

6D Pose Estimation for Textureless Objects on RGB Frames using Multi-View Optimization

Jun Yang*, Wenjie Xue†, Sahar Ghavidel†, and Steven L. Waslander*

Abstract—6D pose estimation of textureless objects is a valuable but challenging task for many robotic applications. In this work, we propose a framework to address this challenge using only RGB images acquired from multiple viewpoints. The core idea of our approach is to decouple 6D pose estimation into a sequential two-step process, first estimating the 3D translation and then the 3D rotation of each object. This decoupled formulation first resolves the scale and depth ambiguities in single RGB images, and uses these estimates to accurately identify the object orientation in the second stage, which is greatly simplified with an accurate scale estimate. Moreover, to accommodate the multi-modal distribution present in rotation space, we develop an optimization scheme that explicitly handles object symmetries and counteracts measurement uncertainties. In comparison to the state-of-the-art multi-view approach, we demonstrate that the proposed approach achieves substantial improvements on a challenging 6D pose estimation dataset for textureless objects.

I. INTRODUCTION

Texture-less rigid objects occur frequently in industrial environments and are of significant interest in many robotic applications. The task of 6D pose estimation aims to detect these objects of known geometry and estimate their 6DoF (Degree of Freedom) poses, i.e., 3D translations and 3D rotations, with respect to a global coordinate frame. In robotic manipulation tasks, accurate object poses are required for path planning and grasp executions [1], [2], [3]. For robotic navigation, 6D object poses provide useful information to the robot for localization and obstacle avoidance [4], [5], [6], [7].

Due to the lack of appearance features, historically, the problem of 6D pose estimation for textureless objects is mainly addressed with depth data [8], [9], [10], [11], [12] or RGB-D images [13], [2], [14], [15], [16]. These approaches can achieve strong pose estimation performance when given high-quality depth data. Despite recent advances in depth acquisition technology, commodity-level depth cameras produce depth maps with low accuracy and missing data when surfaces are too glossy or dark [17], [18], or the object is transparent [19], [20]. Hence, in the past decade, RGB-based solutions have received a lot of attention as an alternative approach [21], [22]. Due to the advancements in deep learning, some learning-based approaches have been recently shown to significantly

boost the object pose estimation performance using only RGB images [23], [24], [25], [26], [27]. However, due to the scale, depth, and perspective ambiguities inherent to a single viewpoint, RGB-based solutions usually have low accuracy for the final estimated 6D poses.

To this end, recent works utilize multiple RGB frames acquired from different viewpoints to enhance their pose estimation results [28], [29], [30], [6], [31], [7]. In particular, these approaches can be further categorized into offline batch-based solutions [28], [30], where all the frames are provided at once, and incremental solutions [29], [6], [31], [7], where frames are provided sequentially. While fusing pose estimates from different viewpoints can improve the overall performance, handling extreme inconsistency, such as appearance ambiguities, rotational symmetries, and possible occlusions, is still challenging. To address these challenges, in this work, we propose a decoupled formulation to factorize the 6D pose estimation problem into a sequential two-step optimization process. Figure 1 shows an overview of the framework. Based on the per-frame predictions of the object’s segmentation mask and 2D center from neural networks, we first optimize the 3D translation and obtain the object’s scale in the image. The acquired scale greatly simplifies the object rotation estimation problem with a template-matching method [21]. A max-mixture formulation [32] is finally adopted to accommodate the multi-modal output distribution present in rotation space. We conduct extensive experiments on the challenging ROBI dataset [33]. In comparison to the state-of-the-art method CosyPose [28], we achieve a substantial improvement with our method (28.5% and 3.4% over its RGB and RGBD version, respectively).

In summary, our key contributions are:

- We propose a novel 6D object pose estimation approach that decouples the problem into a sequential two-step process. This process resolves the depth ambiguities from individual RGB frames and greatly improves the estimate of rotation parameters.
- To deal with the multi-modal uncertainties of object rotation, we develop a rotation optimization scheme that explicitly handles the object symmetries and counteracts measurement ambiguities.

