space inputs, fully encoded ground-truth world-space coordinates Q_o . In this way, we believe that F_{Q_o} shares the same feature space with $F_{\tilde{Q}_o}$ but is more accurate, compact, and friendly for pose estimation.

In conclusion, this module generates a standard latent feature that further provides restriction from a high-level aspect, regularizing $F_{\tilde{Q}_o}$ with implicit geometric information. In conjunction with constraint L_{rec} in Sec. 4.2, the space transformation is guided to learn to transform camera-space features to world-space.

4.5. Overall Loss Function

In summary, the overall loss function is as below:

$$L = L_{main} + L_{aux1} + L_{aux2} + \lambda_f L_{feat} + \lambda_r L_{rec}, \quad (4)$$

where λ_f and λ_r are hyper-parameters which used to balance the individual loss contributions. L_{main} , L_{aux1} , and L_{aux2} refer to the supervision for the outputs of main pose estimator and two feature enhancers which share the same loss format as:

$$L_{pose} = \|R_{pred} - R_{gt}\|_2 + \|t_{pred} - t_{gt}\|_2 + \|s_{pred} - s_{gt}\|_2, \quad (5)$$

where subscripts ‘‘pred’’ and ‘‘gt’’ denote the predicted and ground truth pose parameters, respectively.

5. Experiments

5.1. Datasets

REAL275 & CAMERA25: Our model is trained on both the virtual dataset, CAMERA25, and the real dataset,

REAL275 [39], and conducted an evaluation on REAL275 test split. CAMERA25 contains 300k synthetic RGB-G images, which are generated by rendering 1,085 synthetic objects with real-world backgrounds. REAL275 includes 8k RGB-D images, where 4300 images are split for training, 950 images for validation, and 2750 images for testing. In both datasets, there are 6 categories, including bottle, bowl, camera, can, laptop, and mug.

Wild6D: Wild6D [10] contains 5,166 videos with 1722 object instances and 5 categories (bottle, bowl, camera, laptop, and mug). Among this data, 486 videos are split for model evaluation.

5.2. Implementation Details

Following previous work [25], we use the off-the-shelf MaskRCNN [11] with a backbone of ResNet101 [12] for generating high-quality instance masks. We adopt a PSP Network [47] based on ResNet-18 [12] for semantic feature extraction and PointNet++ [32] for point-level feature extraction. The number of object point N_o is set as 1024 and the size of the RGB image is resized to 192×192 . We adopt several commonly used data augmentation, including random uniform noise, random rotational and translational perturbations, and bounding box-based adjustment, which is proposed by FS-Net [4]. For the issue of symmetry, we follow previous work [31], mapping ambiguous rotations in training images with the same object appearance to a canonical rotation. Given an input image, the network predicts the canonical rotation, which is the best that can be done in the presence of e of symmetries. The hyper-parameters of the loss weights λ_f and λ_r are set to 10 and 1, respectively. All the experiments are conducted on 2 RTX3090Ti GPUs with a batch size of 24, and the ratio of real data to synthetic data is 1:3. For a fair comparison, the total training epoch is fixed to 30 epochs, and all the modules are trained in an end-

Method	Data	$3D_{25}$	$3D_{50}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
CASS [2]	C+R	19.8	1.0	0.0	0.0	0.0	0.0
SPD [35]	C+R	55.5	32.5	2.6	3.5	9.7	13.9
DualPoseNet [26]	C+R	90.0	70.0	17.8	22.8	26.3	36.5
GPV-Pose [7]	C+R	91.3	67.8	14.1	21.5	23.8	41.1
RePoNet [10]	C+W*	84.7	70.3	29.5	34.4	35.0	42.5
Self-Pose [44]	W*	92.3	68.2	32.7	35.3	38.3	45.3
IST-Net (Ours)	C+R	93.4	79.6	30.7	35.8	37.1	43.7

Table 2. Comparison with state-of-art methods on Wild6D dataset. The ‘‘Data’’ column refers to the data type for training. C=CAMERA25, R=REAL275, W=Wild6D, *=not using pose annotation.

to-end manner. During inference, only the feature extractor and implicit space transformation module are preserved.

