

Tabletop Transparent Scene Reconstruction via Epipolar-Guided Optical Flow with Monocular Depth Completion Prior

Xiaotong Chen¹, Zheming Zhou², Zhuo Deng², Omid Ghasemalizadeh²,
Min Sun², Cheng-Hao Kuo², Arnie Sen²

Abstract—Reconstructing transparent objects using affordable RGB-D cameras is a persistent challenge in robotic perception due to inconsistent appearances across views in the RGB domain and inaccurate depth readings in each single-view. We introduce a two-stage pipeline for reconstructing transparent objects tailored for mobile platforms. In the first stage, off-the-shelf monocular object segmentation and depth completion networks are leveraged to predict the depth of transparent objects, furnishing single-view shape prior. Subsequently, we propose Epipolar-guided Optical Flow (EOF) to fuse several single-view shape priors from the first stage to a cross-view consistent 3D reconstruction given camera poses estimated from opaque part of the scene. Our key innovation lies in EOF which employs boundary-sensitive sampling and epipolar-line constraints into optical flow to accurately establish 2D correspondences across multiple views on transparent objects. Quantitative evaluations demonstrate that our pipeline significantly outperforms baseline methods in 3D reconstruction quality, paving the way for more adept robotic perception and interaction with transparent objects.

I. INTRODUCTION

Transparent objects are prevalent in daily life, from small tabletop objects like cups, bottles, and bowls, to large glass windows and doors. They possess challenges to robotics visual perception mainly two-fold. In the RGB domain, transparent objects lack distinctive visual features, and their appearance is highly correlated to the background and environmental lighting, aggravating the problem. In the depth domain, the non-Lambertian surface cannot produce reliable structure light patterns nor has a consistent refractive index of the medium which leads to missing depth readings from the conventional depth camera as shown in Figure 2. As a result, it remains effectively difficult for robots with RGB-D sensors to correctly reconstruct and interact with transparent objects in household settings.

Most existing works [1]–[6] on table-top transparent object reconstruction are performed under the lab environment with specialized hardware setup. These approaches have strict prerequisites: known material index of refraction, one or two times of light refraction through an isolated solid object, external light sources, or a background with coded patterns. Apart from those prerequisites, These approaches often acquire object silhouettes from multiple views, perform

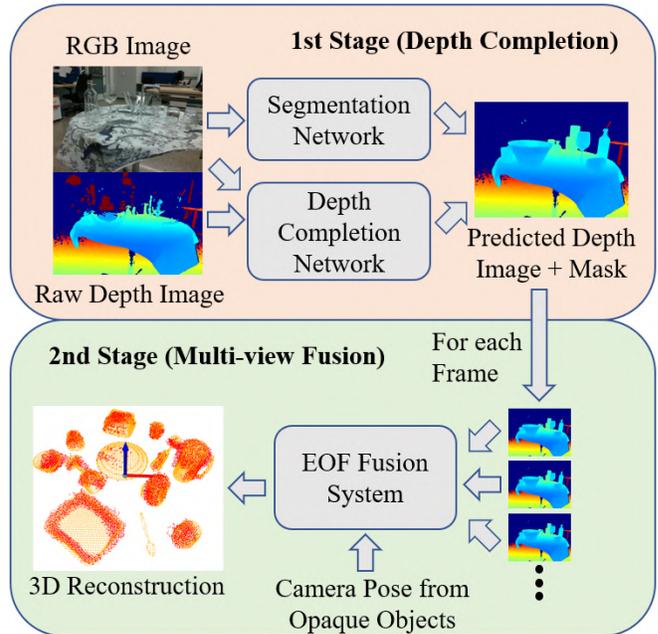


Fig. 1: An overview of *D-EOF* transparent object 3D reconstruction pipeline. The first stage predicts depth and segmentation mask from single viewpoints. The second stage generates 2D correspondences and refines through the epipolar-line-guided optical flow for transparent surfaces. The finalized 2D correspondence and camera pose from opaque objects will be jointly optimized through Bundle Adjustment to output the reconstructed point cloud.

a space carving/visual hull to initialize a 3D convex shape, and then use the consistency between position and normal to refine the 3D shape further.

By contrast, in this paper, we build a tabletop transparent object scene reconstruction pipeline (Figure 1) using a RGB-D camera without any extra external device or strict assumptions on background and lighting conditions. The proposed approach serves as an upstream scene reconstruction module to provide reliable point clouds for downstream robotics reasoning and action over transparent objects.

More specifically, we propose Monocular Depth Prior-based Epipolar-Guided Optical Flow (*D-EOF*) as a two-stage tabletop transparent scene reconstruction approach using a sequence of RGB-D images as input. In the first stage, off-the-shelf single-view depth completion networks and a segmentation network are employed to acquire transparent object masks and approximate depth predictions. However, due to errors in depth predictions, inconsistencies across views are present at this stage, leading to a noisy reconstruction of

¹ X. Chen is with the Department of Robotics, University of Michigan, Ann Arbor, MI, USA. cxt@umich.edu

² Z. Zhou, Z. Deng, O. Ghasemalizadeh, M. Sun, and C.H. Kuo are with Amazon Lab126, Sunnyvale, CA, USA. [{zhemiz, zhuod, ghasemal, minnsun, chkuo, senarnie}@amazon.com](mailto)

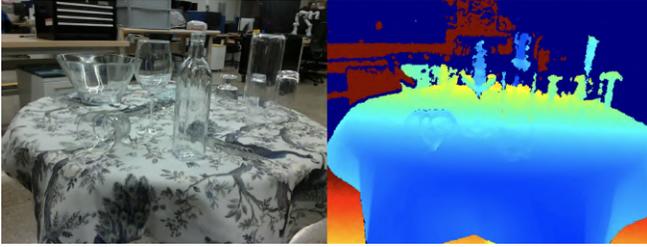


Fig. 2: (Left) an RGB image of a transparent object scene. (Right) raw depth image from an Intel RealSense L515 camera. Compared to ground truth, most pixel values in the transparent region in the raw depth image are missing or inaccurate.

the scene. Addressing this, the second stage formalizes the problem as a Bundle Adjustment (BA) optimization process, focusing on the identification of reliable 2D correspondences on transparent surfaces within the RGB domain. Recognizing the inherent challenges in feature matching on transparent surfaces due to intricate light interactions, we introduce Epipolar-guided Optical Flow (*EOF*). *EOF* integrates two components: (1) Boundary-inspired 2D landmarks generation, which judiciously selects features within transparent regions by capitalizing on boundary information, and (2) Epipolar-guided optical flow correspondence estimation, which monitors the selected landmarks to establish, correct, and refine correspondences utilizing epipolar constraints. This synergistic pairing is pivotal in securing reliable feature matches, which are indispensable for the precise reconstruction of transparent objects.