II. RELATED WORK

A. Object Pose Estimation from a Single RGB Image

Many approaches have been presented in recent years to address the pose estimation problem for texture-less

This work was supported by Epson Canada Ltd.

*Jun Yang and Steven L. Waslander are with University of Toronto Institute for Aerospace Studies and Robotics Institute. {jun.yang, steven.waslander}@robotics.utias.utoronto.ca

†Wenjie Xue and Sahar Ghavidel are with Epson Canada {mark.xue, sahar.ghavidel}@ea.epson.com

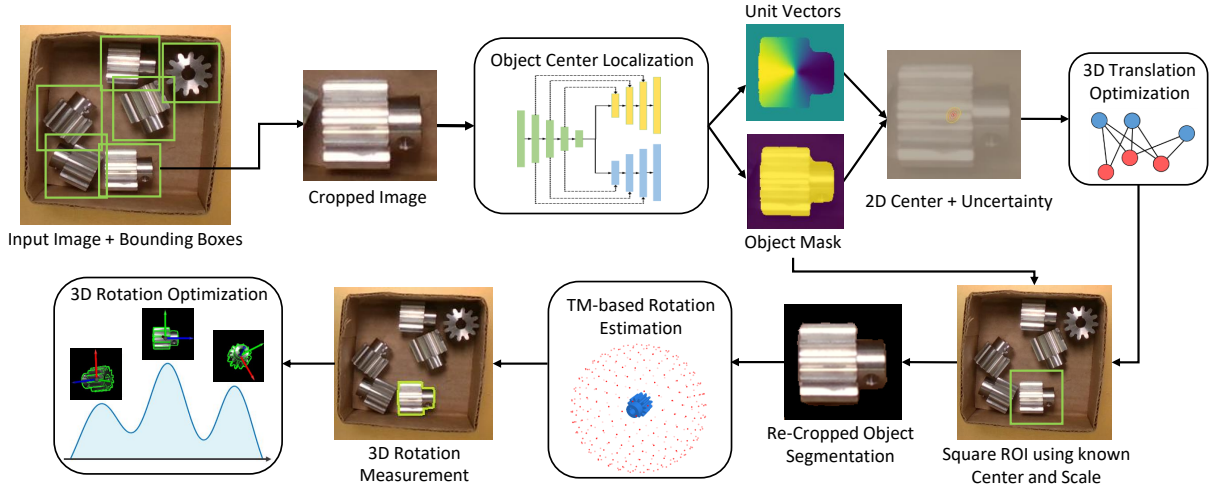


Fig. 1: An overview of the proposed multi-view object pose estimation pipeline with a two-step optimization formulation.

objects with only RGB images. Due to the lack of appearance features, traditional methods usually tackle the problem via holistic template matching techniques [21], [34], [35], but are susceptible to scale change and cluttered environments. More recently, deep learning techniques, such as convolutional neural networks (CNNs), have been employed to overcome these challenges. As two pioneering methods, SSD-6D [23] and PoseCNN [24] developed the CNN architectures to estimate the 6D object poses from a single RGB image. In comparison, some recent works leverage CNNs to first predict 2D object keypoints [36], [37], [26] or dense 2D-3D correspondences [38], [39], [27], [40], and then compute the pose through 2D-3D correspondences with a PnP algorithm [41]. Although these methods show good 2D detection results, the accuracy in the final 6D poses is generally low.

B. Object Pose Estimation from Multiple Viewpoints

Multi-view approaches aim to resolve the scale and depth ambiguities that commonly occur in the single viewpoint setting and improve the accuracy of the estimated poses. Traditional works utilize local features [42], [43] and cannot handle textureless objects. Recently, the multi-view object pose estimation problem has been revisited with neural networks. These approaches used an offline, batch-based optimization formulation, where all the frames are given at once, to obtain a single consistent scene interpretation [44], [28], [20], [30]. Compared to batch-based methods, other works solve the multi-view pose estimation problem in an online manner. These works estimate camera poses and object poses simultaneously, known as object-level SLAM [5], [6], [45], [7], or estimate object poses with known camera poses [1], [29], [31]. Although these methods show performance improvements with only RGB images, they still face difficulty in dealing with object scales, rotational symmetries, and measurement uncertainties.

