

# StereoPose: Category-Level 6D Transparent Object Pose Estimation from Stereo Images via Back-View NOCS

Kai Chen<sup>1</sup>, Stephen James<sup>2,3</sup>, Congying Sui<sup>4</sup>, Yun-Hui Liu<sup>1</sup>, Pieter Abbeel<sup>2</sup> and Qi Dou<sup>1</sup>

**Abstract**—Most existing methods for category-level pose estimation rely on object point clouds. However, when considering transparent objects, depth cameras are usually not able to capture high-quality data, resulting in point clouds with severe artifacts. Without a complete point cloud, existing methods are not applicable to challenging transparent objects. To tackle this problem, we present StereoPose, a novel stereo image based framework for category-level object pose estimation, ideally suited for transparent objects. For a robust estimation from pure stereo images, we develop a pipeline that decouples category-level pose estimation into object size estimation, initial pose estimation, and pose refinement. StereoPose then estimates object pose based on representation in the normalized object coordinate space (NOCS). To address the issue of image content aliasing, we further define a back-view NOCS map for the transparent object. The back-view NOCS aims to reduce the network learning ambiguity caused by content aliasing, and leverage informative cues on the back of the transparent object for more accurate pose estimation. To further improve the performance of the stereo framework, StereoPose is equipped with a parallax attention module for stereo feature fusion and an epipolar loss for improving the stereo-view consistency of network predictions. Extensive experiments on the public TOD dataset demonstrate the superiority of the proposed StereoPose framework for category-level 6D transparent object pose estimation. Code and demos will be available on the project homepage: [www.cse.cuhk.edu.hk/~kaichen/stereopose.html](http://www.cse.cuhk.edu.hk/~kaichen/stereopose.html).

## I. INTRODUCTION

Transparent objects are common in both daily life and industry. Parsing their poses before grasping or manipulating them using robots is of great importance for ensuring accuracy and safety. Because of the generalization ability to novel objects, category-level object pose estimation [1]–[4] has received increasing attention recently. However, this problem is especially challenging for transparent objects. First of all, transparent objects usually have specular and reflective materials. These factors make acquiring high-quality depth maps of transparent objects challenging with commonly used depth sensors. As shown in Fig. 1 (a), their depth maps often exhibit severe artifacts, such as noise and large missing areas. Unfortunately, existing state-of-the-art methods [3], [5] rely on high-quality depth maps and point clouds to handle intra-class shape variation of different objects for accurate object pose estimation. The point cloud artifacts caused by the transparent material are hard to mitigate even with recent

This work was supported by InnoHK of the Government of Hong Kong via the Hong Kong Centre for Logistics Robotics.

<sup>1</sup> K. Chen, Y. H. Liu and Q. Dou are with The Chinese University of Hong Kong. <sup>2</sup> S. James and P. Abbeel are with University of California, Berkeley. <sup>3</sup> S. James is also with Dyson Robot Learning Lab. <sup>4</sup> C. Sui is with Hong Kong Centre for Logistics Robotics.

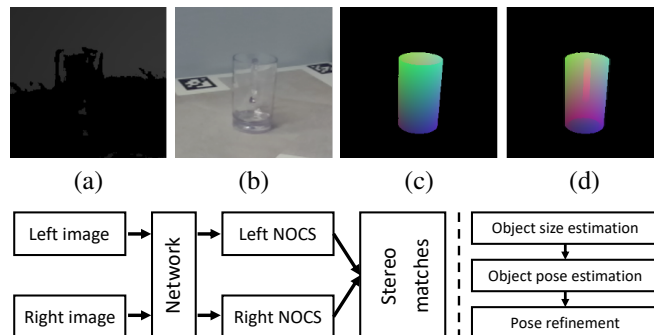


Fig. 1. Challenges of category-level transparent object pose estimation and an overview of the proposed StereoPose framework. (a) The depth map of a transparent mug with severe artifacts. (b) The RGB image patch of a transparent mug with content aliasing. (c) The conventional front-view NOCS map. (d) The proposed back-view NOCS map.

techniques such as depth completion and reconstruction [6]–[9]. Without a means to guarantee point cloud quality, these point-cloud-based pose estimation methods perform poorly on transparent objects.

While depth data fail to fully capture transparent objects, RGB data can often exhibit severe content aliasing caused by the transparent material. As shown in Fig. 1 (b), the handle of the mug could be observed on the image, despite being on the back of the mug. The mixed image content brings challenges to predicting dense correspondences of image pixels for object pose estimation. It is because conventional coordinate representations (e.g., instance-level voting field [10] and category-level NOCS map [1]) used on non-transparent objects only describe dense correspondences for the front-view of the object. As shown in Fig. 1 (b) and (c), the RGB image and the conventional front-view NOCS map are highly inconsistent in appearance. Learning this NOCS map from the aliased image has a relatively large ambiguity, which decreases the quality of the predicted NOCS map and makes the category-level pose estimation results not accurate.

To address the above issues, we present StereoPose, a novel stereo image framework for category-level 6D transparent object pose estimation. Instead of explicitly using the object point cloud, StereoPose exploits stereo images to implicitly model the object shape information, which is important for category-level object pose estimation. As shown in Fig. 1, StereoPose estimates NOCS maps for each view of the stereo images. An individual NOCS map provides intra-view correspondences between the camera and the object frame, while the left and right NOCS maps provide inter-view correspondences for the stereo views. Based on these correspondences, StereoPose decouples the estimation of category-level object pose into object size estimation,

initial pose estimation, and pose refinement, which achieves a robust object pose estimation without using the depth map.