5.3. Evaluation Metrics

We follow [39, 7] and utilize the widely adopted metrics for evaluation, including the mean precision of 3D intersection over union (IoU) to jointly evaluate rotation, translation, and size. Besides, the $5^{\circ}2cm$, $5^{\circ}5cm$, $10^{\circ}2cm$, $10^{\circ}5cm$ and $10^{\circ}10cm$ are used to evaluate the rotation and translation error directly, specifically, only the prediction error under both thresholds can be considered correct.

5.4. Comparison with State of the Arts

We present the results of IST-Net with state-of-the-art methods on REAL275 [39], as shown in Tab. 1. For comparison with prior-free methods, we surpass others with a large gap on all evaluation metrics, *e.g.*, we reach 47.5 and 76.6 on $5^{\circ}2cm$ and D_{75} , which outperform GPV-Pose by 19.4% and 41%. As for prior-based methods, compared with the current most powerful method DPDN [25], we still perform significant improvements in most of the metrics. *e.g.*, 76.6 vs 76 on D_{75} , 47.5 vs 46.0 on $5^{\circ}2cm$, 53.4 vs 50.7 on $5^{\circ}5cm$, 80.5 vs 78.4 on $10^{\circ}5cm$ and 82.6 vs 80.4. Notably, it is the first time for the prior-free methods to achieve comparable or even higher performance compared with prior-based methods on REAL275. In addition, we present a per-class comparison between IST-Net and DPDN, as shown in Fig. 7. Notably, our method performs better on the prediction of rotation, the error curve is steeper on the geometrically complex object, *e.g.*, camera, which clearly proves the effectiveness of our contributions. Apart from the performance, model efficiency is also worthy of attention, we list the inference time in the last column of Tab. 1. IST-Net reaches top inference speed which far exceeds other methods by more than **25% acceleration**.

We conduct experiments on a larger dataset, Wild6D [10], to further verify the effectiveness of the proposed method. We directly test our model which are trained on REAL275 and CAMERA25 datasets without extra fine-tuning. The results are reported in Tab 2. It can be observed that IST-Net is much better than those designed for the REAL275 dataset and has an obvious improvement

	IST	CE	WE	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
E1				70.7	35.2	42.9	58.8	72.2
E2	✓			72.2	42.9	48.5	68.7	78.1
E3	✓	✓		73.9	44.2	52.5	68.5	79.2
E4	✓		✓	75.9	43.5	48.9	70.4	80.4
E5	✓	✓	✓	76.6	47.5	53.4	72.1	80.5

Table 3. Ablation on different configurations of network architectures. IST refers to the implicit space transformation, CE and WE denote the camera and world-space enhancement. Note: In E1 we just remove the supervision of IST, not the module itself.

PE	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
w	82.5	76.6	47.5	53.4	72.1	80.5
w/o	81.3	73.0	41.5	47.4	68.8	78.8

Table 4. Ablation on position encoding term. PE refers to the position encoding term.

in various matrices. Compared with RePoNet [10] and Self-Pose [44] both of *which are trained upon Wild6D, our method shows good generalization without fine-tuning on the target dataset*. We attribute this to the fact that our model is designed to learn the correspondence between camera and world space. We achieve the highest performance of 93.4 and 79.6 on $3D_{25}$ and $3D_{50}$. As for other metrics, IST-Net can also reach a similar results to Self-Pose. These analysis and results verify the potential of our method.