With the per-frame neural network predictions as measurements, our work resolves the depth and scale am-

biguities by a decoupled formulation. It also explicitly handles rotational symmetries and measurement uncertainties within an incremental online framework.

III. APPROACH OVERVIEW AND PROBLEM FORMULATION

Given the 3D object model and multi-view images, the goal of 6D object pose estimation is to estimate the rigid transformation $\mathbf{T}_{wo} \in SE(3)$ from the object model frame O to a global (world) frame W . We assume that we know the camera poses $\mathbf{T}_{wc} \in SE(3)$ with respect to the world frame. This can be done by robot forward kinematics and eye-in-hand calibration when the camera is mounted on the end-effector of a robotic arm [46], or off-the-shelf SLAM methods for a hand-held camera [47], [48].

Given measurements $\mathbf{Z}_{1:k}$ up to viewpoint k , we aim to estimate the posterior distribution of the 6D object pose $P(\mathbf{R}_{wo}, \mathbf{t}_{wo} | \mathbf{Z}_{1:k})$. The direct computation of this distribution is generally not feasible since object translation \mathbf{t}_{wo} and rotation \mathbf{R}_{wo} have distinct distributions. Specifically, the translation distribution $P(\mathbf{t}_{wo})$ is straightforward and expected to be unimodal. In contrast, the distribution for object rotation $P(\mathbf{R}_{wo})$ is less obvious due to complex uncertainties arising from shape symmetries, appearance ambiguities, and possible occlusions. Inspired by [29], we decouple the pose posterior $P(\mathbf{R}_{wo}, \mathbf{t}_{wo} | \mathbf{Z}_{1:k})$ into:

$$P(\mathbf{R}_{wo}, \mathbf{t}_{wo} | \mathbf{Z}_{1:k}) = P(\mathbf{R}_{wo} | \mathbf{Z}_{1:k}, \mathbf{t}_{wo}) P(\mathbf{t}_{wo} | \mathbf{Z}_{1:k}) \quad (1)$$

where $P(\mathbf{t}_{wo} | \mathbf{Z}_{1:k})$ can be formulated as a unimodal Gaussian distribution $\mathcal{N}(\mathbf{t}_{wo} | \boldsymbol{\mu}, \boldsymbol{\Sigma})$. $P(\mathbf{R}_{wo} | \mathbf{Z}_{1:k}, \mathbf{t}_{wo})$ is the rotation distribution conditioned on the input images $\mathbf{Z}_{1:k}$ and the 3D translation \mathbf{t}_{wo} . To represent the complex rotation uncertainties, similar to [42], we formulate $P(\mathbf{R}_{wo} | \mathbf{Z}_{1:k}, \mathbf{t}_{wo})$ as the mixture of Gaussian distribution:

$$P(\mathbf{R}_{wo} | \mathbf{Z}_{1:k}, \mathbf{t}_{wo}) = \sum_{i=1}^N w_i \mathcal{N}(\mathbf{R}_{wo} | \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \quad (2)$$

which consists of N Gaussian components. The coefficient w_i denotes the weight of the mixture component. $\boldsymbol{\mu}_i$

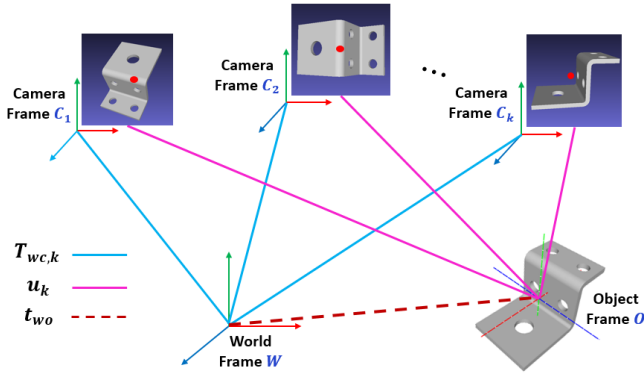


Fig. 2: Illustration of the object, world, and camera coordinates. The 3D translation \mathbf{t}_{wO} is the coordinate of the object model origin in the world coordinate. We can estimate the translation by localizing the per-frame 2D center of the object, \mathbf{u}_k , and minimizing the re-projection errors with known camera poses $\mathbf{T}_{wc,k}$.

