



# TransNet: Category-Level Transparent Object Pose Estimation

Huijie Zhang<sup>(✉)</sup> , Anthony Opipari , Xiaotong Chen , Jiyue Zhu ,  
Zeren Yu , and Odest Chadwicke Jenkins 

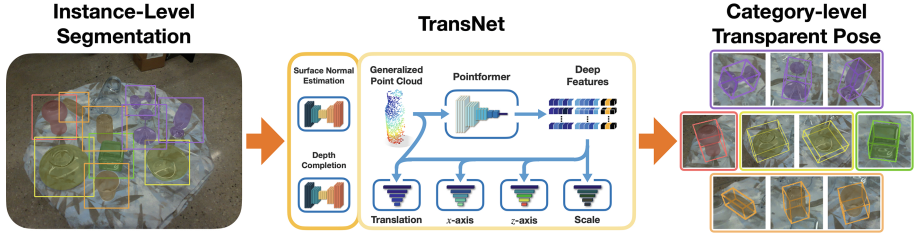
University of Michigan, Ann Arbor, MI 48109, USA  
{huijiezhang, topipari, cxt, cormacz, yuzeren, ocj}@umich.edu

**Abstract.** Transparent objects present multiple distinct challenges to visual perception systems. First, their lack of distinguishing visual features makes transparent objects harder to detect and localize than opaque objects. Even humans find certain transparent surfaces with little specular reflection or refraction, e.g. glass doors, difficult to perceive. A second challenge is that common depth sensors typically used for opaque object perception cannot obtain accurate depth measurements on transparent objects due to their unique reflective properties. Stemming from these challenges, we observe that transparent object instances within the same category (e.g. cups) look more similar to each other than to ordinary opaque objects of that same category. Given this observation, the present paper sets out to explore the possibility of category-level transparent object pose estimation rather than instance-level pose estimation. We propose *TransNet*, a two-stage pipeline that learns to estimate category-level transparent object pose using localized depth completion and surface normal estimation. TransNet is evaluated in terms of pose estimation accuracy on a recent, large-scale transparent object dataset and compared to a state-of-the-art category-level pose estimation approach. Results from this comparison demonstrate that TransNet achieves improved pose estimation accuracy on transparent objects and key findings from the included ablation studies suggest future directions for performance improvements. The project webpage is available at: <https://progress.eecs.umich.edu/projects/transnet/>.

**Keywords:** Transparent objects · Category-level object pose estimation · Depth completion · Surface normal estimation

## 1 Introduction

From glass doors and windows to kitchenware and all kinds of containers, transparent materials are prevalent throughout daily life. Thus, perceiving the pose (position and orientation) of transparent objects is a crucial capability for autonomous perception systems seeking to interact with their environment. However, transparent objects present unique perception challenges both in the RGB and depth domains. As shown in Fig. 2, for RGB, the color appearance



**Fig. 1.** Overview of TransNet, a pipeline for category-level transparent object pose estimation. Given instance-level segmentation masks as input, TransNet estimates the 6 degrees of freedom pose and scale for each transparent object in the image. Internally, TransNet uses surface normal estimation, depth completion, and a transformer-based architecture for accurate pose estimation despite noisy sensor data.

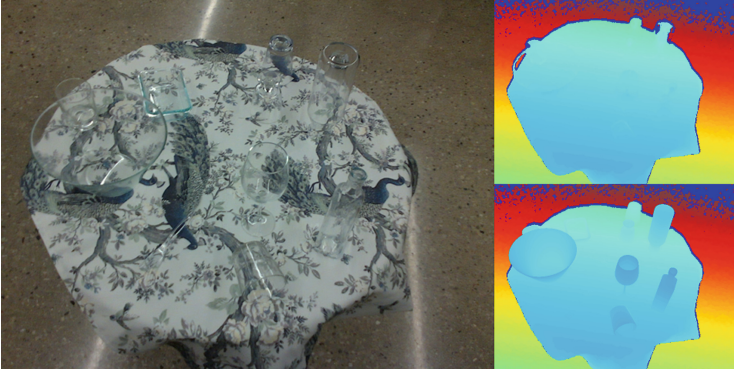
of transparent objects is highly dependent on the background, viewing angle, material, lighting condition, etc. due to light reflection and refraction effects. For depth, common commercially available depth sensors record mostly invalid or inaccurate depth values within the region of transparency. Such visual challenges, especially missing detection in the depth domain, pose severe problems for autonomous object manipulation and obstacle avoidance tasks. This paper sets out to address these problems by studying how category-level transparent object pose estimation may be achieved using end-to-end learning.

Recent works have shown promising results on grasping transparent objects by completing the missing depth values followed by the use of a geometry-based grasp engine [9, 12, 29], or transfer learning from RGB-based grasping neural networks [36]. For more advanced manipulation tasks such as rigid body pick-and-place or liquid pouring, geometry-based estimations, such as symmetrical axes, edges [27] or object poses [26], are required to model the manipulation trajectories. Instance-level transparent object poses could be estimated from keypoints on stereo RGB images [23, 24] or directly from a single RGB-D image [38] with support plane assumptions. Recently emerged large-scale transparent object datasets [6, 9, 23, 29, 39] pave the way for addressing the problem using deep learning.

In this work, we aim to extend the frontier of 3D transparent object perception with three primary contributions.

- First, we explore the importance of depth completion and surface normal estimation in transparent object pose estimation. Results from these studies indicate the relative importance of each modality and their analysis suggests promising directions for follow-on studies.
- Second, we introduce *TransNet*, a category-level pose estimation pipeline for transparent objects as illustrated in Fig. 1. It utilizes surface normal estimation, depth completion, and a transformer-based architecture to estimate transparent objects’ 6D poses and scales.

- Third, we demonstrate that TransNet outperforms a baseline that uses a state-of-the-art opaque object pose estimation approach [7] along with transparent object depth completion [9].



**Fig. 2.** Challenge for transparent object perception. Images are from Clearpose dataset [6]. The left is an RGB image. The top right is the raw depth image and the bottom right is the ground truth depth image.