We assessed *D-EOF* against a state-of-the-art end-to-end transparent object reconstruction method, two state-of-the-art depth completion networks (utilized as our first stage), a general-purpose reconstruction method, and conducted two ablation studies. The comparisons underscored the advantages of our *EOF* approach. In terms of overall 3D reconstruction accuracy and completeness metrics, *D-EOF* surpassed these methods, underlining its effectiveness in reconstructing transparent objects.

II. RELATED WORK

A. Transparent Surface Reconstruction

Reconstruction of transparent surfaces or objects is a challenging problem. Except for intrusive techniques that require physical contact and might destroy the object, most approaches utilize the light reflection or refraction rules. For reflection-based methods, one approach is to estimate the 3D geometry based on the specular highlights from a spotlight [7]–[9], structured light [10] or polarization [11], [12]. Another line of works developed the ‘scanning from heating’ pipeline that determined the shape through laser surface heating and thermal imaging [13], [14]. For refraction-based methods, the shape is solved by analyzing the distortion of background patterns using a time-of-flight camera [15] or light-field camera [16], [17]. Some recent works used the RGB camera to capture multiple views to solve light path triangulation [18] and do environment matting [19] with the help of a turntable [2], [3] or AprilTags [5] to get accurate camera poses.

Without the dependency on external hardware, Li *et al.* [4] designed a physically-based deep rendering network that learned the normal of front and back surfaces and fused features from multiple viewpoints to reconstruct the 3D point cloud. However, this work still requires carefully placed isolated objects and ground truth segmentation. Zhu *et al.* [20] built a transparent SLAM system on a mixed transparent-opaque object scene, where they trained a segmentation network on transparency, used the opaque parts to solve camera trajectory, and reconstructed the transparent objects using visual hull space carving. Compared to this work, we used single-view depth completion plus multi-view optimization that could deal with concave shapes.

B. Transparent Scene Depth Estimation

Transparent objects cause missing or inaccurate values on transparency pixels. Before diving into the depth completion task, object detection or segmentation is a prior task to locate where the transparency lies on the image. Lai *et al.* and Khaing *et al.* [21], [22] pioneered on transparent object detection using CNNs. Xie *et al.* [23] developed a deep segmentation model that outperformed other baselines and a large-scale segmentation dataset. ClearGrasp [24] tackled the problem of depth completion for transparent object scene. They employed depth completion on robotic grasping tasks, where they trained three DeepLabv3+ [25] models to perform image segmentation, surface normal estimation, and boundary segmentation. Recently, various different approaches have been explored to improve the speed and accuracy of the depth completion task, including implicit functions with ray-voxel pairs [26], neural radiance field for rendering [27], combined point cloud and depth estimation [28], adversarial learning [29], multi-view input [30], network structure from RGB image completion [31], and sim2real transfer [32]. 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 [23], [24], [28], [31]–[36].

Our system will utilize the advance in deep networks on single-view depth completion and take their output as a prior for further optimization.

III. TRANSPARENT SCENE RECONSTRUCTION PIPELINE

Given a sequence of RGB-D input \mathcal{O} and corresponding camera poses \mathcal{T} initialized by opaque objects, the objective of *D-EOF* is to infer a point cloud represented by a set of 3D points $\{X_i | i=1, \dots, n\}$ of the transparent objects.

The two-stage transparent scene reconstruction pipeline is shown in Figure 1. The first stage consists of a single-view depth completion neural network and transparent objects segmentation estimation neural network. The second stage is a transparency-aware 2D landmarks sampling followed by *EOF* Bundle Adjustment.

A. Single-view Depth Completion and Segmentation

After evaluating recent advancements in depth completion neural networks [24], [26], [28], [31], we chose to build

the first stage on top of TransCG [31], which offers a good balance between processing efficiency and prediction accuracy. TransCG employs a four-layer U-Net architecture, comprising blocks with intertwined convolutional and fully-connected layers. The network ingests an RGB image alongside a depth image, and outputs a refined depth image with matching resolution. To regulate both the value and gradient of the predicted depth, the network employs a combined loss \mathcal{L} of depth and surface normal:

$$\mathcal{L} = \|\hat{\mathcal{D}} - \mathcal{D}^*\|_2 + \beta(1 - \cos\langle \Delta\hat{\mathcal{D}}_w \times \Delta\hat{\mathcal{D}}_h, \Delta\mathcal{D}_w^* \times \Delta\mathcal{D}_h^* \rangle) \quad (1)$$

where $\hat{\mathcal{D}}$ and \mathcal{D}^* denote predicted and ground truth depth, and $\Delta\mathcal{D}_w, \Delta\mathcal{D}_h$ denote normalized gradient vectors along depth width-axis and height-axis, respectively. β is the weighting factor between two losses.

As the depth completion network is designed to estimate the depth value at transparent pixels specifically, it is common practice that they require a transparent mask as input. For this purpose, we utilize the widely-adopted Mask R-CNN [37] architecture to train a binary segmentation model for transparent objects. Importantly, we fine-tune both TransCG and Mask R-CNN on a dataset specifically consisting of transparent objects to enhance their efficacy. The transparency mask obtained in this stage is also instrumental in the second stage for sampling 2D correspondences.