Due to the transparent property, the back side of the object can always be observed from the image. Different from opaque object pose estimation, some object structures on the back (e.g., bottle rim or mug handle) can also be informative cues for determining object pose but are neglected in existing methods. To tackle this issue, as shown in Fig. 1 (d), we define a back-view NOCS map for the transparent object. Given an aliased image, we jointly estimate the conventional front-view NOCS map and the back-view NOCS map from the image. Jointly estimating two NOCS maps reduces the learning ambiguity caused by the image content aliasing, which helps to produce a NOCS map more accurately. Moreover, the back-view NOCS map captures informative cues on the back of the object, which also reduces the pose estimation ambiguity caused by object self-occlusion and further improves the pose accuracy for transparent objects. To make full use of the stereo features, we further propose a parallax attention module to fuse stereo image features through the guidance of parallax information. An epipolar geometry loss is applied to constrain the consistency between NOCS maps of left and right views. We summarize our main contributions as follows:

- We propose StereoPose, a novel stereo image based framework for category-level transparent object pose estimation. StereoPose is equipped with a new parallax attention module for stereo image feature fusion and an elaborated pipeline for robust object pose estimation with pure stereo images.
- We define the back-view NOCS map for the transparent objects. It reduces the negative effect of image content aliasing on transparent object pose estimation. It also captures informative structures on the back of the object for more accurate object pose estimation.
- We conduct extensive experiments on the public TOD dataset, which covers three typical categories of transparent objects with about 36K stereo images. StereoPose achieves dramatic performance improvements over other existing methods for category-level 6D transparent object pose estimation.

## II. RELATED WORKS

### A. Transparent Object Pose Estimation

Transparent objects [11] often exhibit sparkling texture, refractive and reflective material [12]. These factors result in a lot of artifacts on their image appearance and point cloud geometry, which makes accurate pose estimation for transparent objects very challenging [13]. Early methods extract edge information from images for transparent object pose estimation. Phillips et al. [14] develop a heuristic method to analyze the shape of the extracted edge curves for pose estimation of rotationally symmetric transparent objects. Lysenkov et al. [15], [16] extract object silhouettes from the edge information. They assume that the CAD model of the transparent object is available and then estimate object pose

by adjusting the pose hypothesis to fit the silhouettes. However, the 2D edge information is ambiguous in determining a complete 6D object pose, which limits the pose accuracy. In order to use the 3D geometry feature for more accurate pose estimation, recent methods [9], [17]–[19] do much effort to recover accurate geometry of transparent objects from their defective RGB-D data. Sajjan et al. [6] present a multi-task learning network to infer surface normal and object boundary of transparent objects for depth refinement. Zhu et al. [7] complete the missing depth with an implicit neural representation. Ichnowski et al. [12] use neural radiance fields to estimate the geometry of transparent objects. Xu et al. [20] develop a two-stage network to jointly refine the depth map and perform pose estimation for transparent objects. However, training an additional network for depth refinement is costly. The refined depth cannot recover the object shape perfectly, and the residual shape distortion would affect the final pose accuracy. In contrast, our method achieves accurate object pose estimation with stereo images.

### B. Category-Level Object Pose Estimation

Category-level object pose estimation aims to predict the pose of novel objects. The object intra-class shape variation [21], [22] is the main challenge of applying the network to novel objects for accurate pose estimation. In this regard, previous approaches [23], [24] extract shape-invariant features from depth map for category-level pose estimation. Wang et al. [1] then present the NOCS representation, which represents objects in a normalized canonical space to reduce the shape difference for object pose estimation. Furthermore, Tian et al. [25] describe a prior-based framework to explicitly reconstruct the novel object shape in the NOCS space. This prior-based framework is widely applied to recent methods with further improvements, such as recurrent NOCS model reconstruction [26], dynamic prior adaptation [5], and deep prior deformation [27]. Although the NOCS representation and the prior-based framework are quite popular and have achieved state-of-the-art performance for category-level object pose estimation, nearly all existing methods [28]–[31] highly rely on point cloud data and are therefore not suited for transparent object category-level pose estimation. Our method is a complement to existing methods and is tailored for transparent object category-level pose estimation.

## III. METHODOLOGY

In this section, we present the proposed StereoPose for category-level transparent object pose estimation. In Sec. III-A, we present the overall framework based on stereo images. A set of steps are elaborated for accurate pose estimation from pure stereo images. In Sec. III-B, we introduce the back-view NOCS map for transparent objects. We propose to jointly estimate the front-view and back-view NOCS maps to address the problem of image content aliasing for transparent object pose estimation. In Sec. III-C, we present a parallax attention module to integrate features from stereo images for more accurate object pose estimation with complementary features of stereo images.