## 2 Related Works

### 2.1 Transparent Object Visual Perception for Manipulation

Transparent objects need to be perceived before being manipulated. Lai *et al.* [18] and Khaing *et al.* [16] developed CNN models to detect transparent objects from RGB images. Xie *et al.* [37] proposed a deep segmentation model that achieved state-of-the-art segmentation accuracy. ClearGrasp [29] employed depth completion for use with pose estimation on robotic grasping tasks, where they trained three DeepLabv3+ [4] models to perform image segmentation, surface normal estimation, and boundary segmentation. Follow-on studies developed different approaches for depth completion, including implicit functions [47], NeRF features [12], combined point cloud and depth features [39], adversarial learning [30], multi-view geometry [1], and RGB image completion [9]. Without completing depth, Weng *et al.* [36] proposed a method to transfer the learned grasping policy from the RGB domain to the raw sensor depth domain. For instance-level pose estimation, Xu *et al.* [38] utilized segmentation, surface normal, and image coordinate UV-map as input to a network similar to [32] that can estimate 6 DOF object pose. Keypose [24] was proposed to estimate 2D keypoints and regress object poses from stereo images using triangulation. For other special sensors, Xu *et al.* [40] used light-field images to do segmentation using a graph-cut-based approach. Kalra *et al.* [15] trained Mask R-CNN [11] using polarization images as input to outperform the baseline that was trained on only RGB

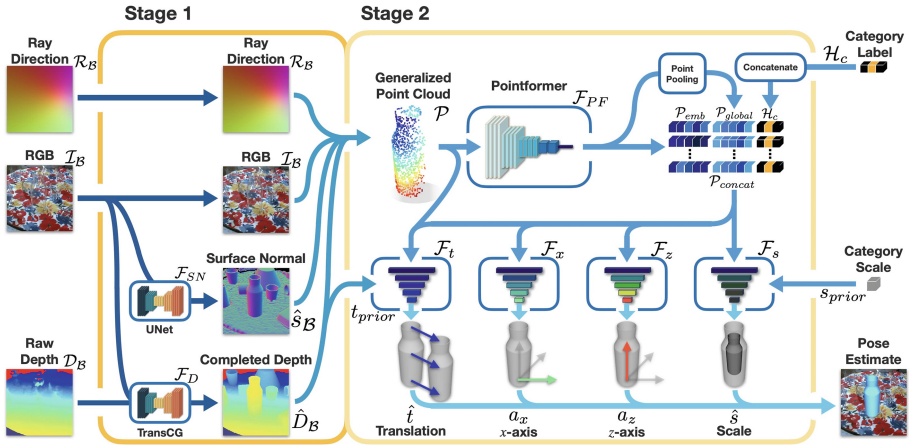
images by a large margin. Zhou *et al.* [44–46] employed light-field images to learn features for robotic grasping and object pose estimation. Along with the proposed methods, massive datasets, across different sensors and both synthetic and real-world domains, have been collected and made public for various related tasks [6, 9, 15, 23, 24, 29, 37, 39, 44, 47]. Compared with these previous works, and to the best of our knowledge we propose the first category-level pose estimation approach for transparent objects. Notably, the proposed approach provides reliable 6D pose and scale estimates across instances with similar shapes.

## 2.2 Opaque Object Category-Level Pose Estimation

Category-level object pose estimation is aimed at estimating unseen objects’ 6D pose within seen categories, together with their scales or canonical shape. To the best of our knowledge, there is not currently any category-level pose estimation works focusing on transparent objects, and the works mentioned below mostly consider opaque objects. They won’t work well for transparency due to their dependence on accurate depth. Wang *et al.* [35] introduced the Normalized Object Coordinate Space (NOCS) for dense 3D correspondence learning, and used the Umeyama algorithm [33] to solve the object pose and scale. They also contributed both a synthetic and a real dataset used extensively by the following works for benchmarking. Later, Li *et al.* [19] extended the idea towards articulated objects. To simultaneously reconstruct the canonical point cloud and estimate the pose, Chen *et al.* [2] proposed a method based on canonical shape space (CASS). Tian *et al.* [31] learned category-specific shape priors from an autoencoder, and demonstrated its power for pose estimation and shape completion. 6D-ViT [48] and ACR-Pose [8] extended this idea by utilizing pyramid visual transformer (PVT) and generative adversarial network (GAN) [10] respectively. Structure-guided prior adaptation (SGPA) [3] utilized a transformer architecture for a dynamic shape prior adaptation. Other than learning a dense correspondence, FS-Net [5] regressed the pose parameters directly, and it proposed to learn two orthogonal axes for 3D orientation. Also, it contributed to an efficient data augmentation process for depth-only approaches. GPV-Pose [7] further improved FS-Net by adding a geometric consistency loss between 3D bounding boxes, reconstruction, and pose. Also with depth as the only input, category-level point pair feature (CPPF) [42] could reduce the sim-to-real gap by learning deep point pairs features. DualPoseNet [20] benefited from rotation-invariant embedding for category-level pose estimation. Differing from other works using segmentation networks to crop image patches as the first stage, CenterSnap [13] presented a single-stage approach for the prediction of 3D shape, 6D pose, and size.

Compared with opaque objects, we find the main challenge to perceive transparent objects is the poor quality of input depth. Thus, the proposed TransNet takes inspiration from the above category-level pose estimation works regarding feature embedding and architecture design. More specifically, TransNet leverages both Pointformer from PVT and the pose decoder from FS-Net and GPV-Pose. In the following section, the TransNet architecture is described, focusing on how to integrate the single-view depth completion module and utilize imperfect depth predictions to learn pose estimates of transparent objects.