B. Multi-view Boundary-Inspired and Epipolar-Guided Optical Flow

Utilizing pre-trained networks that have been fine-tuned on a transparent dataset for single-view depth prediction and transparency segmentation, we are able to directly concatenate the masked predicted depth from multiple views to create an initial 3D shape. However, this direct fusion fails to mitigate the noisy depth predictions and inconsistencies across different views (Figure 6 direct concatenation methods). In this section, we introduce the boundary-inspired 2D landmarks generation along with the Epipolar-guided Optical Flow (EOF) Bundle Adjustment algorithm for transparent object point cloud reconstruction. In this process, we initialize the camera poses from SLAM or dense reconstruction methods over the opaque background, similar to [20], [36].

1) *Boundary-Inspired 2D Landmarks Generation*: Inspired by how humans perceive transparent objects, we exploit the boundary feature of transparent objects to find reliable visual landmark correspondence between images. As shown in Figure 3 (Right), the optical flow estimation error is lower in the region near the object boundary while larger in the central areas. Our explanation is shown in Figure 3 (Left). Firstly, for a table-top transparent object, typically a container within the camera field of view, the near-boundary area appears much thicker, which behaves more like opaque surfaces and provides consistent color intensity. Therefore, it provides more reliable correspondence across frames compared to the central area of transparent

objects, which corresponds to a relatively thinner surface and is more sensitive to changes in lighting and background. Secondly, the near-boundary area has a distinct appearance compared to the background, which naturally bounds the area in the 2D landmark correspondence search step. Based on this observation, we use the pixels near the transparent segmentation mask boundary to generate 2D landmarks. Even though this operation will filter out all 2D landmarks located near the center of the transparent object in a single view, our approach can still obtain enough coverage of objects by incorporating different view angles from a sequence of the RGB-D inputs.



Fig. 3: An example that inspires transparent mask boundary-based filtering. (Left) a diagram of camera light rays through a transparent surface, where the effective length is shown as dotted lines. (Middle) an RGB image example of a transparent object scene. (Right) a heatmap corresponding to the RGB image, showing the error of a common dense optical flow estimation between its neighbor frames, measured in pixels.

2) *Epipolar-Guided Optical Flow Correspondence Estimation*: We estimate the 2D correspondences within the local epipolar line segment projected from the predicted depth with uncertainty. As shown in Figure 4, given a 2D pixel sample \mathbf{p} from one keyframe, its ground truth depth point lies near the 3D points projected using the depth prediction \mathbf{x}^* on the camera ray. When the depth prediction error is within a certain distance δ_d , the ground truth depth is between the ‘Near End’ \mathbf{x}^- and ‘Far End’ \mathbf{x}^+ . Accordingly, the ground truth 2D correspondence lies on the projected epipolar line segment l between \mathbf{p}^- and \mathbf{p}^+ projected by \mathbf{x}^- and \mathbf{x}^+ , respectively, in a neighbor frame when the transform \mathbf{T} between the keyframe and this frame is known:

$$\begin{aligned} \mathbf{x}^\pm &= (\hat{\mathcal{D}} \pm \delta_d) \mathbf{K}^{-1} [\mathbf{p} \ 1]^T \\ [\mathbf{p}^\pm \ 1] &= s \cdot \mathbf{K} (\mathbf{R}\mathbf{x}^\pm + \mathbf{t}) \end{aligned} \quad (2)$$

where $\hat{\mathcal{D}}$, \mathbf{K} are the single-view depth prediction, camera intrinsic matrix, and \mathbf{R} , \mathbf{t} are rotation and translation part of transformation \mathbf{T} , respectively. s is the scale factor (inverse depth) to get the pixel value of \mathbf{p}^\pm .

We then use the dense optical flow estimation \mathbf{v}_0^\pm to get a prediction pixel

$$\mathbf{p}_O = \mathbf{p} + \mathbf{v}_0^\pm$$

that might not lie on the epipolar line, and take its orthogonal projection towards the line as the final correspondence estimation \mathbf{p}_{EO} . This correspondence will be ignored when the projected point is outside this line segment. The calculation

can be formulated as

$$p_{EO} = p^- + \frac{(p_O - p^-) \cdot \vec{n}}{|\vec{n}|^2} \vec{n} \quad (3)$$

where

$$\vec{n} = p^+ - p^-$$

is the 2D vector of epipolar line segment.

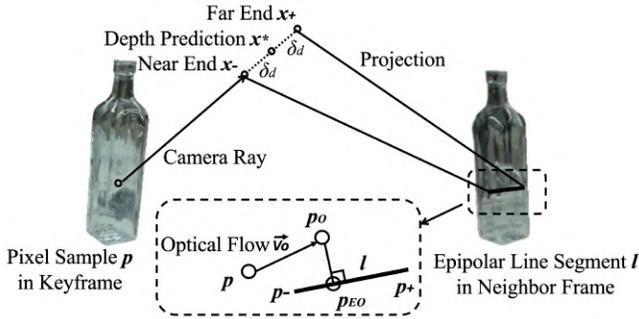


Fig. 4: 2D correspondence estimation based on epipolar line search and optical flow estimation.

3) *Bundle Adjustment Formulation*: We initialize the estimated camera poses from off-the-shelf SLAM systems and formulate a structure-only Bundle Adjustment. The 2D pixel samples p in keyframes with valid correspondences p_{EO} in neighbor frames are set as the 2D landmarks. Their according 3D landmarks x are initialized as x^* and optimized by minimizing the reprojection error e . The optimized 3D points compose the output 3D shape of the *EOF* module.

$$\begin{aligned} x^* &= \hat{D}K^{-1} [p \quad 1]^T \\ e &= p - \pi(T, x) \end{aligned} \quad (4)$$

where π is the projection function same as Equation (2).

IV. EXPERIMENTS

A. Dataset and Network Training

We evaluate our method using the comprehensive ClearPose dataset [36], a recent collection of continuous RGB-D video frames featuring static transparent object scenes. The training set comprises 25 scenes with 200K images, while the test set spans eight scenes, each containing 1,700 images. The latter encompasses challenges such as novel backgrounds, mixed scenes with opaque objects, and significant occlusions. Figure 5 illustrates examples from the test samples. Notably, even though the training and testing sets include identical object instances, our initial stage is ego-centric, employing a depth completion network that remains independent of the object’s specific instance.