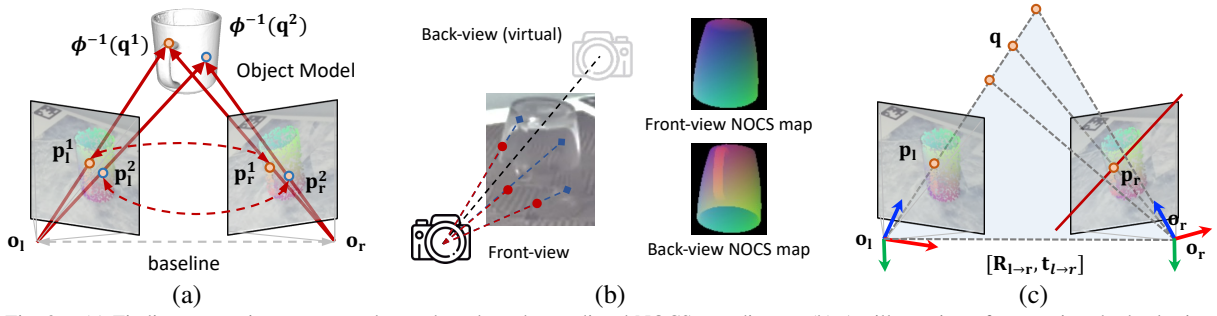


Fig. 2. (a) Finding cross-view correspondences based on the predicted NOCS coordinates. (b) An illustration of generating the back-view NOCS map. (c) Epipolar geometry loss for constraining the consistency of left-view and right-view NOCS maps.

### A. Overview of StereoPose

Different from existing point-cloud-based methods, our proposed StereoPose leverages stereo image features of the transparent object to account for the inherent shape variation of novel objects and predicts NOCS maps to construct dense correspondences for category-level object pose estimation. Following [25], StereoPose is developed based on a similar prior-based framework. Let  $\mathbf{F}_l, \mathbf{F}_r \in \mathbb{R}^{H \times W \times C}$  be stereo features extracted from stereo images,  $\mathcal{P} \in \mathbb{R}^{N \times 3}$  be the average prior model for the corresponding object category. StereoPose predicts a deformation field  $\mathcal{D} \in \mathbb{R}^{N \times 3}$  to reconstruct the shape of the observed transparent object by  $\mathcal{P}' = \mathcal{P} + \mathcal{D}$ . Subsequently, StereoPose predicts two matching matrices  $\mathcal{A}_l, \mathcal{A}_r \in \mathbb{R}^{M \times N}$  to predict NOCS coordinates for  $M$  sampled image pixels of the target transparent object as  $\mathcal{M}_{(l,r)} = \mathcal{A}_{(l,r)} \times \mathcal{P}'$ .

For one sampled image pixel  $\mathbf{p} \in \mathbb{R}^2$ , its NOCS coordinates on the corresponding NOCS map could be denoted as  $\mathbf{q} \in \mathbb{R}^3$ . According to the definition of NOCS in [1],  $\mathbf{p}$  and  $\mathbf{q}$  should satisfy:

$$\begin{pmatrix} \mathbf{p} \\ 1 \end{pmatrix} = \frac{1}{Z} \mathbf{K} \begin{bmatrix} s\mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \begin{pmatrix} \mathbf{q} \\ 1 \end{pmatrix}, \quad (1)$$

where  $\mathbf{K}$  denotes the camera intrinsic matrix,  $s$  denotes the isotropic scale factor related to the object size, and  $[\mathbf{R}, \mathbf{t}]$  denotes the transformation matrix corresponding to the 6D object pose. Given a calibrated camera with known intrinsic parameters, category-level object pose estimation aims to recover  $s$  and  $[\mathbf{R}, \mathbf{t}]$  simultaneously. Conventional methods first recover the object point cloud in the camera frame using a depth map, and then jointly estimate  $s, \mathbf{R}$  and  $\mathbf{t}$  by 3D-3D registration. However, it is not applicable to transparent objects with low-quality depth maps.

StereoPose therefore proposes to decouple the estimation of  $s$  and  $[\mathbf{R}, \mathbf{t}]$ . An elaborated pipeline then is developed for a robust and accurate category-level pose estimation of transparent objects with pure stereo images. Firstly, we estimate  $s$  based on correspondences between  $\mathcal{M}_l$  and  $\mathcal{M}_r$ . As shown in Fig. III (a), we propose to efficiently find the cross-view correspondences based on the predicted NOCS coordinates. Let  $\mathbf{p}_l, \mathbf{p}_r \in \mathbb{R}^2$  be two points on the left and right image. Intuitively,  $\mathbf{p}_l$  and  $\mathbf{p}_r$  are a pair of correspondences if they have the same NOCS coordinates. For a stereo camera with a known baseline length  $b$ , we can measure the depth value of the corresponding 3D point as  $d = \frac{f}{\|\mathbf{p}_l - \mathbf{p}_r\|} \times b$ . For arbitrary

two pairs of points, we can estimate one scale factor based on the ratio of the distance in the camera frame to the distance in the NOCS space. For a robust estimation, we extract all matched points from  $\mathcal{M}_l$  and  $\mathcal{M}_r$  and average the estimation results across all pairs.

Given the estimated  $s$ , StereoPose further solves the remaining  $\mathbf{R}$  and  $\mathbf{t}$  by the Perspective-n-Point (PnP) algorithm [32]. Our method directly leverages 2D-3D correspondences for object pose estimation. Using the RANSAC algorithm [33] would further improve the robustness of the estimation results. We also compare our proposed decoupled pose estimation pipeline with the conventional scheme, in which we first recover 3D positions for all matched points and then jointly estimate  $s, \mathbf{R}$  and  $\mathbf{t}$  by fitting a 3D-3D similarity transformation. Results indicate that our proposed method could achieve higher performance for category-level transparent object pose estimation. To further improve the pose accuracy, similar to [34], we propose to rescale  $\mathbf{t}$ , which aims to align the average object depth corresponding to the estimated object pose to the average depth value recovered from matched points of  $\mathcal{M}_l$  and  $\mathcal{M}_r$ .