### 3 TransNet



**Fig. 3.** Architecture for TransNet. TransNet is a two-stage deep neural network for category-level transparent object pose estimation. The first stage uses an object instance segmentation (from Mask R-CNN [11], which is not included in the diagram) to generate patches of RGB-D then used as input to a depth completion and a surface normal estimation network (RGB only). The second stage uses randomly sampled pixels within the objects’ segmentation mask to generate a generalized point cloud formed as the per-pixel concatenation of ray direction, RGB, surface normal, and completed depth features. Pointformer [48], a transformer-based point cloud embedding architecture, transforms the generalized point cloud into high-dimensional features. A concatenation of embedding features, global features, and a one-hot category label (from Mask R-CNN) is provided for the pose estimation module. The pose estimation module is composed of four decoders, one each for translation,  $x$ -axis,  $z$ -axis, and scale regression respectively. Finally, the estimated object pose is recovered and returned as output.

Given an input RGB-D pair  $(\mathcal{I}, \mathcal{D})$ , our goal is to predict objects’ 6D rigid body transformations  $[\mathbf{R}|\mathbf{t}]$  and 3D scales  $\mathbf{s}$  in the camera coordinate frame, where  $\mathbf{R} \in SO(3)$ ,  $\mathbf{t} \in \mathbb{R}^3$  and  $\mathbf{s} \in \mathbb{R}_+^3$ . In this problem, inaccurate/invalid depth readings exist within the image region corresponding to transparent objects (represented as a binary mask  $\mathcal{M}_t$ ). To approach the category-level pose estimation problem along with inaccurate depth input, we propose a novel two-stage deep neural network pipeline, called **TransNet**.

#### 3.1 Architecture Overview

Following recent work in object pose estimation [5, 7, 34], we first apply a pre-trained instance segmentation module (Mask R-CNN [11]) that has been fine-tuned on the pose estimation dataset to extract the objects’ bounding box

patches, masks, and category labels to separate the objects of interest from the entire image.

The first stage of TransNet takes the patches as input and attempts to correct the inaccurate depth posed by transparent objects. Depth completion (TransCG [9]) and surface normal estimation (U-Net [28]) are applied on RGB-D patches to obtain estimated depth-normal pairs. The estimated depth-normal pairs, together with RGB and ray direction patches, are concatenated to feature patches, followed by a random sampling strategy within the instance masks to generate generalized point cloud features.

In the second stage of TransNet, the generalized point cloud is processed through Pointformer [48], a transformer-based point cloud embedding module, to produce concatenated feature vectors. The pose is then separately estimated in four decoder modules for object translation,  $x$ -axis,  $z$ -axis, and scale respectively. The estimated rotation matrix can be recovered using the estimated two axes. Each component is discussed in more detail in the following sections.

### 3.2 Object Instance Segmentation

Similar to other categorical pose estimation work [7], we train a Mask R-CNN [11] model on the same dataset used for pose estimation to obtain the object’s bounding box  $\mathcal{B}$ , mask  $\mathcal{M}$  and category label  $\mathcal{H}_c$ . Patches of ray direction  $\mathcal{R}_\mathcal{B}$ , RGB  $\mathcal{I}_\mathcal{B}$  and raw depth  $\mathcal{D}_\mathcal{B}$  are extracted from the original data source following bounding box  $\mathcal{B}$ , before inputting to the first stage of TransNet.

### 3.3 Transparent Object Depth Completion

Due to light reflection and refraction on transparent material, the depth of transparent objects is very noisy. Therefore, depth completion is necessary to reduce the sensor noise. Given the raw RGB-D patch  $(\mathcal{I}_\mathcal{B}, \mathcal{D}_\mathcal{B})$  pair and transparent mask  $\mathcal{M}_t$  (a intersection of transparent objects’ masks within bounding box  $\mathcal{B}$ ), transparent object depth completion  $\mathcal{F}_D$  is applied to obtain the completed depth of the transparent region  $\{\hat{\mathcal{D}}_{(i,j)} | (i, j) \in \mathcal{M}_t\}$ .

Inspired by one state-of-the-art depth completion method, TransCG [9], we incorporate a similar multi-scale depth completion architecture into TransNet.

$$\hat{\mathcal{D}}_\mathcal{B} = \mathcal{F}_D(\mathcal{I}_\mathcal{B}, \mathcal{D}_\mathcal{B}) \quad (1)$$

We use the same training loss as TransCG:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_d + \lambda_{smooth} \mathcal{L}_s \\ \mathcal{L}_d &= \frac{1}{N_p} \sum_{p \in \mathcal{M}_t \cap \mathcal{B}} \left\| \hat{\mathcal{D}}_p - \mathcal{D}_p^* \right\|^2 \\ \mathcal{L}_s &= \frac{1}{N_p} \sum_{p \in \mathcal{M}_t \cap \mathcal{B}} \left( 1 - \cos \left\langle \mathcal{N}(\hat{\mathcal{D}}_p), \mathcal{N}(\mathcal{D}_p^*) \right\rangle \right) \end{aligned} \quad (2)$$

where  $\mathcal{D}^*$  is the ground truth depth image patch,  $p \in \mathcal{M}_t \cap \mathcal{B}$  represents the transparent region in the patch,  $\langle \cdot, \cdot \rangle$  denotes the dot product operator and  $\mathcal{N}(\cdot)$  denotes the operator to calculate surface normal from depth.  $\mathcal{L}_d$  is  $L_2$  distance between estimated and ground truth depth within the transparency mask.  $\mathcal{L}_s$  is the cosine similarity between surface normal calculated from estimated and ground truth depth.  $\lambda_{smooth}$  is the weight between the two losses.

### 3.4 Transparent Object Surface Normal Estimation

Surface normal estimation  $\mathcal{F}_{SN}$  estimates surface normal  $\mathcal{S}_{\mathcal{B}}$  from RGB image  $\mathcal{I}_{\mathcal{B}}$ . Although previous category-level pose estimation works [5,7] show that depth is enough to obtain opaque objects' pose, experiments in Sect. 4.3 demonstrate that surface normal is not a redundant input for transparent object pose estimation. Here, we slightly modify U-Net [28] to perform the surface normal estimation.

$$\hat{\mathcal{S}}_{\mathcal{B}} = \mathcal{F}_{SN}(\mathcal{I}_{\mathcal{B}}) \quad (3)$$

We use the cosine similarity loss:

$$\mathcal{L} = \frac{1}{N_p} \sum_{p \in \mathcal{B}} \left( 1 - \cos \langle \hat{\mathcal{S}}_p, \mathcal{S}_p^* \rangle \right) \quad (4)$$

where  $p \in \mathcal{B}$  means the loss is applied for all pixels in the bounding box  $\mathcal{B}$ .