5.5. Ablation Study

Effects of Proposed Modules. We ablate the combination of different modules of the proposed method, the results are shown in Tab 3. Firstly we present the effectiveness of the implicit space transformation module (IST). By adding this module, we can easily observe that the baseline is greatly lifted, suggesting that transforming camera-space features to world-space counterparts and building the correspondence between them in an implicit manner indeed can benefit the pose estimation. Besides this, after adding the camera-space enhancer (CE), the precision on $5^{\circ}5cm$ increases from 48.5 to 52.5, the reason is that with this auxiliary module, the feature extractors are enriched with more pose-sensitive information, which is beneficial to the quality of feature transformation and to improve the accuracy of final pose estimation. In addition, we show the advantage of world-space enhancer (WE), by combining it with IST. The results (E4) show that WE can further extend the performance, especially on $10^{\circ}2cm$ and $10^{\circ}5cm$, which indicates that high-level supervision provides additional information different from low-level constraint. Finally, by combining all modules, we reach relatively competitive performance.

Effects of Position Encoding Term. In this part, we verify the effect of the positional encoding (PE) term, the results are shown in Tab. 4. Obviously, without the PE module, the performance drops significantly on $5^{\circ}2cm$ and $5^{\circ}5cm$. This

	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$	Param.	FPS
Implicit	73.9	44.2	52.5	68.5	79.2	21M	34
Explicit	75.1	45.0	50.7	69.6	80.0	24M	22

Table 5. Comparison with explicit space transformation. Param. refers to the number of parameters.

Method	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
SGPA [3]	80.1	61.9	35.9	39.6	61.3	70.7
IST-Net ^M	82.5	72.7	35.8	38.4	64.2	72.4

Table 6. Ablation on the predicted world coordinate. M is our matching-based variant, using the Umeyama algorithm for testing.

	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
DPDN	78.8	48.6	11.1	13.1	31.6	45.3
IST-Net	81.0	53.9	25.5	29.7	44.9	53.9

Table 7. Performance with limited training data (REAL275 only).

	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
DPDN	83.0 (-0.4)	75.4 (-0.6)	41.5(-3.7)	44.9 (-5.8)	69.2 (-1.2)	76.5(-1.9)
IST-Net	82.2 (-0.2)	75.9(-1.0)	41.0(-4.2)	46.7(-5.2)	69.6(-1.3)	79.0(-1.9)

Table 8. Robustness test on REAL275 dataset with Gaussian noise.

Unseen class	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$
Bowl	100.0 / 100.0	99.4 / 22.0	36.1 / 0.0	49.1 / 7.1	59.7 / 0.0	97.6 / 35.1
Laptop	17.7 / 18.2	0.4 / 0.7	0.0 / 0.0	0.0 / 0.0	0.0 / 0.0	0.0 / 0.0

Table 9. Results on unseen class. Color map: Black for IST-Net, Gray for DPDN (prior-based).

λ_f	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}10cm$	$10^{\circ}5cm$
1	81.6	72.6	40.8	47.0	68.0	77.7	79.8
3	82.7	76.1	42.9	48.9	70.1	80.1	82.1
5	83.2	76.1	43.6	50.6	69.1	79.7	81.8
10	82.5	76.6	47.5	53.4	72.1	80.5	82.6
20	82.0	75.2	45.0	51.5	68.0	77.9	80.0
50	82.8	75.3	41.5	47.2	68.8	77.9	80.0
100	83.1	76.3	45.3	50.4	70.4	78.9	81.1

Table 10. Ablation on the loss weight λ_f .

in turn proves that the PE module makes up for the position lost by the feature extractor, which benefits pose regression.

Comparison with Explicit Space Transformation. To further verify the effectiveness of the proposed implicit space transformation, we set up an experiment with its explicit counterpart. For reaching a fair comparison, only WE and IST are included in the implicit candidate. From the results shown in Tab. 5, we can easily find the results of the two methods are very close. However, our method yields obvious superiority in speed (34Hz vs 22Hz) and parameter quantity (21M vs 24M) which attributes to the succinct feature space transformation instead of introducing repetitive modules for extracting features from coordinates. This further indicates the potential of the proposed modules.