The network models are implemented in Pytorch. For binary segmentation, we train a Mask R-CNN model¹ on an RTX 2080 SUPER GPU for 5 epochs, with batch size 5 and a learning rate of 0.005 using the SGD optimizer. After training, Mask R-CNN gets an accuracy of Intersection over Union (IoU) as 0.9542 on the training set and 0.8850 on the test set.

¹<https://github.com/opipari/ClearPose>

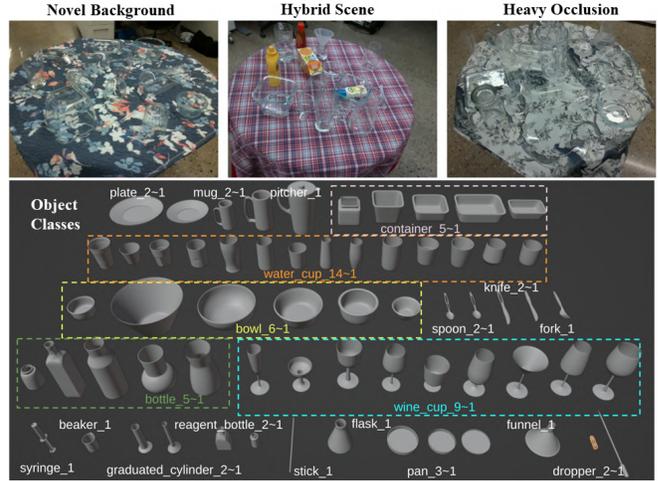


Fig. 5: Examples of testing scenes and objects. (Top row) three challenging test scene categories. (Bottom row) object instances that is used in the testing.

TransCG² is trained on an 8G RTX-3070 GPU for around 300K iterations. β is set to 0.001 following the original paper. We use the L_2 loss to train the TransCG network and improve the depth image resizing by using nearest neighbor replacement to maintain the depth discontinuity property near object boundary. After training, TransCG gets the mean absolute error (MAE) as 0.044m and the root of mean squared error (RMSE) as 0.057m inside the transparent mask on the test set.

B. Evaluation Metrics of 3D Reconstruction

We use a combined metrics of Chamfer distance and inlier ratio threshold to evaluate the quality of 3D reconstruction comprehensively. We define 4 values: *Accuracy*, *Completeness*, *Precision*, and *Recall* to evaluate the difference between the predicted 3D point cloud P_{pred} and the ground truth point cloud P_{gt} . Chamfer distance is defined as the average of the nearest distance from the points in one point cloud to the points in another point cloud. The inlier ratio is defined as the percentage of points in one point cloud within a certain radius d near the other point cloud. Among the 4 values, *Accuracy* and *Precision* evaluate how close the prediction point cloud P_{pred} is close to the ground truth point cloud, while *Completeness* and *Recall* focus on how much of P_{gt} is covered by P_{pred} .

$$\begin{aligned} \text{ChamferDist}(S_1, S_2) &= \frac{1}{|S_1|} \sum_{p_1 \in S_1} \min_{p_2 \in S_2} \|p_1 - p_2\| \\ \text{Accuracy} &= \text{ChamferDist}(P_{pred}, P_{gt}) \\ \text{Completeness} &= \text{ChamferDist}(P_{gt}, P_{pred}) \\ \text{InlierRatio}_d(S_1, S_2) &= \frac{1}{|S_1|} \sum_{p_1 \in S_1} \mathbb{I} \left(\min_{p_2 \in S_2} \|p_1 - p_2\| < d \right) \\ \text{Precision} &= \text{InlierRatio}_d(P_{pred}, P_{gt}) \\ \text{Recall} &= \text{InlierRatio}_d(P_{gt}, P_{pred}) \end{aligned} \quad (5)$$

²<https://github.com/Galaxies99/TransCG>

C. Implementation Details and Baselines

The *EOF* module is implemented as follows. We build optical flow based on Farneback [38] method. For Bundle Adjustment, we implement it using `g2opy`³ library with robust Huber loss. The 2D landmark points are sampled within the transparent masks on the keyframe RGB images. The 3D landmarks in BA are optimized with 30 iterations to get the final 3D point cloud set. The unoptimized implementation of the 2D landmark generation and refinement took approximately 0.3 seconds per frame to execute on a system equipped with an NVIDIA RTX 2080 GPU and an Intel i7-10875H CPU.

To evaluate the accuracy of the proposed correspondence matching method, we compare it with the SIFT feature detection and matching and optical flow estimation. To assess the effectiveness of our second stage (*EOF*), we integrate it with two state-of-the-art single-frame depth completion approaches: SeeingGlass [28] and ImplicitDepth [26]. For these methods, the baseline directly concatenates the network’s outputs across multiple frames. To evaluate our full pipelines, considering the limited work in multi-view transparent reconstruction, we compare it against one state-of-the-art method on single tabletop transparent object reconstruction: ThroughLookingGlass [4] and a general-purpose 3D reconstruction method, TSDF, implemented with PyCUDA [39]. To make TSDF compatible with transparent objects, we feed the predicted masked depth images from our first-stage network into the TSDF pipeline sequentially and then infer a 3D point cloud surface using the ray marching algorithm⁴. We fine-tune the parameters of TSDF fusion to achieve its best performance. Furthermore, for full pipeline evaluation, we conduct two ablation studies to evaluate how correspondences would influence the final 3D reconstruction output. Specifically, we use SIFT and optical-flow-only methods to refine the correspondences before feeding them into our BA pipeline for 3D reconstruction.

D. Results and Discussions

1) *2D Correspondence Estimation Evaluation*: To validate the reliability of near-boundary pixels for tracking, we measured the 2D correspondence accuracy against the ground truth by projecting its depth onto various frames. Average pixel errors for different frame intervals are presented in Table I.

Frame Interval / Pixel Error	2	4	6	8	10	12
SIFT Feature Only	2.380	4.192	6.184	8.364	10.782	13.418
Optical Flow Only	2.634	4.894	6.881	8.907	10.717	12.321
EOF (ours)	0.700	1.254	1.868	2.536	3.124	3.520

TABLE I: 2D correspondence search accuracy result.