and Σ_i are the mean and covariance of i^{th} component, respectively.

Our proposed decoupling formulation implies a useful correlation between translation and rotation in the image domain. The 3D translation estimation \mathbf{t}_{wO} is independent of the object's rotation and encodes the center and scale information of the object. By applying the camera pose $\mathbf{T}_{wc,k}$ at frame k , the estimated 3D translation $\mathbf{t}_{co,k}$ under the camera coordinate provides the scale and 2D center of the object in the image. Based on it, the per-frame object rotation measurement $\mathbf{R}_{co,k}$ can be estimated from its visual appearance in the image. With this formulation, our multi-view framework comprises two main steps, summarized in Figure 1. In the first step (Section IV), we estimate the 3D translation \mathbf{t}_{wO} by integrating the per-frame neural network outputs into an optimization formulation. The network outputs the segmentation mask and the 2D projection \mathbf{u}_k of the object's 3D center, and the object's 3D translation \mathbf{t}_{wO} is estimated by minimizing the 2D re-projection error across views. Given the estimated 3D translation \mathbf{t}_{wO} and segmentation mask, in the second step (Section V), we re-crop a rotation-independent Region of Interest (RoI) for each object with the estimated scale. We then feed the RoI into a rotation estimator to get per-frame 3D rotation measurement $\mathbf{R}_{co,k}$. The final object rotation \mathbf{R}_{wO} is obtained by an optimization approach with the explicitly handling of shape symmetries, and a max-mixture formulation [32] to counteract measurement uncertainties.

IV. 3D TRANSLATION ESTIMATION

As illustrated in Figure 2, the 3D translation \mathbf{t}_{wO} is the coordinate of the object model origin in the world frame. Since the camera pose \mathbf{T}_{wc} is known, it is equivalent to solve the translation from the object model origin to the camera optical center, $\mathbf{t}_{co} = [t_x, t_y, t_z]^T$. Given an RGB image from an arbitrary camera viewpoint, the translation

\mathbf{t}_{co} can be recovered by the following back-projection assuming a pinhole camera model,

$$\begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} = \begin{bmatrix} \frac{u_x - c_x}{f_x} t_z \\ \frac{u_y - c_y}{f_y} t_z \\ t_z \end{bmatrix} \quad (3)$$

where f_x and f_y denote the camera focal lengths, and $[c_x, c_y]^T$ is the principal point. We define $\mathbf{u} = [u_x, u_y]^T$ as the projection of the object model origin O and call it the 2D center of the object in the rest of the paper. We can see that if we can localize object center \mathbf{u} in the image and estimate the depth t_z , then \mathbf{t}_{co} (or \mathbf{t}_{wO}) is solved. In our framework, we predict per-frame 2D object center \mathbf{u} using the power of the neural network and estimate the depth t_z with a multi-view optimization formulation.

Our per-frame 2D object center localization network is shown in the upper part of Figure 1. Our network architecture is based on PVNet [26]. To deal with multiple instances in the scene, we first use off-the-shelf YOLOv5 [49] to detect 2D bounding boxes of the objects. The detections are then cropped and resized to 128x128 before being fed into the network. The network predicts pixel-wise binary labels and a 2D vector field towards the object center. A RANSAC-based voting scheme is finally utilized to estimate the mean \mathbf{u}_k and covariance Σ_k of the object center at frame k . For more details of the object center localization prediction, we refer the reader to [26].