Ablations on Predicted World Coordinate. Considering that we predict the coordinate of observed points in world-space for supervising implicit space transformation from a low-level perspective. Therefore, the quality of the gener-

Type	$3D_{50}$	$3D_{75}$	$5^{\circ}2cm$	$5^{\circ}5cm$	$10^{\circ}2cm$	$10^{\circ}5cm$	$10^{\circ}10cm$
MSE	82.5	76.6	47.5	53.4	72.1	80.5	82.6
L1	82.8	74.6	44.6	50.4	69.8	78.8	81.1

Table 11. Ablation on the loss type of L_{feat}

ated coordinate can also reflect the effectiveness of the proposed method. To verify this, we use the predicted coordinate and observed points in camera space for solving the pose parameters by Umeyama algorithm [36]. As shown in Tab. 6, our method achieves comparable results with SGPA, even with significant improvement on $3D_{75}$ (72.7 vs 61.9), indicating that the network can reconstruct the perspective in world-space without introducing shape prior.

Performance with limited training data. We train IST-Net and DPDN only on REAL275 with the results shown in Tab. 7. IST-Net performs better with limited training data. We attribute this to the fact that the prior-based method is more complex with a larger number of parameters, which requires more data to prevent overfitting.

Robustness test. We investigate the model robustness in Tab.R. 8. Please note that there is no dedicated design for addressing robustness in all methods. We apply Gaussian noise with a maximum amplitude of 3mm on test data, which is stronger than that in train augmentation (1mm) for DPDN and IST-Net. The result shows the two methods perform similarly in robustness evaluations.

Generalization to unseen objects. We evaluate whether the trained model can be directly transferred to unseen categories, as shown in Tab. 9. We take bowl and laptop as unseen classes and others left for training. It can be observed that IST-Net generalizes much better than DPDN on the simple object (bowl). However, both of them struggle to generalize shapes with complicated structures (laptop).

Ablate on Loss Weight. We further ablate the effect of different choices of λ_f on pose accuracy. We gradually enlarge the λ_f from 1 to 100. The comparative results are shown in Tab. 10. When λ_f is too small, the supervision is limited, and when it is too large, the supervision from ground truth will be weakened. Overall, when λ_f is set as 10, we reach the best performance.

Ablate on Loss Type. In this experiment, we ablate the effect of different loss types of L_{feat} . We present the results of MSE loss and L1 loss in Tab. 11. In contrast, MSE Loss has more advantages. It does not require two features to be completely similar but imposes strong constraints on places with large differences, which makes the imitation between features easier to learn.

6. Qualitative Analysis

In Fig 6, we visually compare our methods and DPDN on the REAL275 dataset. It clearly shows the superiority of our method. As highlighted in blue boxes, DPDN