In the *SIFT Feature Only* method, we utilized the DoG detector for 2D landmark generation and the SIFT descriptor for matching. The *Optical Flow Only* [38] method initialized

³<https://github.com/uoip/g2opy>

⁴https://scikit-image.org/docs/stable/auto_examples/edges/plot_marching_cubes.html

Accuracy (cm)↓	Novel Background	Hybrid Scenes	Heavy Occlusion	All Scenes
SeeingGlass [28] + Direct Concatenation	1.88	3.80	1.85	2.35
SeeingGlass + EOF	1.44	1.14	2.19	1.28
ImplicitDepth [26] + Direct Concatenation	2.10	4.02	2.53	2.69
ImplicitDepth + EOF	1.00	0.86	1.25	1.04
TransCG + Direct Concatenation	2.63	2.06	3.96	2.45
TransCG + EOF (D-EOF)	0.97	1.52	0.92	1.11
Completeness (cm)↓				
SeeingGlass [28] + Direct Concatenation	0.66	0.72	0.83	0.72
SeeingGlass + EOF	1.01	1.01	0.83	1.18
ImplicitDepth [26] + Direct Concatenation	0.65	0.67	0.83	0.70
ImplicitDepth + EOF	0.92	1.03	0.66	0.94
TransCG + Direct Concatenation	0.51	0.48	0.50	0.59
TransCG + EOF (D-EOF)	0.93	0.65	0.95	0.87
Precision (% ≤2cm)↑				
SeeingGlass [28] + Direct Concatenation	64.21	44.98	62.74	59.04
SeeingGlass + EOF	78.13	82.90	69.12	77.62
ImplicitDepth [26] + Direct Concatenation	62.72	43.22	53.20	55.47
ImplicitDepth + EOF	86.01	89.31	80.60	84.82
TransCG + Direct Concatenation	56.75	64.00	43.97	55.04
TransCG + EOF (D-EOF)	87.38	78.05	88.41	85.09
Recall (% ≤2cm)↑				
SeeingGlass [28] + Direct Concatenation	85.55	76.46	54.03	75.40
SeeingGlass + EOF	78.57	81.31	82.89	68.78
ImplicitDepth [26] + Direct Concatenation	84.55	82.08	59.21	77.61
ImplicitDepth + EOF	86.33	86.76	91.67	80.12
TransCG + Direct Concatenation	95.91	96.07	99.06	92.43
TransCG + EOF (D-EOF)	86.68	89.48	79.58	86.07

TABLE II: Comparison of 3D reconstruction at the scene level, considering results with and without the *EOF*. All methods involving *EOF* are bolded against direct concatenation approaches.

the 2D landmarks within the transparency mask and used dense optical flow for matching.

Our proposed method consistently outperforms both baseline approaches across all frame intervals. Traditional techniques tend to diverge rapidly when applied to 2D landmarks on transparent surfaces due to errors from inconsistent RGB observations. In contrast, our method’s correction and pruning steps ensure more accurate correspondences.

2) *EOF Evaluation*: We further evaluate the efficacy of our proposed second stage – *EOF* with boundary-inspired sampling on improving overall 3D reconstruction performance with different state-of-the-art single view depth completion networks. The results are shown in Table II. The evaluation is conducted on 8 different scenes, which covers 3 different categories and 14 object classes. We evaluate the performance of different baselines using the metrics described in Section IV-B. The results illustrate a notable enhancement in both accuracy and precision upon the introduction of *EOF*, while maintaining comparable recall results. It is important to note that the direct concatenation method, despite having a higher point count and thus better completeness, introduces excessive noise which obscures the true shape of the objects.

3) *End-to-end Evaluation*: We evaluate the performance of our proposed *D-EOF* pipeline with baseline end-to-end methods using the metrics described in Section IV-B. The results are analyzed on both per object and per scene basis and are presented in Table III and Table IV, respectively.

D-EOF outperforms all baseline methods in terms of 3D reconstruction accuracy, precision, completeness and recall at the scene level, as demonstrated by our quantitative results. Regarding object-level evaluation, our method performs significantly better than all other baselines overall. We also observed that our method performs better for objects with significant height dimension, such as containers, wine cups, and mugs, compared to flat objects like spoons, knives, and