### B. Back-View NOCS Map

Though our proposed StereoPose manages to estimate pose for transparent objects from pure stereo images with an elaborated pose estimation pipeline, the problem of image content aliasing has not been addressed for transparent objects. This unresolved problem would significantly affect the pose accuracy. It is because (1) Estimating a conventional front-view NOCS map from the aliased image is of a large ambiguity, which affects the accuracy of the predicted NOCS map; (2) Though the back side could be observed on the RGB image of a transparent object, the conventional front-view NOCS map fails to leverage this information to improve the pose estimation accuracy.

In order to tackle these problems, we define the back-view NOCS map for transparent objects. Fig. III (b) illustrates the generation process. Let  $\mathbf{o} \in \mathbb{R}^3$  be the camera center, and  $\mathbf{r}$  be a ray marching from  $\mathbf{o}$ . Let  $\mathcal{C}$  be the point set that  $\mathbf{r}$  intersects with the object model. The conventional front-view NOCS could be denoted as:

$$\mathbf{c}_{front} = \phi(\arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{o}\|), \quad (2)$$

where  $\phi(\cdot)$  returns the NOCS coordinates for the corresponding point on the object model. For the back view, instead of really moving the camera to the back side for capturing a

back-view NOCS map, we take the NOCS coordinates of the object point along  $\mathbf{r}$  that is the farthest from the front-view camera as the corresponding back-view NOCS coordinates. Specifically, the back-view NOCS map is defined as:

$$\mathbf{c}_{back} = \phi(\arg \max_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{o}\|). \quad (3)$$

According to our definition, the back-view NOCS map would be ideally aligned with the front-view NOCS map as well as the RGB image. It facilitates estimating the back-view NOCS map from the RGB image. Fig. III (b) presents an example for the generated back-view NOCS map. Our proposed back-view NOCS map exhibits the informative structure cues on the back of the object. We then distill the front-view feature and the back-view feature from the aliased image to jointly estimate the front-view and the back-view NOCS maps. It decreases the ambiguity of network learning and helps to improve the NOCS map accuracy. In addition, the proposed StereoPose adopts the PnP solver for object pose estimation. It is based on 2D-3D correspondences, which means that the predicted back-view NOCS map could be directly applied to the PnP scheme for object pose estimation. In other words, the introduction of the back-view NOCS map excavates the structure information on the back of the transparent object for more accurate pose estimation.

### C. Parallax Attention for Stereo Feature Fusion

Stereo images provide informative cross-view features for the target object. Properly fusing them enhances the image feature for representing the shape of the transparent object for more accurate category-level object pose estimation. The parallax<sup>1</sup> magnitude of an object exhibited on stereo images would vary with the shape and depth of the object. In other words, the object parallax implicitly indicates this useful information, which could be used to guide the stereo feature fusion for improving the object pose estimation performance.

In this regard, we propose a parallax attention module for stereo feature fusion. Given the stereo image features  $\mathbf{F}_l, \mathbf{F}_r \in \mathbb{R}^{H \times W \times C}$ , we compute the pixel-wise feature correlation between  $\mathbf{F}_l$  and  $\mathbf{F}_r$  to softly model the parallax information of an object with respect to the camera. We refer to this pixel-wise correlation as a parallax attention map and use it for stereo feature fusion. Specifically, we leverage the disentangled non-local block [35], [36] to compute the parallax attention map. At first, two whitening layers are used to normalize  $\mathbf{F}_l$  and  $\mathbf{F}_r$  to  $\mathbf{F}'_l$  and  $\mathbf{F}'_r$ . After that, the attention map could be computed as:

$$A_{r \rightarrow l} = \sigma(\mathbf{F}'_l \otimes \mathbf{F}'_r{}^\top), A_{l \rightarrow r} = \sigma(\mathbf{F}'_r \otimes \mathbf{F}'_l{}^\top), \quad (4)$$

where  $\otimes$  denotes the batch-wise matrix multiplication,  $\sigma(\cdot)$  is the operation of softmax normalization along the last dimension. The resulted  $A_{r \rightarrow l}, A_{l \rightarrow r} \in \mathbb{R}^{H \times W \times W}$  are parallax attention maps from the right view to the left view and from the left view to the right view, respectively. To fuse the right-view feature with the left-view one, we convert  $\mathbf{F}_r$  to the left-view based on  $A_{r \rightarrow l}$ :

$$\mathbf{F}_{r \rightarrow l} = A_{r \rightarrow l} \otimes \mathbf{F}_r. \quad (5)$$

<sup>1</sup>Parallax is a position displacement of matched points on stereo images.

We can follow a similar way to obtain  $\mathbf{F}_{l \rightarrow r}$ . The converted features are fused with the original feature by an MLP layer for subsequent NOCS map prediction.