### 3.5 Generalized Point Cloud

As input to the second stage, generalized point cloud  $\mathcal{P} \in \mathbb{R}^{N \times d}$  is a stack of  $d$ -dimensional features from the first stage taken at  $N$  sample points, inspired from [38]. To be more specific,  $d = 10$  in our work. Given the completed depth  $\hat{\mathcal{D}}_{\mathcal{B}}$  and predicted surface normal  $\hat{\mathcal{S}}_{\mathcal{B}}$  from Eq. (1), (3), together with RGB patch  $\mathcal{I}_{\mathcal{B}}$  and ray direction patch  $\mathcal{R}_{\mathcal{B}}$ , a concatenated feature patch is given as  $[\mathcal{I}_{\mathcal{B}}, \hat{\mathcal{D}}_{\mathcal{B}}, \hat{\mathcal{S}}_{\mathcal{B}}, \mathcal{R}_{\mathcal{B}}] \in \mathbb{R}^{H \times W \times 10}$ . Here the ray direction  $\mathcal{R}$  represents the direction from camera origin to each pixel in the camera frame. For each pixel  $(u, v)$ :

$$\begin{aligned} p &= [u \ v \ 1]^T \\ \mathcal{R} &= \frac{K^{-1}p}{\|K^{-1}p\|^2} \end{aligned} \quad (5)$$

where  $p$  is the homogeneous UV coordinate in the image plane and  $K$  is the camera intrinsic. The UV mapping itself is an important cue when estimating poses from patches [14], as it provides information about the relative position and size of the patches within the overall image. We use ray direction instead of UV mapping because it also contains camera intrinsic information.

We randomly sample  $N$  pixels within the transparent mask of the feature patch to obtain the generalized point cloud  $\mathcal{P} \in \mathbb{R}^{N \times 10}$ . A more detailed experiment in Sect. 4.3 explores the best choice of the generalized point cloud.

### 3.6 Transformer Feature Embedding

Given generalized point cloud  $\mathcal{P}$ , we apply an encoder and multi-head decoder strategy to get objects’ poses and scales. We use Pointformer [48], a multi-stage transformer-based point cloud embedding method:

$$\mathcal{P}_{emb} = \mathcal{F}_{PF}(\mathcal{P}) \quad (6)$$

where  $\mathcal{P}_{emb} \in \mathbb{R}^{N \times d_{emb}}$  is a high-dimensional feature embedding. During our experiments, we considered other common point cloud embedding methods such as 3D-GCN [21] demonstrating their power in many category-level pose estimation methods [5, 7]. During feature aggregation for each point, they use the nearest neighbor algorithm to search nearby points within coordinate space, then calculate new features as a weighted sum of the features within surrounding points. Due to the noisy input  $\hat{D}$  from Eq. (1), the nearest neighbor may become unreliable by producing noisy feature embeddings. On the other hand, Pointformer aggregates feature by a transformer-based method. The gradient back-propagates through the whole point cloud. More comparisons and discussions in Sect. 4.2 demonstrate that transformer-based embedding methods are more stable than nearest neighbor-based methods when both are trained on noisy depth data.

Then we use a Point Pooling layer (a multilayer perceptron (MLP) plus max-pooling) to extract the global feature  $\mathcal{P}_{global}$ , and concatenate it with local feature  $\mathcal{P}_{emb}$  and the one-hot category  $\mathcal{H}_c$  label from instance segmentation for the decoder:

$$\begin{aligned} \mathcal{P}_{global} &= \text{MaxPool}(\text{MLP}(\mathcal{P}_{emb})) \\ \mathcal{P}_{concat} &= [\mathcal{P}_{emb}, \mathcal{P}_{global}, \mathcal{H}_c] \end{aligned} \quad (7)$$

### 3.7 Pose and Scale Estimation

After we extract the feature embeddings from multi-modal input, we apply four separate decoders for translation,  $x$ -axis,  $z$ -axis, and scale estimation.

**Translation Residual Estimation.** As demonstrated in [5], residual estimation achieves better performance than direct regression by learning the distribution of the residual between the prior and actual value. The translation decoder  $\mathcal{F}_t$  learns a 3D translation residual from the object translation prior  $t_{prior}$  calculated as the average of predicted 3D coordinate over the sampled pixels in  $\mathcal{P}$ . To be more specific:

$$\begin{aligned} t_{prior} &= \frac{1}{N_p} \sum_{p \in N} K^{-1} [u_p \ v_p \ 1]^T \hat{\mathcal{D}}_p \\ \hat{t} &= t_{prior} + \mathcal{F}_t([\mathcal{P}_{concat}, \mathcal{P}]) \end{aligned} \quad (8)$$

where  $K$  is the camera intrinsic and  $u_p, v_p$  are the 2D pixel coordinate for the selected pixel. We also use the  $L_1$  loss between the ground truth and estimated position:

$$\mathcal{L}_t = |\hat{t} - t^*| \quad (9)$$



**Pose Estimation.** Similar to [5], rather than directly regress the rotation matrix  $R$ , it is more effective to decouple it into two orthogonal axes and estimate them separately. As shown in Fig. 3, we decouple  $R$  into the  $z$ -axis  $a_z$  (red axis) and  $x$ -axis  $a_x$  (green axis). Following the strategy of confidence learning in [7], the network learns confidence values to deal with the problem that the regressed two axes are not orthogonal:

$$\begin{aligned} [\hat{a}_i, c_i] &= \mathcal{F}_i(\mathcal{P}_{concat}), i \in \{x, z\} \\ \theta_z &= \frac{c_x}{c_x + c_z} \left( \theta - \frac{\pi}{2} \right) \\ \theta_x &= \frac{c_z}{c_x + c_z} \left( \theta - \frac{\pi}{2} \right) \end{aligned} \quad (10)$$

where  $c_x, c_z$  denote the confidence for the learned axes.  $\theta$  represents the angle between  $a_x$  and  $a_z$ .  $\theta_x, \theta_z$  are obtained by solving an optimization problem and then used to rotate the  $a_x$  and  $a_z$  within their common plane. More details can be found in [7]. For the training loss, first, we use  $L_1$  loss and cosine similarity loss for axis estimation:

$$\mathcal{L}_{r_i} = |\hat{a}_i - a_i^*| + 1 - \langle \hat{a}_i, a_i^* \rangle, i \in \{x, z\} \quad (11)$$