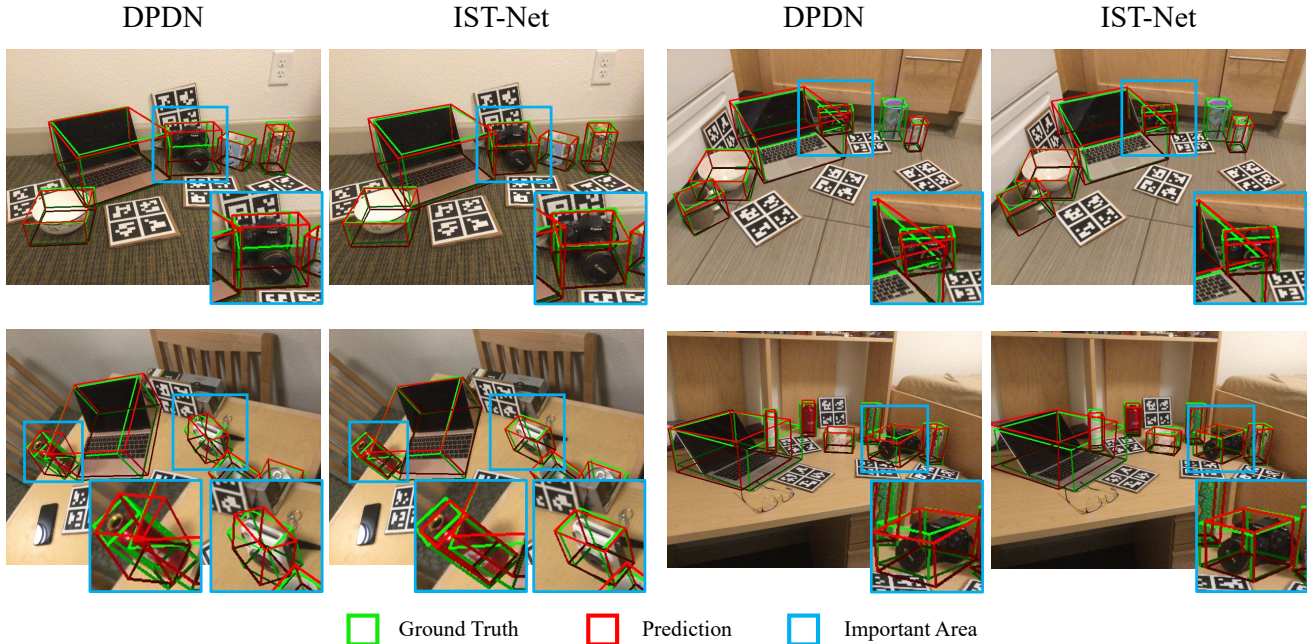


Figure 6. Qualitive comparison between IST-Net and DPDN on REAL275 dataset.

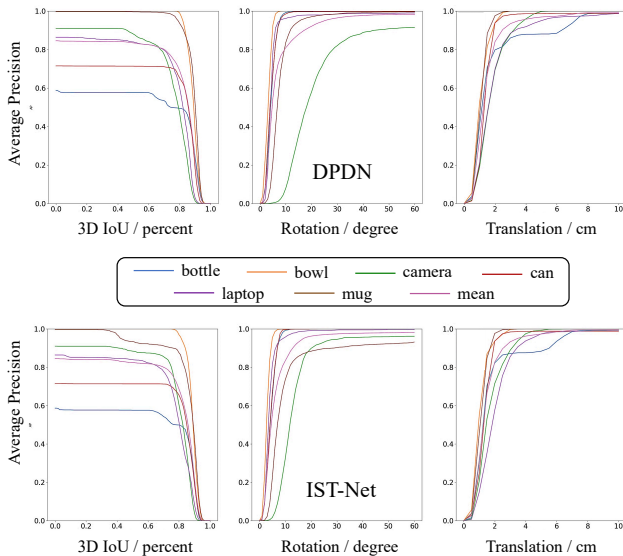


Figure 7. Quantitative comparison with DPDN [25] on REAL275 in terms of average precision in 3D IoU, Rotation and Translation.

easily gets stuck in the object with complex structure, *e.g.*, camera, which presents as apparent deviations of predicted boxes. This reflects the prior deformation has a poor capability for modeling challenging cases. By contrast, IST-Net demonstrates strong performance in predicting accurate rotation and translation estimations. The reason is that with the implicit transformation, the geometric structures are transmitted to the world-space with the feature, which ensures the sensitivity of the network to complex structures.

7. Conclusion

In this paper, we analyze the overlooked issues in prior-based pose estimation methods and empirically find that shape prior does not contribute to performance boosts. The keypoint is actually the deformation process, which builds correspondence between camera and world coordinates by reconstructing the object shape in the world space. Inspired by this, we design an implicit space transformation network (IST-Net) to transform the camera-space features to world space in an implicit manner. It builds the space correspondence without requiring 3D priors or ground-truth 3D models of target objects. Besides, we design two independent feature enhancers to further enhance the features from both camera- and world-space, which enriches them with more pose-sensitive information and geometrical constraints. Extensive experiments on the challenging benchmark show the effectiveness of our method in both efficiency and accuracy. We hope our investigation can provide new insights for future research in the community.