Given a sequence of measurements, we can estimate the object 3D translation \mathbf{t}_{wO} based on the maximum likelihood estimation (MLE) formulation. By assuming the uni-modal Gaussian error model, we solve it with a nonlinear least squares (NLLS) optimization approach. The optimization is formulated by creating measurement residuals that constrain the object translation \mathbf{t}_{wO} with the object center \mathbf{u}_k, Σ_k and known camera pose $\mathbf{T}_{wc,k}$ at viewpoint k ,

$$\mathbf{r}_k(\mathbf{t}_{wO}) = \pi(\mathbf{T}_{wc,k}^{-1} \mathbf{t}_{wO}) - \mathbf{u}_k \quad (4)$$

where π is the perspective projection function. The full problem becomes the minimization of the cost L across all the viewpoints,

$$L = \sum_k \rho_H(\mathbf{r}_k^T \Sigma_k^{-1} \mathbf{r}_k) \quad (5)$$

where Σ_k is the covariance matrix estimated by the localization network for the object center \mathbf{u}_k , and ρ_H is the Huber norm to reduce the impact of outliers for the optimization. We initialize each object's translation \mathbf{t}_{wO} using the diagonal length of its 2D bounding box, similar to [23], from the first frame. With the known camera pose $\mathbf{T}_{wc,k}$, we perform object association based on epipolar geometry constraints and the estimated translation $\mathbf{t}_{wO,1:k-1}$ up to viewpoint $k-1$. Detections that are not associated with any existing objects are initialized as new objects.

We solve the NLLS problem (Equation 4 and 5) in an iterative Gauss-Newton procedure:

$$(\mathbf{J}_{\mathbf{t}_{wO}}^T \Sigma_z^{-1} \mathbf{J}_{\mathbf{t}_{wO}}) \delta \mathbf{t}_{wO} = \mathbf{J}_{\mathbf{t}_{wO}}^T \Sigma_z^{-1} \mathbf{r}(\mathbf{t}_{wO}) \quad (6)$$

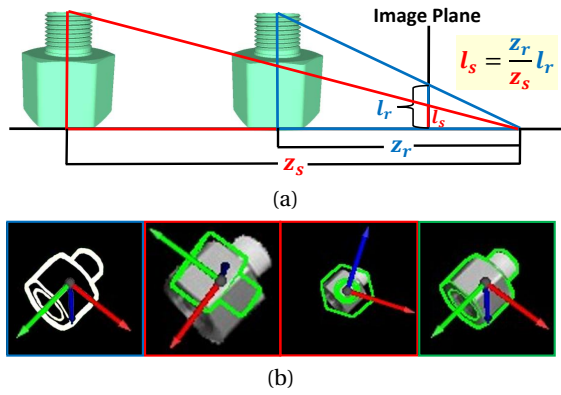


Fig. 3: (a). The inference of object size l_s from their projective ratio. (b) Left: the rendered object template at a canonical distance. Middle: incorrect rotation estimates due to the scale change. Right: re-cropped object RoI using the translation estimate, leading to the correct result.

where Σ_z is the stacked measurement covariance matrix up to the current frame and obtained from the object center localization network (upper part of Figure 1). The over Jacobian, $J_{t_{wo}}$, is stacked up by the per-frame Jacobian matrix $J_{t_{wo},k}$.

V. 3D ROTATION ESTIMATION

The procedure of estimating the object rotation R_{wo} is shown in the lower part of Figure 1. We first adopt a template-matching (TM)-based approach, LINE-2D [21], for obtaining the per-frame rotation measurement $R_{co,k}$. The acquired measurements from multiple viewpoints are then integrated into an optimization scheme. We handle the rotational symmetries explicitly given the object CAD model. To counteract the measurement uncertainties (e.g., from appearance ambiguities), a max-mixture formulation [32] is used to recover a globally consistent set of object pose estimates. Note that the acquisition of the rotation measurement $R_{co,k}$ is not limited to the LINE-2D [21] or TM-based approaches and can be superseded by other holistic-based methods [50], [34], [23], [25].