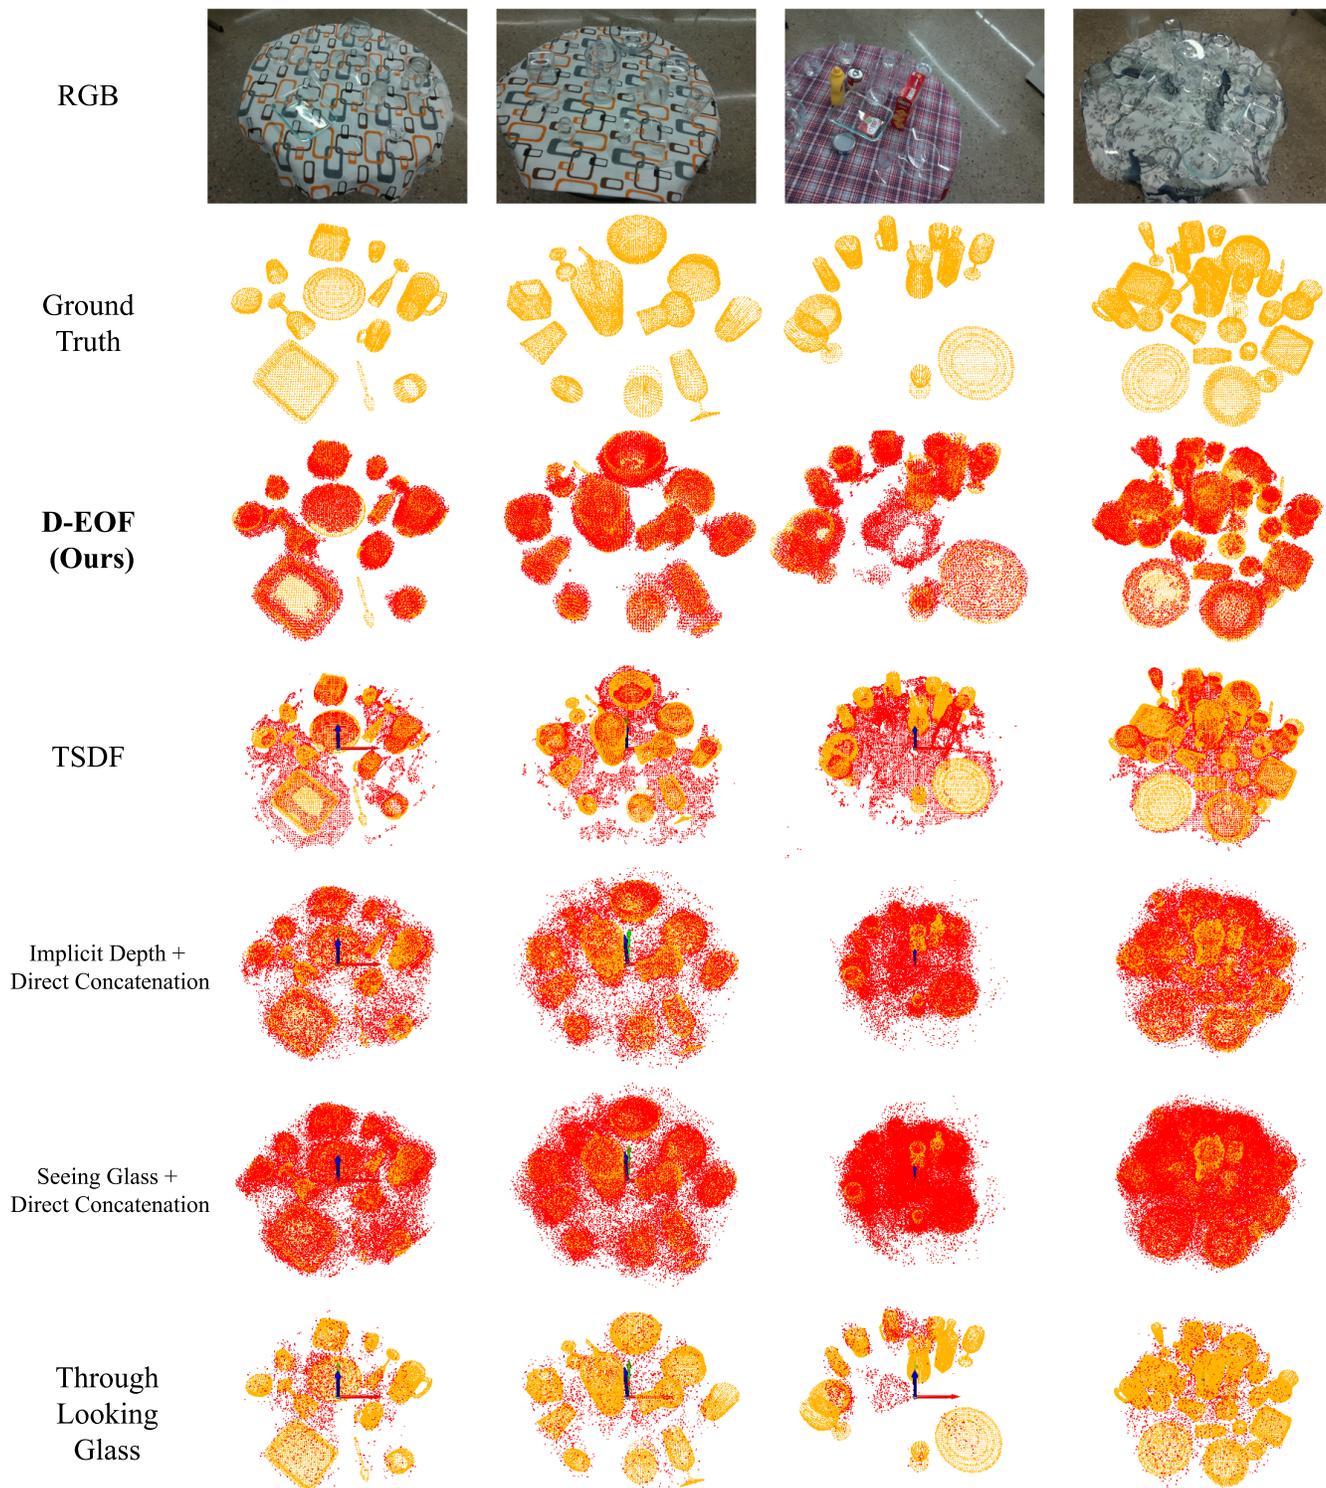


Fig. 6: Visualization comparison as 3D point clouds. The orange point clouds represent the ground truth, while the red point clouds depict the predictions of the transparent scenes, corresponding to the RGB images displayed above. It is evident from the visualization that *D-EOF* produces the most consistent and clean results, aligning closely with the ground truth. In contrast, TSDF incorporates numerous points from the background table and opaque distractor objects, substantially reducing its performance. When employing direct concatenation with Implicit Depth and SeeingGlass, the results lack distinction in the object shapes and blanket the entire space. This illustrates that merely concatenating monocular depth priors results in inconsistency and noise across different viewpoints. ThroughLookingGlass generates very sparse points, which can be attributed to its primary focus on single solid convex transparent objects. It struggles to handle cluttered scenes effectively.