### D. Epipolar Geometry Loss Function

Since the proposed StereoPose would separately predict NOCS maps for the left view and the right view of the stereo camera, to constrain the inter-view consistency of left-view and right-view NOCS maps, we propose an epipolar geometry loss in this section. Let  $\mathbf{K}_l$  and  $\mathbf{K}_r$  be the intrinsic matrices of the left camera and the right camera respectively. Let  $\mathbf{R}_{l \rightarrow r}$  and  $\mathbf{t}_{l \rightarrow r}$  be the relative rotation and translation from the left camera to the right camera,  $\mathcal{M}_l$  and  $\mathcal{M}_r$  be the left-view NOCS map and the right-view NOCS map. As shown in Fig. III (c), based on the epipolar geometry [37], if we follow the scheme mentioned in Sec. III-A to find a pair of matched points ( $\mathbf{p}_l, \mathbf{p}_r$ ) that have the same coordinates in the NOCS space,  $\mathbf{p}_l$  and  $\mathbf{p}_r$  should satisfy:

$$\hat{\mathbf{p}}_l \mathbf{K}_l^{-\top} [\mathbf{t}_{l \rightarrow r}]_{\times} \mathbf{R}_{l \rightarrow r} \mathbf{K}_r^{-1} \hat{\mathbf{p}}_r{}^\top = 0, \quad (6)$$

where  $\hat{\mathbf{p}}_l, \hat{\mathbf{p}}_r \in \mathbb{R}^{1 \times 3}$  are homogeneous image pixel coordinates,  $[\cdot]_{\times}$  denotes computing the skew symmetric matrix of  $\mathbf{t}_{l \rightarrow r}$ . The epipolar geometry loss  $l_{ep}$  therefore could be directly defined as  $l_{ep} = \hat{\mathbf{p}}_l \mathbf{K}_l^{-\top} [\mathbf{t}_{l \rightarrow r}]_{\times} \mathbf{R}_{l \rightarrow r} \mathbf{K}_r^{-1} \hat{\mathbf{p}}_r{}^\top$ .

To train the StereoPose model, apart from  $l_{ep}$ , similar to [25], [26], we use the complete loss function defined as:

$$L = \lambda_1 l_{ep} + \lambda_2 l_{cd} + \lambda_3 l_{nocs} + \lambda_4 l_a + \lambda_5 l_d, \quad (7)$$

in which  $l_{cd}$  is a reconstruction loss for computing the Chamfer Distance between  $\mathcal{P}'$  and  $\mathcal{P}_{gt}$  as:

$$l_{cd} = \sum_{x \in \mathcal{P}'} \min_{y \in \mathcal{P}_{gt}} \|x - y\|_2^2 + \sum_{y \in \mathcal{P}_{gt}} \min_{x \in \mathcal{P}'} \|x - y\|_2^2. \quad (8)$$

The NOCS loss  $l_{nocs}$  is used to minimize the  $L_1$  distance between  $\mathcal{M}$  and  $\mathcal{M}_{gt}$ . In addition, we regularize the deformation field with  $l_d = \frac{1}{N} \sum_{\mathbf{d} \in \mathcal{D}} \|\mathbf{b}\|_2$ , and constrain the distribution of  $\mathcal{A}$  by  $l_a = \frac{1}{M} \sum_i \sum_j -a_{i,j} \log a_{i,j}$ .

## IV. EXPERIMENTS

### A. Datasets and Evaluation Metrics

We evaluate our proposed StereoPose on the TOD dataset [38]. This dataset contains three categories of transparent objects: bottle (3 instances), mug (7 instances), and cup (2 instances), with 33989 stereo image pairs for 12 different object instances in 10 different environment backgrounds. We follow the setting in [38] to conduct category-level experiments across three different category splits: ‘mug’, ‘bottle’, and ‘bottle and cup’. In each split, we train StereoPose on different instances and evaluate the model on novel instances that are not used during training.

We quantitatively evaluate the performance of category-level object pose estimation following the widely used metrics in [5], [25], [39]. Specifically, the metric of 3D IoU computes the overlapping ratio of 3D bounding boxes corresponding to the ground-truth pose and the predicted pose. We report the mAP with respect to 3D<sub>25</sub> and 3D<sub>50</sub> in our experiments. In addition, we directly measure the rotation

TABLE I

QUANTITATIVE COMPARISON WITH STATE-OF-THE-ART METHODS ON TOD DATASET. THE LARGER THE METRIC, THE HIGHER THE POSE ACCURACY.

Method	Bottle				Bottle and Cup				Mug			
	$3D_{25}$	$3D_{50}$	$10^\circ 5cm$	$10^\circ 10cm$	$3D_{25}$	$3D_{50}$	$10^\circ 5cm$	$10^\circ 10cm$	$3D_{25}$	$3D_{50}$	$10^\circ 5cm$	$10^\circ 10cm$
SPD [25]	44.3	7.4	11.5	17.8	41.8	6.5	7.8	12.2	63.6	19.7	2.3	4.2
SGPA [5]	46.9	9.6	13.3	22.5	48.1	12.5	9.5	14.6	64.4	19.6	2.8	5.1
KeyPose [38]	-	-	52.7	62.3	-	-	30.2	35.1	-	-	24.6	25.1
StereoPose	<b>85.4</b>	<b>22.2</b>	<b>57.8</b>	<b>70.3</b>	<b>85.7</b>	<b>17.8</b>	<b>34.7</b>	<b>41.3</b>	<b>97.9</b>	<b>77.4</b>	<b>34.4</b>	<b>38.2</b>

and translation error between the ground-truth pose and the predicted pose. We report mAP with respect to  $10^\circ 5cm$  and  $10^\circ 10cm$  in our experiments.