A. Per-Frame Rotation Measurement

Given the object 3D model, LINE-2D renders object templates from a view sphere in the offline training stage (bottom middle in Figure 1). At run-time, it utilizes the gradient response on input RGB or grayscale images for template matching. A confidence score is provided based on the matching quality. In general, the TM-based approach suffers from scale change issues, and the object templates need to be generated at multiple distances and scales. In our work, instead of training the multi-scale templates, which increases the run-time complexity, we fix the 3D translation to a canonical centroid distance $t_r = [0, 0, z_r]$. At run-time, given the 3D translation hypothesis $t_{co} = [x_s, y_s, z_s]$ from object origin to camera center (obtained from t_{wo} and camera pose T_{wc}), we can re-crop the RoI from the image. The RoI size l_s is determined by $l_s = \frac{z_r}{z_s} l_r$, where l_r and z_r are the RoI size and canonical

distance at training time, respectively. This process is illustrated in Figure 3a. Note that the RoI is a square region here and is independent of the object's rotation. To further reduce the gap between rendered templates and RoI images, we take the segmentation mask from the object center localization network (upper part of Figure 1) and then feed the re-cropped object RoI into the LINE-2D estimator to get a per-frame rotation measurement $R_{co,k}$, as shown in Figure 3b.

B. Optimization formulation

Generally, estimating object 3D rotation from a sequence of measurements can also be formulated as an MLE problem:

$$\hat{X} = \operatorname{argmax}_X \prod_k p(z_k | X) \quad (7)$$

where X denotes the object 3D rotation R_{wo} to be estimated. The measurement z_k here is the object's rotation with respect to the camera coordinate $R_{co,k}$, obtained from Section V-A. The measurement model is a function of camera pose (rotation part) $R_{wc,k}$ and object rotation R_{wo} in world frame:

$$h(R_{wo}, R_{wc,k}) = R_{wc,k}^{-1} R_{wo} \quad (8)$$

We formulate the optimization problem by creating the residual between R_{wo} and per-frame measurement $R_{co,k}$:

$$r_k(R_{wo}) = \log(R_{co,k} h(R_{wo}, R_{wc,k})^{-1})^V \quad (9)$$

where $r_k(R_{wo})$ is expressed by Lie algebra $\mathfrak{so}(3)$. To handle rotational symmetries, we consider them explicitly together with the measurement $R_{co,k}$ in Equation 9. Generally, when an object has symmetry, there exist a set of rotations that leave the object's appearance unchanged:

$$S(R_{co}) = \{R'_{co} \in SO(3) \text{ s.t. } \forall \mathcal{G}(R_{co}) = \mathcal{G}(R'_{co})\} \quad (10)$$

where $\mathcal{G}(R_{co})$ is the rendered image of object under rotation R_{co} (assuming the same object translation). We can update the measurement $R_{co,k}$ in Equation 9 to $\bar{R}_{co,k}$:

$$\bar{R}_{co,k} = \operatorname{argmin}_{R'_{co,k} \in S(R_{co,k})} \left\| \log\left(\left(R'_{co,k}\right) h(R_{wo}, R_{wc,k})^{-1}\right)^V \right\| \quad (11)$$

where $\|\cdot\|$ denotes the absolute angle for a 3D rotation vector ϕ , and $\bar{R}_{co,k}$ is the updated rotation measurement that has the minimal loss relative to R_{wo} .

C. Measurement ambiguities

Due to complex uncertainties, such uni-modal estimates are still not sufficient to adequately represent the rotation uncertainties. To this end, we now consider the sum-mixture of Gaussians as the likelihood function:

$$p(\bar{z}_k | X) = \sum_{i=1}^N w_i \mathcal{N}(\mu_i, \Sigma_i) \quad (12)$$

where \bar{z}_k is the updated measurement (using Equation 11), and each $\mathcal{N}(\mu_i, \Sigma_i)$ represents a distinct Gaussian distribution, and w_i is the weight for component i .