Accuracy (cm)↓	Container	Wine Cup	Water Cup	Mug	Plate	Bowl	Bottle	Spoon	Knife	Fork	Pitcher	Funnel	Syringe	Pan
ThroughLookingGlass [4]	1.83	1.88	1.78	1.50	2.08	1.85	1.75	2.95	2.97	2.43	1.66	1.26	-	-
TSDF	1.27	1.62	1.48	1.09	1.49	1.60	1.47	3.51	3.43	3.59	0.62	1.00	2.49	2.42
SIFT-BA	1.56	1.76	1.77	1.49	1.95	1.67	1.58	2.68	2.95	3.01	1.22	1.49	3.05	2.47
OF-BA	1.64	1.63	1.68	1.44	1.88	1.74	1.57	2.36	2.70	2.78	1.11	1.84	4.07	2.97
D-EOF	0.98	1.05	1.04	0.83	0.77	1.04	0.95	0.93	0.86	0.78	0.97	1.47	0.53	1.40
Completeness (cm)↓														
ThroughLookingGlass [4]	2.11	1.88	1.73	1.74	2.14	2.38	2.49	1.92	2.99	3.21	1.65	1.27	-	-
TSDF	1.46	2.58	1.60	1.27	2.94	1.58	1.95	2.62	4.86	2.82	0.83	1.01	1.13	4.67
SIFT-BA	1.15	1.22	1.45	1.36	1.48	1.33	1.13	2.83	3.21	2.72	0.74	1.72	3.35	1.54
OF-BA	1.19	1.15	1.43	1.00	1.23	1.35	1.12	4.08	3.38	1.84	0.87	0.96	4.42	2.51
D-EOF	0.68	0.66	0.62	0.48	1.08	0.66	0.56	2.53	3.10	2.14	0.67	0.33	0.96	1.79
Precision (% ≤2cm)↑														
ThroughLookingGlass [4]	59.37	58.28	61.12	72.33	50.95	62.39	61.42	47.12	67.14	48.43	63.64	76.46	-	-
TSDF	75.91	65.71	69.23	80.08	71.62	64.82	69.71	39.39	36.72	6.60	91.75	86.13	39.53	30.89
SIFT-BA	67.66	62.51	61.41	70.14	55.28	64.58	67.66	22.92	19.39	12.34	79.39	70.10	4.64	35.74
OF-BA	64.28	65.40	63.94	70.90	56.39	61.95	67.52	28.02	21.16	29.83	83.80	57.73	00.00	16.77
D-EOF	86.48	84.72	85.19	91.70	94.49	84.68	87.57	87.66	98.62	94.64	85.95	71.16	100.00	93.46
Recall (% ≤2cm)↑														
ThroughLookingGlass [4]	19.89	27.67	30.06	29.56	18.90	19.58	28.76	9.05	04.62	17.38	15.67	65.78	-	-
TSDF	63.08	54.94	60.95	72.46	33.33	56.61	55.61	10.20	5.83	3.46	86.54	81.15	59.95	2.30
SIFT-BA	66.19	66.59	58.07	59.36	67.07	63.25	68.24	30.69	26.16	11.18	80.48	32.66	00.00	45.33
OF-BA	68.52	64.29	58.24	73.84	69.35	60.46	67.28	30.02	27.20	23.35	75.42	76.38	00.00	10.07
D-EOF	87.35	88.45	90.17	94.42	70.48	89.81	91.28	33.95	50.00	25.10	86.42	100.0	81.23	23.90

TABLE III: Object-level comparison of 3D reconstruction results across different end-to-end methods.

Accuracy (cm)↓	Novel Background	Hybrid Scenes	Heavy Occlusion	All Scenes
ThroughLookingGlass [4]	2.38	2.75	2.75	2.56
TSDF	2.95	4.42	1.12	2.86
SIFT-BA	2.15	3.06	1.43	1.94
OF-BA	1.95	2.58	1.66	2.02
D-EOF	0.97	1.52	0.92	1.11
Completeness (cm)↓				
ThroughLookingGlass [4]	1.76	3.19	1.93	2.16
TSDF	1.66	2.59	1.80	1.93
SIFT-BA	1.54	1.68	1.15	1.21
OF-BA	1.54	1.23	1.32	1.39
D-EOF	0.93	0.65	0.95	0.87
Precision (% ≤2cm)↑				
ThroughLookingGlass [4]	55.33	56.37	44.96	53.00
TSDF	51.75	36.72	80.42	55.16
SIFT-BA	54.57	48.75	73.31	63.63
OF-BA	58.58	59.83	66.97	61.78
D-EOF	87.38	78.05	88.41	85.09
Recall (% ≤2cm)↑				
ThroughLookingGlass [4]	37.88	37.87	4.92	29.64
TSDF	57.69	43.07	59.26	54.43
SIFT-BA	59.38	54.70	69.80	68.18
OF-BA	57.74	67.18	62.43	62.06
D-EOF	86.68	89.48	79.58	86.07

TABLE IV: Scene-level comparison of 3D reconstruction results across different end-to-end methods.

forks.

Figure 6 visually demonstrates the performance differences among various 3D reconstruction methods across different scenarios. The first two scenes in each subfigure belong to the *New background* category, while the 3rd and 4th scenes represent *With opaque objects* and *Heavy occlusion* categories. ThroughLookingGlass achieves relatively poor results across all four metrics compared to our method. This is because it focuses more on reconstructing a single solid convex object, which does not work well for household objects in a tabletop setup that often have concave shapes and numerous occlusions, leading to missing areas in the reconstruction and very sparse construction. The TSDF method is not specific to transparent 3D construction and produces relatively noisy results, containing many points of the background table and opaque distractor objects, which

significantly diminishes its performance. In comparison, our method establishes clear object shapes with fewer noisy points from the background and displays great robustness across heavy occlusions and opaque distractor objects compared to all other baselines.

There are still limitations in our approach. The 2D landmark generation, which uses heuristic filtering based on object boundaries, presumes thicker surfaces near these boundaries. This may not apply to flat surfaces like windows and doors. Also, the approach relies on segmentation masks and demands dense and complete view angle sampling. A confidence map prediction would be preferable to sample with importance and focus on the area with lower error. For 2D correspondence matching, the current optical flow-based implementation hasn't incorporated the recent advances such as deep learning-based methods [40]. Another point to improve is to consider more than two frames or to involve other modalities, such as surface normal [4] to refine the correspondence selection.

V. CONCLUSION

We presented a two-stage framework, *D-EOF*, for reconstructing transparent objects in cluttered tabletop scenes. Our evaluations demonstrate that the proposed *D-EOF* pipeline yields reliable point clouds, adept at handling the challenges posed by tabletop transparent scenes. For further performance enhancement, integrating learning-based approaches for optical flow estimation and incorporating additional modalities, such as surface normals, are promising directions. Additionally, taking into account instance-level object detection and utilizing shape priors could bring about finer details in reconstructions. This development marks a promising step forward in advancing robotic perception and interaction within environments containing transparent objects.