### B. Implementation Details

StereoPose uses a ResNet-18 [40] with PSP network [41] to extract features for stereo images. Similar to most category-level object pose estimation methods, the target object region is first cropped based on the segmentation result. We train a TransLab [42] model on Trans10K dataset and then directly use the trained model for the transparent object segmentation of the TOD dataset. The cropped object image patch is resized to  $224 \times 224$  and then fed into the backbone network for image feature extraction. We set  $M = N = 1024$  when reconstructing the NOCS object model and predicting the NOCS map for the observed image under the prior-based framework. For the hyperparameters of the loss function, we empirically use  $\lambda_1 = 0.01$ ,  $\lambda_2 = 5.0$ ,  $\lambda_3 = 1.0$ ,  $\lambda_4 = 0.0001$ , and  $\lambda_5 = 0.01$ .

### C. Comparison with State-of-the-Art Methods

Since there are few available methods designed for stereo category-level transparent object pose estimation, we adapt existing methods to transparent objects and compare our proposed framework with three state-of-the-art methods. On the one hand, we compare our method with the paradigm of ‘depth completion + opaque object pose estimation’. We first train TransCG [9] on the TOD dataset and use it to improve the depth map. After that, we use the refined depth map for category-level object pose estimation. We choose two state-of-the-art methods (SPD [25] and SGPA [5]) for comparisons. On the other hand, we compare our method with one keypoint-based approach (KeyPose [38]), in which stereo images are used to predict category-level object key points. These key points are then used to recover the complete object pose. Noting that KeyPose does not need depth/point cloud, thus could be applied to transparent object pose estimation without any modification.

Tab. I presents the comparative results. StereoPose significantly outperforms other competing methods in all evaluation metrics. First, StereoPose achieves higher accuracy than ‘depth completion + opaque object pose estimation’. Although the depth accuracy indeed gets apparent improvement after the refinement of TransCG (as shown in Tab. II), the point cloud recovered from the refined depth map still exhibits severe distortion, which cannot represent the shape of the transparent object precisely. The SPD and SGPA are easily affected by the point cloud artifacts. As shown in Fig. 3, without a good point cloud to represent the transparent

TABLE II

EVALUATION OF THE DEPTH ACCURACY BEFORE AND AFTER TRANSCG REFINEMENT. FOR BOTH ‘RMSE’ AND ‘MAE’, THE SMALLER THE VALUE, THE HIGHER THE ACCURACY.

	Bottle		Cup		Mug	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
Original	0.506	0.450	0.516	0.466	0.564	0.515
Refined	<b>0.071</b>	<b>0.065</b>	<b>0.075</b>	<b>0.068</b>	<b>0.080</b>	<b>0.072</b>




Fig. 3. Qualitative comparisons with ‘depth completion + opaque object pose estimation’. (a) Ground-truth point cloud. (b) Recovered point cloud by TransCG. (c) Results of SGPA. (d) Results of StereoPose.

object shape, their pose estimation accuracy is inferior to StereoPose by a large margin.

In addition, StereoPose also outperforms KeyPose, a keypoint-based method for object pose estimation. As shown in Fig. IV-C, although the 2D positions of the predicted object key points on stereo images are visually accurate, the recovered object pose still exhibits a relatively large error in both rotation and translation. It is because the sparse object key points could provide only very limited correspondence information between the object frame and the camera frame. As a result, the recovered object pose is sensitive to the position error of predicted object key points. In contrast, the proposed StereoPose uses dense correspondences by predicting the NOCS map for robust object pose estimation. Moreover, the introduction of the back-view NOCS map, parallax-aware stereo feature fusion, and epipolar geometry loss function further improve the quality of the predicted NOCS map for a high pose accuracy.

### D. Ablation Studies

**Back-view NOCS Map.** We remove the back-view NOCS branch from the framework and evaluate the network on the same testing dataset. Tab. III presents the experiment results. Without the back-view NOCS map, all metrics suffer a severe drop. As shown in Fig. IV-D, when the handle is on the back of the mug, the model without using the back-view NOCS map cannot localize the handle accurately and produces a large rotation error. In contrast, using the back-view NOCS map could leverage this informative structure cue on the back




Fig. 4. **Qualitative comparison with KeyPose.** (a) Predicted left-view keypoints. (b) Predicted right-view keypoints. (c) Pose results computed from key points. (d) Pose results w/o the proposed decoupled estimation scheme. (e) Pose results of the proposed method. (f) Ground-truth.

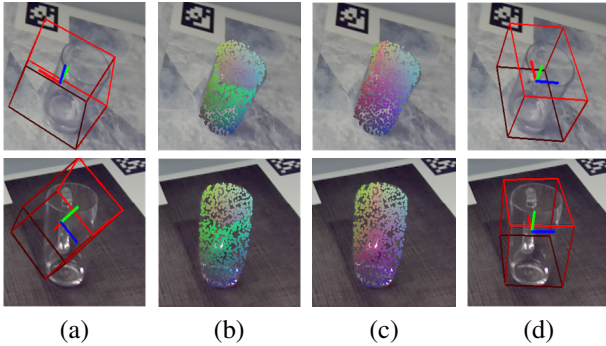


Fig. 5. **Effect of the proposed back-view NOCS map.** (a) Results w/o back-view NOCS map. (b) Predicted front-view NOCS map. (c) Predicted back-view NOCS map. (d) Results w/ back-view NOCS map.

of the object and obtain a very accurate pose estimation result even in this challenging scenario. These results demonstrate the effectiveness of our proposed back-view NOCS map for transparent object pose estimation.