Then to constrain the perpendicular relationship between two axes, we add the angular loss:

$$\mathcal{L}_a = \langle \hat{a}_x, \hat{a}_z \rangle \quad (12)$$

To learn the axis confidence, we add the confidence loss, which is the  $L_1$  distance between estimated confidence and exponential  $L_2$  distance between the ground truth and estimated axis:

$$\mathcal{L}_{con_i} = |c_i - \exp(\alpha \|\hat{a}_i - a_i^*\|_2)|, i \in \{x, z\} \quad (13)$$

where  $\alpha$  is a constant to scale the distance.

Thus the overall loss for the second stage is:

$$\begin{aligned} \mathcal{L} &= \lambda_s \mathcal{L}_s + \lambda_t \mathcal{L}_t + \lambda_{r_x} \mathcal{L}_{r_x} + \lambda_{r_z} \mathcal{L}_{r_z} \\ &+ \lambda_{r_a} \mathcal{L}_a + \lambda_{con_x} \mathcal{L}_{con_x} + \lambda_{con_z} \mathcal{L}_{con_z} \end{aligned} \quad (14)$$

To deal with object symmetry, we apply specific treatments for different symmetry types. For axial symmetric objects (those that remain the same shape when rotating around one axis), we ignore the loss for the  $x$ -axis, *i.e.*,  $\mathcal{L}_{con_x}, \mathcal{L}_{r_x}$ . For planar symmetric objects (those that remain the same shape when mirrored about one or more planes), we generate all candidate  $x$ -axis rotations. For example, for an object symmetric about the  $x-z$  plane and  $y-z$  plane, rotating the  $x$ -axis about the  $z$ -axis by  $\pi$  radians will not affect the object's shape. The new  $x$ -axis is denoted as  $a_{x_\pi}$  and the loss for the  $x$ -axis is defined as the minimum loss of both candidates:

$$\mathcal{L}_x = \min(\mathcal{L}_x(a_x), \mathcal{L}_x(a_{x_\pi})) \quad (15)$$

**Scale Residual Estimation.** Similar to the translation decoder, we define the scale prior  $s_{prior}$  as the average of scales of all object 3D CAD models within each category. Then the scale of a given instance is calculated as follows:

$$\hat{s} = s_{prior} + \mathcal{F}_s(\mathcal{P}_{concat}) \quad (16)$$

The loss function is defined as the  $L_1$  loss between the ground truth scale and estimated scale:

$$\mathcal{L}_s = |\hat{s} - s^*| \quad (17)$$

## 4 Experiments

**Dataset.** We evaluated TransNet and baseline models on the Clearpose Dataset [6] for categorical transparent object pose estimation. The Clearpose Dataset contains over 350K real-world labeled RGB-D frames in 51 scenes, 9 sets, and around 5M instance annotations covering 63 household objects. We selected 47 objects and categorize them into 6 categories, *bottle*, *bowl*, *container*, *tableware*, *water cup*, *wine cup*. We used all the scenes in set2, set4, set5, and set6 for training and scenes in set3 and set7 for validation and testing. The division guaranteed that there were some unseen objects for testing within each category. Overall, we used 190K images for training and 6K for testing. For training depth completion and surface normal estimation, we used the same dataset split.

**Implementation Details.** Our model was trained in several stages. For all the experiments in this paper, we were using the ground truth instance segmentation as input, which could also be obtained by Mask R-CNN [11]. The image patches were generated from object bounding boxes and re-scaled to a fixed shape of  $256 \times 256$  pixels. For TransCG, we used AdamW optimizer [25] for training with  $\lambda_{smooth} = 0.001$  and the overall learning rate is 0.001 to train the model till converge. For U-Net, we used the Adam optimizer [17] with a learning rate of  $1e^{-4}$  to train the model until convergence. For both surface normal estimation and depth completion, the batch size was set to 24 images. The surface normal estimation and depth completion model were frozen during the training of the second stage.

For the second stage, the training hyperparameters for Pointformer followed those used in [48]. We used data augmentation for RGB features and instance mask for sampling generalized point cloud. A batch size of 18 was used. To balance sampling distribution across categories, 3 instance samples were selected randomly for each of 6 categories. We followed GPV-Pose [7] on training hyperparameters. The learning rate for all loss terms were kept the same during training,  $\{\lambda_{r_x}, \lambda_{r_z}, \lambda_{r_a}, \lambda_t, \lambda_s, \lambda_{con_x}, \lambda_{con_z}\} = \{8, 8, 4, 8, 8, 1, 1\} \times 0.0001$ . We used the Ranger optimizer [22, 41, 43] and used a linear warm-up for the first 1000 iterations, then used a cosine annealing method at the 0.72 anneal point. All the experiments for pose estimation were trained on a 16G RTX3080 GPU for 30 epochs with 6000 iterations each. All the categories were trained on the same model, instead of one model per category.

**Evaluation Metrics.** For category-level pose estimation, we followed [5, 7] using 3D intersection over union (IoU) between the ground truth and estimated 3D bounding box (we used the estimated scale and pose to draw an estimated 3D bounding box) at 25%, 50% and 75% thresholds. Additionally, we used 5°2 cm, 5°5 cm, 10°5 cm, 10°10 cm as metrics. The numbers in the metrics represent the percentage of the estimations with errors under such degree and distance. For Sect. 4.4, we also used separated translation and rotation metrics: 2 cm, 5 cm, 10 cm, 5°, 10° that calculate percentage with respect to one factor.