8. Acknowledgements

This work has been supported by Hong Kong Research Grant Council - Early Career Scheme (Grant No. 27209621), General Research Fund Scheme (Grant no. 17202422), and RGC matching fund scheme (RMGS). Part of the described research work is conducted in the JC STEM Lab of Robotics for Soft Materials funded by The Hong Kong Jockey Club Charities Trust.

References

- [1] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 3
- [2] Dengsheng Chen, Jun Li, and Kai Xu. Learning canonical shape space for category-level 6d object pose and size estimation. *CoRR*, abs/2001.09322, 2020. 1, 2, 6, 7
- [3] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2773–2782, 2021. 1, 3, 6, 8
- [4] Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, Linlin Shen, and Ales Leonardis. Fs-net: Fast shape-based network for category-level 6d object pose estimation with decoupled rotation mechanism. *CoRR*, abs/2103.07054, 2021. 1, 2, 6
- [5] Xu Chen, Zijian Dong, Jie Song, Andreas Geiger, and Otmar Hilliges. Category level object pose estimation via neural analysis-by-synthesis. *CoRR*, abs/2008.08145, 2020. 1, 2
- [6] Xinke Deng, Yu Xiang, Arsalan Mousavian, Clemens Eppner, Timothy Bretl, and Dieter Fox. Self-supervised 6d object pose estimation for robot manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3665–3671. IEEE, 2020. 1
- [7] Yan Di, Ruida Zhang, Zhiqiang Lou, Fabian Manhardt, Xiangyang Ji, Nassir Navab, and Federico Tombari. Gpv-pose: Category-level object pose estimation via geometry-guided point-wise voting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 6771–6781. IEEE, 2022. 1, 2, 6, 7
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [9] Zhaoxin Fan, Zhengbo Song, Jian Xu, Zhicheng Wang, Kejian Wu, Hongyan Liu, and Jun He. Acr-pose: Adversarial canonical representation reconstruction network for category level 6d object pose estimation. *arXiv preprint arXiv:2111.10524*, 2021. 3
- [10] Yang Fu and Xiaolong Wang. Category-level 6d object pose estimation in the wild: A semi-supervised learning approach and a new dataset. *arXiv preprint arXiv:2206.15436*, 2022. 1, 2, 6, 7
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [13] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11632–11641, 2020. 1
- [14] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6423–6432, 2021. 1
- [15] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Proceedings of the IEEE international conference on computer vision*, pages 1521–1529, 2017. 1
- [16] Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In *International conference on machine learning*, pages 1945–1954. PMLR, 2017. 2
- [17] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 574–591. Springer, 2020. 1
- [18] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17545–17555, 2023. 1
- [19] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8500–8509, 2022. 1
- [20] Xin Lai, Yuhui Yuan, Ruihang Chu, Yukang Chen, Han Hu, and Jiaya Jia. Mask-attention-free transformer for 3d instance segmentation. In *ICCV*, 2023. 1
- [21] Taeyeop Lee, Byeong-Uk Lee, Inkyu Shin, Jaesung Choe, Ukcheol Shin, In So Kweon, and Kuk-Jin Yoon. Uda-coppe: unsupervised domain adaptation for category-level object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14891–14900, 2022. 1
- [22] Yiming Li, Tao Kong, Ruihang Chu, Yifeng Li, Peng Wang, and Lei Li. Simultaneous semantic and collision learning for 6-dof grasp pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3571–3578. IEEE, 2021. 1
- [23] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–698, 2018. 1
- [24] Haitao Lin, Zichang Liu, Chilam Cheang, Lingwei Zhang, Yanwei Fu, and Xiangyang Xue. Donet: Learning category-level 6d object pose and size estimation from depth observation. *CoRR*, abs/2106.14193, 2021. 6
- [25] Jiehong Lin, Zewei Wei, Changxing Ding, and Kui Jia. Category-level 6d object pose and size estimation using self-supervised deep prior deformation networks. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 19–34. Springer, 2022. 1, 2, 3, 4, 6, 7, 9