**Stereo Feature Fusion.** We remove the stereo feature fusion module from the framework and evaluate the network on the same testing dataset. Note that we still use both left and right images for network training and only remove the cross-view feature interaction module. Therefore, networks with and without stereo feature fusion are trained on the same number of images. As shown in Tab. III, the pose accuracy of the network without the stereo feature fusion module is significantly inferior to the result of our complete model. These results demonstrate the effectiveness of our proposed stereo fusion module in integrating the cross-view features for more accurate object pose estimation.

**Epipolar Geometry Loss.** We remove  $l_{ep}$  when training the network and evaluate the model on the same testing dataset. Tab. III concludes the experiment results. Removing the epipolar loss hurts the pose accuracy in terms of nearly all evaluation metrics. Fig. IV-D further presents the qualitative comparison on the NOCS map. We can find that the epipolar loss acting as a regularization term is effective in improving the quality of the predicted NOCS map. On the one hand, it improves the consistency between the left-view NOCS and the right-view NOCS. On the other hand, the epipolar loss also helps to produce a more accurate NOCS map when compared with the ground-truth NOCS map. These results demonstrate the effectiveness of the proposed epipolar loss




Fig. 6. **Qualitative comparison of stereo NOCS maps predicted by models w/ and w/o the proposed epipolar loss.** (a) Ground-truth NOCS maps. (b) NOCS maps w/o  $l_{ep}$ . (c) NOCS maps w/  $l_{ep}$ . The top row is for the left view and the bottom row is for the right view.

for the stereo object pose estimation framework.

TABLE III

ABLATION STUDY ON CATEGORY OF MUG. 'A': BACK-VIEW NOCS MAP. 'B': PARALLAX ATTENTION FUSION. 'C': EPIPOLAR LOSS.

	A	B	C	$3D_{25}$	$3D_{50}$	$10^\circ 5cm$	$10^\circ 10cm$
1	-	-	-	88.1	46.5	16.8	25.4
2	✓	-	-	93.3	55.4	19.8	29.3
3	-	✓	-	89.9	48.5	18.7	27.9
4	-	-	✓	96.4	63.2	21.3	24.2
5	✓	✓	-	95.5	75.2	26.8	31.0
6	✓	-	✓	93.6	57.7	21.6	28.8
7	-	✓	✓	93.0	57.2	21.1	31.7
8	✓	✓	✓	<b>97.9</b>	<b>77.4</b>	<b>34.4</b>	<b>38.2</b>

## V. CONCLUSION

In this work, we propose StereoPose, a novel stereo framework for category-level transparent object pose estimation. StereoPose decouples category-level pose estimation into object size estimation, initial pose estimation, and pose refinement, achieving a robust and accurate pose estimation for transparent objects with pure stereo images. The back-view NOCS map is defined for transparent objects, which excavates the structure information on the back of the transparent object for more accurate pose estimation. To further improve the performance of the stereo framework, we propose a parallax attention module for stereo feature fusion. An epipolar geometry loss is put forward to improve the consistency of stereo-view NOCS maps. Extensive experiments on the public TOD dataset demonstrate the superiority of the proposed StereoPose for category-level 6D transparent object pose estimation, which shows huge potential for robotic applications related to transparent objects.