REFERENCES

- [1] Y. Qian, M. Gong, and Y. H. Yang, "3d reconstruction of transparent objects with position-normal consistency," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4369–4377.
- [2] B. Wu, Y. Zhou, Y. Qian, M. Gong, and H. Huang, "Full 3d reconstruction of transparent objects," *arXiv preprint arXiv:1805.03482*, 2018.
- [3] J. Lyu, B. Wu, D. Lischinski, D. Cohen-Or, and H. Huang, "Differentiable refraction-tracing for mesh reconstruction of transparent objects," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–13, 2020.
- [4] Z. Li, Y.-Y. Yeh, and M. Chandraker, "Through the looking glass: Neural 3d reconstruction of transparent shapes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1262–1271.
- [5] J. Xu, Z. Zhu, H. Bao, and W. Xu, "A hybrid mesh-neural representation for 3d transparent object reconstruction," *arXiv preprint arXiv:2203.12613*, 2022.
- [6] Z. Zhou, T. Pan, S. Wu, H. Chang, and O. C. Jenkins, "Glass-loc: plenoptic grasp pose detection in transparent clutter," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4776–4783.
- [7] N. J. Morris and K. N. Kutulakos, "Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography," in *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 2007, pp. 1–8.
- [8] S.-K. Yeung, T.-P. Wu, C.-K. Tang, T. F. Chan, and S. Osher, "Adequate reconstruction of transparent objects on a shoestring budget," in *CVPR 2011*. IEEE, 2011, pp. 2513–2520.
- [9] K. He, C. Sui, T. Huang, R. Dai, C. Lyu, and Y.-H. Liu, "3d surface reconstruction of transparent objects using laser scanning with ltrf method," *Optics and Lasers in Engineering*, vol. 148, p. 106774, 2022.
- [10] D. Liu, X. Chen, and Y.-H. Yang, "Frequency-based 3d reconstruction of transparent and specular objects," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 660–667.
- [11] X. Xu, Y. Qiao, and B. Qiu, "Reconstructing the surface of transparent objects by polarized light measurements," *Optics Express*, vol. 25, no. 21, pp. 26 296–26 309, 2017.
- [12] S. Mingqi, X. Chongkun, Y. Zhendong, H. Junnan, and W. Xueqian, "Transparent shape from single polarization images," *arXiv preprint arXiv:2204.06331*, 2022.
- [13] G. Eren, O. Aubreton, F. Meriaudeau, L. S. Secades, D. Fofi, A. T. Naskali, F. Truchetet, and A. Ercil, "Scanning from heating: 3d shape estimation of transparent objects from local surface heating," *Optics express*, vol. 17, no. 14, pp. 11 457–11 468, 2009.
- [14] M. Landmann, H. Speck, P. Dietrich, S. Heist, P. Kühmstedt, A. Tünnermann, and G. Notni, "High-resolution sequential thermal fringe projection technique for fast and accurate 3d shape measurement of transparent objects," *Applied Optics*, vol. 60, no. 8, pp. 2362–2371, 2021.
- [15] K. Tanaka, Y. Mukaigawa, H. Kubo, Y. Matsushita, and Y. Yagi, "Recovering transparent shape from time-of-flight distortion," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4387–4395.
- [16] Z. Zhou, Z. Sui, and O. C. Jenkins, "Plenoptic monte carlo object localization for robot grasping under layered translucency," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–8.
- [17] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar, "Refractive shape from light field distortion," in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 1180–1186.
- [18] K. N. Kutulakos and E. Steger, "A theory of refractive and specular 3d shape by light-path triangulation," *International Journal of Computer Vision*, vol. 76, no. 1, pp. 13–29, 2008.
- [19] D. E. Zongker, D. M. Werner, B. Curless, and D. H. Salesin, "Environment matting and compositing," in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 1999, pp. 205–214.
- [20] 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.
- [21] M. P. Khaing and M. Masayuki, "Transparent object detection using convolutional neural network," in *International Conference on Big Data Analysis and Deep Learning Applications*. Springer, 2018, pp. 86–93.
- [22] P.-J. Lai and C.-S. Fuh, "Transparent object detection using regions with convolutional neural network," in *IPPR Conference on Computer Vision, Graphics, and Image Processing*, vol. 2, 2015.
- [23] 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.
- [24] 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.
- [25] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 801–818.
- [26] 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.
- [27] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, "Dex-nerf: Using a neural radiance field to grasp transparent objects," *arXiv preprint arXiv:2110.14217*, 2021.
- [28] 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," *arXiv preprint arXiv:2110.00087*, 2021.
- [29] 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.
- [30] J. Chang, M. Kim, S. Kang, H. Han, S. Hong, K. Jang, and S. Kang, "Ghostpose*: Multi-view pose estimation of transparent objects for robot hand grasping," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5749–5755.
- [31] H. Fang, H.-S. Fang, S. Xu, and C. Lu, "Transcg: A large-scale real-world dataset for transparent object depth completion and grasping," *arXiv preprint arXiv:2202.08471*, 2022.
- [32] 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," in *European Conference on Computer Vision (ECCV)*, 2022.
- [33] 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.
- [34] Z. Zhou, X. Chen, and O. C. Jenkins, "Lit: Light-field inference of transparency for refractive object localization," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4548–4555, 2020.
- [35] X. Liu, S. Iwase, and K. M. Kitani, "Stereoobj-1m: Large-scale stereo image dataset for 6d object pose estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 870–10 879.
- [36] X. Chen, H. Zhang, Z. Yu, A. Opipari, and O. C. Jenkins, "Clearpose: Large-scale transparent object dataset and benchmark," in *European Conference on Computer Vision*, 2022.
- [37] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [38] G. Farnéback, "Two-frame motion estimation based on polynomial expansion," in *Scandinavian conference on Image analysis*. Springer, 2003, pp. 363–370.
- [39] S. Lu, R. Wang, Y. Miao, C. Mitash, and K. Bekris, "Online object model reconstruction and reuse for lifelong improvement of robot manipulation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1540–1546.
- [40] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *European conference on computer vision*. Springer, 2020, pp. 402–419.