- [26] Jiehong Lin, Zewei Wei, Zhihao Li, Songcen Xu, Kui Jia, and Yuanqing Li. Dualposenet: Category-level 6d object pose and size estimation using dual pose network with refined learning of pose consistency. *CoRR*, abs/2103.06526, 2021. [1](#), [2](#), [6](#), [7](#)
- [27] Chao Liu, Shuai Yu, Min Yu, Baole Wei, Boquan Li, Gang Li, and Weiqing Huang. Adaptive smooth l1 loss: A better way to regress scene texts with extreme aspect ratios. In *2021 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–7. IEEE, 2021. [5](#)
- [28] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2901–2910, 2019. [1](#)
- [29] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 55–64, 2020. [1](#)
- [30] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019. [1](#)
- [31] Giorgia Pitteri, Michaël Ramamonjisoa, Slobodan Ilic, and Vincent Lepetit. On object symmetries and 6d pose estimation from images. In *2019 International conference on 3D vision (3DV)*, pages 614–622. IEEE, 2019. [6](#)
- [32] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *CoRR*, abs/1706.02413, 2017. [6](#)
- [33] Caner Sahin and Tae-Kyun Kim. Category-level 6d object pose recovery in depth images. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. [2](#)
- [34] Yongzhi Su, Jason Rambach, Nareg Minaskan, Paul Lesur, Alain Pagani, and Didier Stricker. Deep multi-state object pose estimation for augmented reality assembly. In *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pages 222–227. IEEE, 2019. [1](#)
- [35] Meng Tian, Marcelo H Ang, and Gim Hee Lee. Shape prior deformation for categorical 6d object pose and size estimation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 530–546. Springer, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)
- [36] S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):376–380, 1991. [3](#), [8](#)
- [37] Chen Wang, Roberto Martín-Martín, Danfei Xu, Jun Lv, Cewu Lu, Li Fei-Fei, Silvio Savarese, and Yuke Zhu. 6-pack: Category-level 6d pose tracker with anchor-based keypoints. *CoRR*, abs/1910.10750, 2019. [2](#)
- [38] Chen Wang, Roberto Martín-Martín, Danfei Xu, Jun Lv, Cewu Lu, Li Fei-Fei, Silvio Savarese, and Yuke Zhu. 6-pack: Category-level 6d pose tracker with anchor-based keypoints. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10059–10066. IEEE, 2020. [1](#)
- [39] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019. [1](#), [2](#), [6](#), [7](#)
- [40] Jiase Wang, Kai Chen, and Qi Dou. Category-level 6d object pose estimation via cascaded relation and recurrent reconstruction networks. *CoRR*, abs/2108.08755, 2021. [1](#), [3](#), [6](#)
- [41] Yijia Weng, He Wang, Qiang Zhou, Yuzhe Qin, Yueqi Duan, Qingnan Fan, Baoquan Chen, Hao Su, and Leonidas J Guibas. Captra: Category-level pose tracking for rigid and articulated objects from point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13209–13218, 2021. [2](#)
- [42] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. [1](#)
- [43] Cheng Zhang, Zhaopeng Cui, Yinda Zhang, Bing Zeng, Marc Pollefeys, and Shuaicheng Liu. Holistic 3d scene understanding from a single image with implicit representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8833–8842, 2021. [1](#)
- [44] Kaifeng Zhang, Yang Fu, Shubhankar Borse, Hong Cai, Fatih Porikli, and Xiaolong Wang. Self-supervised geometric correspondence for category-level 6d object pose estimation in the wild. *arXiv preprint arXiv:2210.07199*, 2022. [2](#), [3](#), [7](#)
- [45] Ruida Zhang, Yan Di, Zhiqiang Lou, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Rbp-pose: Residual bounding box projection for category-level pose estimation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 655–672. Springer, 2022. [1](#), [3](#), [6](#)
- [46] Zhicheng Zhang, Song Chen, Zichuan Wang, and Jufeng Yang. Planeseg: Building a plug-in for boosting planar region segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. [1](#)
- [47] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [6](#)