## REFERENCES

- [1] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. 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*, 2019, pp. 2642–2651.
- [2] X. Deng, J. Geng, T. Bretl, Y. Xiang, and D. Fox, "icaps: Iterative category-level object pose and shape estimation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1784–1791, 2022.
- [3] X. Liu, G. Wang, Y. Li, and X. Ji, "Catre: Iterative point clouds alignment for category-level object pose refinement," *arXiv preprint arXiv:2207.08082*, 2022.
- [4] M. Z. Irshad, S. Zakharov, R. Ambrus, T. Kollar, Z. Kira, and A. Gaidon, "Shapo: Implicit representations for multi-object shape, appearance, and pose optimization," *arXiv preprint arXiv:2207.13691*, 2022.
- [5] K. Chen and Q. Dou, "Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2773–2782.
- [6] S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, "Clear grasp: 3d shape estimation of transparent objects for manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 3634–3642.
- [7] L. Zhu, A. Mousavian, Y. Xiang, H. Mazhar, J. van Eenbergen, S. Debnath, and D. Fox, "Rgb-d local implicit function for depth completion of transparent objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4649–4658.
- [8] H. Xu, Y. R. Wang, S. Eppel, A. Aspuru-Guzik, F. Shkurti, and A. Garg, "Seeing glass: Joint point-cloud and depth completion for transparent objects," in *5th Annual Conference on Robot Learning*, 2021.
- [9] H. Fang, H.-S. Fang, S. Xu, and C. Lu, "Transcg: A large-scale real-world dataset for transparent object depth completion and a grasping baseline," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7383–7390, 2022.
- [10] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao, "Pvnet: Pixel-wise voting network for 6dof pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4561–4570.
- [11] X. Chen, H. Zhang, Z. Yu, A. Opipari, and O. C. Jenkins, "Clearpose: Large-scale transparent object dataset and benchmark," *arXiv preprint arXiv:2203.03890*, 2022.
- [12] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, "Dex-nerf: Using a neural radiance field to grasp transparent objects," in *5th Annual Conference on Robot Learning*, 2021.
- [13] Y. Zhu, J. Qiu, and B. Ren, "Transfusion: A novel slam method focused on transparent objects," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6019–6028.
- [14] C. J. Phillips, M. Lecce, and K. Daniilidis, "Seeing glassware: from edge detection to pose estimation and shape recovery," in *Robotics: Science and Systems*, vol. 3. Michigan, USA, 2016, p. 3.
- [15] I. Lysenkov, V. Eruhmov, and G. Bradski, "Recognition and pose estimation of rigid transparent objects with a kinect sensor," *Robotics*, vol. 273, no. 273-280, p. 2, 2013.
- [16] I. Lysenkov and V. Rabaud, "Pose estimation of rigid transparent objects in transparent clutter," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 162–169.
- [17] Y. Tang, J. Chen, Z. Yang, Z. Lin, Q. Li, and W. Liu, "Depthgrasp: Depth completion of transparent objects using self-attentive adversarial network with spectral residual for grasping," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5710–5716.
- [18] Q. Dai, J. Zhang, Q. Li, T. Wu, H. Dong, Z. Liu, P. Tan, and H. Wang, "Domain randomization-enhanced depth simulation and restoration for perceiving and grasping specular and transparent objects," *arXiv preprint arXiv:2208.03792*, 2022.
- [19] J. Jiang, G. Cao, T.-T. Do, and S. Luo, "A4t: Hierarchical affordance detection for transparent objects depth reconstruction and manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9826–9833, 2022.
- [20] C. Xu, J. Chen, M. Yao, J. Zhou, L. Zhang, and Y. Liu, "6dof pose estimation of transparent object from a single rgb-d image," *Sensors*, vol. 20, no. 23, p. 6790, 2020.
- [21] D. Chen, J. Li, Z. Wang, and K. Xu, "Learning canonical shape space for category-level 6d object pose and size estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 973–11 982.
- [22] X. Li, Y. Weng, L. Yi, L. J. Guibas, A. Abbott, S. Song, and H. Wang, "Leveraging se (3) equivariance for self-supervised category-level object pose estimation from point clouds," *Advances in Neural Information Processing Systems*, vol. 34, pp. 15 370–15 381, 2021.
- [23] C. Sahin and T.-K. Kim, "Category-level 6d object pose recovery in depth images," in *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018, pp. 0–0.
- [24] C. Sahin, G. Garcia-Hernando, J. Sock, and T.-K. Kim, "Instance-and category-level 6d object pose estimation," in *RGB-D Image Analysis and Processing*. Springer, 2019, pp. 243–265.
- [25] M. Tian, M. H. Ang, and G. H. Lee, "Shape prior deformation for categorical 6d object pose and size estimation," in *European Conference on Computer Vision*. Springer, 2020, pp. 530–546.
- [26] J. Wang, K. Chen, and Q. Dou, "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)*. IEEE, 2021, pp. 4807–4814.
- [27] J. Lin, Z. Wei, C. Ding, and K. Jia, "Category-level 6d object pose and size estimation using self-supervised deep prior deformation networks," *arXiv preprint arXiv:2207.05444*, 2022.
- [28] M. Z. Irshad, T. Kollar, M. Laskey, K. Stone, and Z. Kira, "Centersnap: Single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation," *arXiv preprint arXiv:2203.01929*, 2022.
- [29] Y. Di, R. Zhang, Z. Lou, F. Manhardt, X. Ji, N. Navab, and F. Tombari, "Gpv-pose: Category-level object pose estimation via geometry-guided point-wise voting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6781–6791.
- [30] H. Lin, Z. Liu, C. Cheang, Y. Fu, G. Guo, and X. Xue, "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*, 2022, pp. 6707–6717.
- [31] R. Zhang, Y. Di, Z. Lou, F. Manhardt, N. Navab, F. Tombari, and X. Ji, "Rbp-pose: Residual bounding box projection for category-level pose estimation," *arXiv preprint arXiv:2208.00237*, 2022.
- [32] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnp: An accurate o (n) solution to the pnp problem," *International journal of computer vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [33] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [34] R. L. Haugaard and A. G. Buch, "Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6749–6758.
- [35] M. Yin, Z. Yao, Y. Cao, X. Li, Z. Zhang, S. Lin, and H. Hu, "Disentangled non-local neural networks," in *European Conference on Computer Vision*. Springer, 2020, pp. 191–207.
- [36] Y. Wang, X. Ying, L. Wang, J. Yang, W. An, and Y. Guo, "Symmetric parallax attention for stereo image super-resolution," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 766–775.
- [37] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [38] X. Liu, R. Jonschkowski, A. Angelova, and K. Konolige, "Keypose: Multi-view 3d labeling and keypoint estimation for transparent objects," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 602–11 610.
- [39] Y. You, R. Shi, W. Wang, and C. Lu, "Cpfp: Towards robust category-level 9d pose estimation in the wild," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6866–6875.
- [40] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [41] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [42] E. Xie, W. Wang, W. Wang, M. Ding, C. Shen, and P. Luo, "Segmenting transparent objects in the wild," in *European conference on computer vision*. Springer, 2020, pp. 696–711.