For depth completion evaluation, we calculated the root of mean squared error (RMSE), absolute relative error (REL) and mean absolute error (MAE), and used  $\delta_{1.05}$ ,  $\delta_{1.10}$ ,  $\delta_{1.25}$  as metrics, while  $\delta_n$  was calculated as:

$$\delta_n = \frac{1}{N_p} \sum_p \mathbf{I} \left( \max \left( \frac{\hat{\mathcal{D}}_p}{\mathcal{D}_p^*}, \frac{\mathcal{D}_p^*}{\hat{\mathcal{D}}_p} \right) < n \right) \quad (18)$$

where  $\mathbf{I}(\cdot)$  represents the indicator function.  $\hat{\mathcal{D}}_p$  and  $\mathcal{D}_p^*$  mean estimated and ground truth depth for each pixel  $p$ .

For surface normal estimation, we calculated RMSE and MAE errors and used 11.25°, 22.5°, and 30° as thresholds. Here 11.25° represents the percentage of estimates with an angular distance less than 11.25° from ground truth surface normal.

#### 4.1 Comparison with Baseline

**Table 1.** Comparison with the baseline on the Clearpose Dataset.

Method	3D <sub>25</sub> ↑	3D <sub>50</sub> ↑	3D <sub>75</sub> ↑	5°2 cm↑	5°5 cm↑	10°5 cm↑	10°10 cm↑
GPV-Pose	<b>93.7</b>	58.3	10.5	0.4	1.5	7.4	9.1
TransNet	90.3	<b>67.4</b>	<b>22.1</b>	<b>2.4</b>	<b>7.5</b>	<b>23.6</b>	<b>27.6</b>

We chose one state-of-the-art categorical opaque object pose estimation model (GPV-Pose [7]) as a baseline, which was trained with estimated depth from TransCG [9] for a fair comparison. From Table 1, TransNet outperformed the baseline in most of the metrics on the Clearpose dataset. 3D<sub>25</sub> is very easy to learn, so there is no huge difference between them. For the rest of the metrics, TransNet achieved around 2× the percentage on 3D<sub>50</sub>, 3× on 10°5 cm, 10°10 cm and 5× on 5°5 cm, 5°2 cm over the baseline. Qualitative results are shown in Fig. 4 for TransNet.

#### 4.2 Embedding Method Analysis

In Table 2, we compared the embedding method between 3D-GCN [21] and Pointformer [48] on TransNet. Modalities for generalized point cloud were depth,



**Fig. 4.** Qualitative results of category-level pose estimates from TransNet. The left column is the original RGB image within our test set and the right column is the pose estimation results. The white bounding box is the ground truth and the colored one is the estimation result. Different colors represent different categories. For axial symmetric objects, because we only care about the scale and z-axis, we use the ground truth x-axis and estimated z-axis to calculate the estimated x-axis, for better visualization. In the figure, there is a pitcher without either ground truth or estimated bounding box because it is not within any of the defined categories, so we ignore it for both training and testing.

RGB and ray direction (without surface normal) for all the trials. The only differences between them were depth type and embedding methods. With ground truth input, 3D-GCN and Pointformer achieved similar results. For some metrics, *i.e.* 5°5 cm, 3D-GCN was even better. But when the ground truth depth was changed to estimated depth (modeling the change from opaque to transparent setting), Pointformer retained much more accuracy than 3D-GCN. Here is our explanation. Like many point cloud embedding methods, 3D-GCN propagates information between nearest neighbors. It is a very efficient method given a point cloud with low noise. But given the completed depth, high noise makes it unstable to pass data among neighbors. While for Pointformer, information is passed through the whole point cloud, no matter how large the noise is. Therefore, given depth information with large uncertainty, the transformer-based embedding method might be more powerful than embedding methods using nearest neighbors.

**Table 2.** Comparison between different embedding methods

Depth type	Embedding	3D <sub>25</sub> ↑	3D <sub>50</sub> ↑	3D <sub>75</sub> ↑	5°2 cm↑	5°5 cm↑	10°5 cm↑	10°10 cm↑
Ground truth	3D-GCN	<b>90.0</b>	<b>84.1</b>	43.0	21.4	<b>48.0</b>	<b>61.8</b>	<b>64.7</b>
	Pointformer	<b>90.0</b>	81.8	<b>56.5</b>	<b>24.1</b>	39.3	59.0	60.7
Estimation	3D-GCN	<b>88.8</b>	59.8	10.4	0.9	3.4	12.3	15.4
	Pointformer	88.5	<b>62.2</b>	<b>17.6</b>	<b>1.6</b>	<b>5.0</b>	<b>17.4</b>	<b>20.9</b>

### 4.3 Ablation Study of Generalized Point Cloud

We explored different combinations of feature inputs for the generalized point cloud to find the one most suitable for TransNet. Results are shown in Table 3. For trials 1 and 2, we compared the effect of adding estimated surface normal to the generalized point cloud. All the metrics demonstrated that the inclusion of surface normal does improve the resulting pose estimation accuracy.

**Table 3.** Ablation study for a different combination of the generalized point cloud. For both trials, we also use RGB as an input feature for the generalized point cloud.

Trial	Depth	Normal	Ray-direction	3D <sub>25</sub> ↑	3D <sub>50</sub> ↑	3D <sub>75</sub> ↑	5°2 cm↑	5°5 cm↑	10°5 cm↑	10°10 cm↑
1	✓		✓	88.5	62.2	17.6	1.6	5.0	17.4	20.9
2	✓	✓	✓	<b>90.3</b>	<b>67.4</b>	<b>22.1</b>	<b>2.4</b>	<b>7.5</b>	<b>23.6</b>	<b>27.6</b>

### 4.4 Depth and Surface Normal Exploration on TransNet

We explored the combination of depth and surface normal with different accuracy. Results in Table 4 and Table 5 show performance for TransCG and U-Net separately. “GT” and “EST” in Table 6 represent ground truth and estimated input for depth and surface normal respectively. From the comparison

**Table 4.** Accuracy for depth completion on Clearpose dataset. All the metrics are calculated within the transparent mask.

Metric	RMSE↓	REL↓	MAE↓	$\delta_{1.05}$ ↑	$\delta_{1.10}$ ↑	$\delta_{1.25}$ ↑
Value	0.055	0.044	0.041	68.93	89.40	98.93

**Table 5.** Accuracy for surface normal estimation on Clearpose dataset.

Metric	RMSE↓	MAE↓	11.25°↑	22.5°↑	30°↑
Value	0.1915	0.1334	56.75	88.45	96.64

**Table 6.** Evaluation for depth and surface normal accuracy on TransNet.

Trial	Depth	Normal	$3D_{25}$ ↑	$3D_{50}$ ↑	$3D_{75}$ ↑	5°2cm↑	5°5cm↑	10°5cm↑	10°10cm↑	5°↑	10°↑	2cm↑	5cm↑	10cm↑
1	GT	GT	<b>95.1</b>	<b>87.7</b>	<b>66.7</b>	<b>31.8</b>	<b>48.4</b>	<b>66.5</b>	<b>66.7</b>	<b>47.3</b>	<b>66.3</b>	<b>63.3</b>	<b>97.9</b>	<b>99.9</b>
2	GT	EST	90.9	82.1	56.3	23.4	36.5	58.0	59.6	37.3	59.6	53.6	97.2	99.9
3	EST	GT	94.0	83.8	34.3	8.1	29.9	47.8	60.3	37.3	61.8	22.2	77.1	97.4
4	EST	EST	90.3	67.4	22.1	2.4	7.5	23.6	27.6	8.8	28.1	16.6	77.4	96.8

of results among trials 1–3, accurate depth is more essential than surface normal for category-level transparent object pose estimation. For instance, as the ground truth depth changes to the estimated depth from trial 1 to trial 3, 5°2cm decreases by 23.7. Compared with surface normal estimation, 5°2cm only decreases by 8.4 between trial 1 and trial 2. More specifically, from decoupled rotation and translation metrics, we can see that 2cm decreases by 41.1 between trial 1 and trial 3 compared to 9.7 between trial 1 and trial 2, meaning that depth accuracy is more important for translation estimation. Focusing on 2cm, 5cm, 10cm between trial 1 and trial 4, the first metric decreases by 46.7 but the latter two lose much less (20.5 for 5cm and 3.1 for 10cm). This can be explained by the result of depth completion accuracy shown in Table 4 (MAE = 0.041m, between 2cm and 5cm). From the comparison of trial 1–4 on metrics 5° and 10°, we can see that either accurate surface normal or accurate depth can support good performance in rotation metrics (for either trial 2 or trial 3, 5° decreases by 10.0 and 10° decreased by around 7). Once we use the estimation version of both, 5° decreases by 38.5 and 10° decreases by 38.2.

## 5 Conclusions

In this paper, we proposed *TransNet*, a two-stage pipeline for category-level transparent object pose estimation. *TransNet* outperformed a baseline by taking advantage of both state-of-the-art depth completion and opaque object category pose estimation. Ablation studies about multi-modal input and feature embedding modules were performed to guide deeper explorations. In the future, we plan to explore how category information can be used earlier in the network for

better accuracy, improve depth completion potentially using additional consistency losses, and extend the model to be category-level across both transparent and opaque instances.

## References

1. Chang, J., et al.: Ghostpose\*: multi-view pose estimation of transparent objects for robot hand grasping. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5749–5755. IEEE (2021)
2. Chen, D., Li, J., Wang, Z., Xu, K.: 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, pp. 11973–11982 (2020)
3. Chen, K., Dou, Q.: SGPA: structure-guided prior adaptation for category-level 6D object pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2773–2782 (2021)
4. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 801–818 (2018)
5. Chen, W., Jia, X., Chang, H.J., Duan, J., Shen, L., Leonardis, A.: FS-NET: fast shape-based network for category-level 6d object pose estimation with decoupled rotation mechanism. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1581–1590 (2021)
6. Chen, X., Zhang, H., Yu, Z., Opipari, A., Jenkins, O.C.: Clearpose: large-scale transparent object dataset and benchmark. arXiv preprint [arXiv:2203.03890](https://arxiv.org/abs/2203.03890) (2022)
7. Di, Y., et al.: GPV-pose: category-level object pose estimation via geometry-guided point-wise voting. arXiv preprint [arXiv:2203.07918](https://arxiv.org/abs/2203.07918) (2022)
8. Fan, Z., et al.: ACR-pose: adversarial canonical representation reconstruction network for category level 6d object pose estimation. arXiv preprint [arXiv:2111.10524](https://arxiv.org/abs/2111.10524) (2021)
9. Fang, H., Fang, H.S., Xu, S., Lu, C.: TRANSCG: a large-scale real-world dataset for transparent object depth completion and grasping. arXiv preprint [arXiv:2202.08471](https://arxiv.org/abs/2202.08471) (2022)
10. Goodfellow, I., et al.: Generative adversarial nets. *Adv. Neural Inf. Process. Syst.* **27** (2014)
11. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2961–2969 (2017)
12. Ichnowski, J., Avigal, Y., Kerr, J., Goldberg, K.: DEX-NERF: using a neural radiance field to grasp transparent objects. arXiv preprint [arXiv:2110.14217](https://arxiv.org/abs/2110.14217) (2021)
13. Irshad, M.Z., Kollar, T., Laskey, M., Stone, K., Kira, Z.: Centersnap: single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation. arXiv preprint [arXiv:2203.01929](https://arxiv.org/abs/2203.01929) (2022)
14. Jiang, X., Li, D., Chen, H., Zheng, Y., Zhao, R., Wu, L.: UNI6D: a unified cnn framework without projection breakdown for 6d pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11174–11184 (2022)
15. Kalra, A., Taamazyan, V., Rao, S.K., Venkataraman, K., Raskar, R., Kadambi, A.: Deep polarization cues for transparent object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8602–8611 (2020)

16. Khaing, M.P., Masayuki, M.: Transparent object detection using convolutional neural network. In: Zin, T.T., Lin, J.C.-W. (eds.) ICBDL 2018. AISC, vol. 744, pp. 86–93. Springer, Singapore (2019). [https://doi.org/10.1007/978-981-13-0869-7\\_10](https://doi.org/10.1007/978-981-13-0869-7_10)
17. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
18. Lai, P.J., Fuh, C.S.: Transparent object detection using regions with convolutional neural network. In: IPPR Conference on Computer Vision, Graphics, and Image Processing, vol. 2 (2015)
19. Li, X., Wang, H., Yi, L., Guibas, L.J., Abbott, A.L., Song, S.: Category-level articulated object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3706–3715 (2020)
20. Lin, J., Wei, Z., Li, Z., Xu, S., Jia, K., Li, Y.: DUALPOSENET: category-level 6D object pose and size estimation using dual pose network with refined learning of pose consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3560–3569 (2021)
21. Lin, Z.H., Huang, S.Y., Wang, Y.C.F.: Convolution in the cloud: learning deformable kernels in 3D graph convolution networks for point cloud analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
22. Liu, L., et al.: On the variance of the adaptive learning rate and beyond. arXiv preprint [arXiv:1908.03265](https://arxiv.org/abs/1908.03265) (2019)
23. Liu, X., Iwase, S., Kitani, K.M.: STEREOBJ-1M: large-scale stereo image dataset for 6d object pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 10870–10879 (2021)
24. Liu, X., Jonschkowski, R., Angelova, A., Konolige, K.: KeyPose: multi-view 3D labeling and keypoint estimation for transparent objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11602–11610 (2020)
25. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint [arXiv:1711.05101](https://arxiv.org/abs/1711.05101) (2017)
26. Lysenkov, I., Eruhimov, V., Bradski, G.: Recognition and pose estimation of rigid transparent objects with a kinect sensor. *Robotics* **273**(273–280), 2 (2013)
27. Phillips, C.J., Lecce, M., Daniilidis, K.: Seeing glassware: from edge detection to pose estimation and shape recovery. In: *Robotics: Science and Systems*, vol. 3, p. 3 (2016)
28. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
29. Sajjan, S., et al.: Clear Grasp: 3D shape estimation of transparent objects for manipulation. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 3634–3642. IEEE (2020)
30. Tang, Y., Chen, J., Yang, Z., Lin, Z., Li, Q., Liu, W.: 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), pp. 5710–5716. IEEE (2021)
31. Tian, M., Ang, M.H., Lee, G.H.: Shape prior deformation for categorical 6D object pose and size estimation. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12366, pp. 530–546. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58589-1\\_32](https://doi.org/10.1007/978-3-030-58589-1_32)



32. Tian, M., Pan, L., Ang, M.H., Lee, G.H.: Robust 6d object pose estimation by learning rgb-d features. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 6218–6224. IEEE (2020)
33. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **13**(04), 376–380 (1991)
34. Wang, C., et al.: Densfusion: 6d object pose estimation by iterative dense fusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3343–3352 (2019)
35. Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.J.: Normalized object coordinate space for category-level 6d object pose and size estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2642–2651 (2019)
36. Weng, T., Pallankize, A., Tang, Y., Kroemer, O., Held, D.: Multi-modal transfer learning for grasping transparent and specular objects. *IEEE Rob. Autom. Lett.* **5**(3), 3791–3798 (2020)
37. Xie, E., Wang, W., Wang, W., Ding, M., Shen, C., Luo, P.: Segmenting transparent objects in the wild. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12358, pp. 696–711. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58601-0\\_41](https://doi.org/10.1007/978-3-030-58601-0_41)
38. Xu, C., Chen, J., Yao, M., Zhou, J., Zhang, L., Liu, Y.: 6DoF pose estimation of transparent object from a single RGB-D image. *Sensors* **20**(23), 6790 (2020)
39. Xu, H., Wang, Y.R., Eppel, S., Aspuru-Guzik, A., Shkurti, F., Garg, A.: Seeing glass: joint point cloud and depth completion for transparent objects. arXiv preprint [arXiv:2110.00087](https://arxiv.org/abs/2110.00087) (2021)
40. Xu, Y., Nagahara, H., Shimada, A., Taniguchi, R.I.: Transcut: transparent object segmentation from a light-field image. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 3442–3450 (2015)
41. Yong, H., Huang, J., Hua, X., Zhang, L.: Gradient centralization: a new optimization technique for deep neural networks. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12346, pp. 635–652. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58452-8\\_37](https://doi.org/10.1007/978-3-030-58452-8_37)
42. You, Y., Shi, R., Wang, W., Lu, C.: CPPF: towards robust category-level 9D pose estimation in the wild. arXiv preprint [arXiv:2203.03089](https://arxiv.org/abs/2203.03089) (2022)
43. Zhang, M., Lucas, J., Ba, J., Hinton, G.E.: Lookahead optimizer: k steps forward, 1 step back. *Adv. Neural Inf. Process. Syst.* **32** (2019)
44. Zhou, Z., Chen, X., Jenkins, O.C.: Lit: Light-field inference of transparency for refractive object localization. *IEEE Rob. Autom. Lett.* **5**(3), 4548–4555 (2020)
45. Zhou, Z., Pan, T., Wu, S., Chang, H., Jenkins, O.C.: Glassloc: plenoptic grasp pose detection in transparent clutter. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4776–4783. IEEE (2019)
46. Zhou, Z., Sui, Z., Jenkins, O.C.: Plenoptic monte carlo object localization for robot grasping under layered translucency. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1–8. IEEE (2018)
47. Zhu, L., et al.: Rgb-d local implicit function for depth completion of transparent objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4649–4658 (2021)
48. Zou, L., Huang, Z., Gu, N., Wang, G.: 6d-vit: category-level 6d object pose estimation via transformer-based instance representation learning. arXiv preprint [arXiv:2110.04792](https://arxiv.org/abs/2110.04792) (2021)