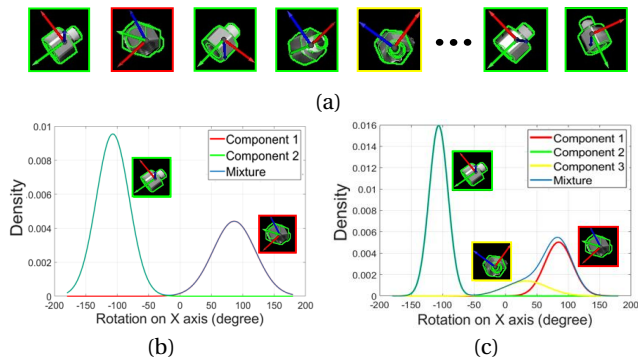


Fig. 4: **Max-mixtures for processing the rotation measurements.** Note that we show the distribution only on one axis for demonstration purposes. (a) Acquired rotation measurements from different viewpoints. (b) Mixture distribution after two viewpoints. (c) Mixture distribution after five viewpoints.

The problem with a sum-mixture is that the MLE solution is no longer simple and falls outside the common NLLS optimization approaches. Instead, we consider the max-marginal and solve the problem with the following max-mixture formulation [32]:

$$p(\bar{z}_k | \mathbf{X}) = \max_{i=1:N} w_i \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \quad (13)$$

The max operator acts as a selector, keeping the problem as a common NLLS optimization. Note that the Max-mixture does not make a permanent choice. In an iteration of optimization, only one of the Gaussian components is selected and optimized. In particular, given a new rotation measurement $\bar{\mathbf{R}}_{co,k}$ at frame k , we actually evaluate each Gaussian component in Equation 13 by computing the absolute rotation angle error $\boldsymbol{\theta}_{k,i}$ between $\bar{\mathbf{R}}_{co,k}$ and $h(\mathbf{R}_{wo,i}, \mathbf{R}_{wc,k})$,

$$\boldsymbol{\theta}_{k,i} = \left\| \log \left(\bar{\mathbf{R}}_{co,k} h(\mathbf{R}_{wo,i}, \mathbf{R}_{wc,k})^{-1} \right)^\vee \right\| \quad (14)$$

and select the one with the minimal angle error. To reduce the impact of outliers, the selected Gaussian component will accept a rotation measurement only if the rotation angle error $\boldsymbol{\theta}_{k,i}$ is less than a pre-defined threshold (30° in our implementation). If the measurement $\bar{\mathbf{R}}_{co,k}$ is not accepted by any Gaussian component, it will be considered as a new component and added to the current Gaussian-mixture model. We optimize the object rotation \mathbf{R}_{wo} within each component by operating on the tangent space $\mathfrak{so}(3)$:

$$\left(\mathbf{J}_{\phi_{wo}}^T \boldsymbol{\Lambda}_z \mathbf{J}_{\phi_{wo}} \right) \delta \phi_{wo} = \mathbf{J}_{\phi_{wo}}^T \boldsymbol{\Lambda}_z \mathbf{r}(\mathbf{R}_{wo}) \quad (15)$$

where $\mathbf{r}(\mathbf{R}_{wo})$ and $\mathbf{J}_{\phi_{wo}}$ are the stacked rotation residual vector and Jacobian matrix, respectively, across multiple viewpoints. We approximate the weight matrix $\boldsymbol{\Lambda}_z$ by placing the LINE-2D confidence score on its diagonal elements.

To compute the weight, w_i , for each Gaussian component, we accumulate the LINE-2D confidence score, c_i , from the rotation measurements within each individual




Fig. 5: Examples of (a) the generated synthetic data and (b) real images from the RealSense Camera.

component across the viewpoints. The weight can be approximated as: $w_i = \frac{c_i}{\sum_i c_i}$. This processing is illustrated in Figure 4. Given the measurements from two viewpoints, the object rotation distribution $P(\mathbf{R}_{wo})$ is represented with two Gaussian components (green and red) with similar weights. The third component (yellow) is added when observing more viewpoints. As a result of receiving more rotation measurements (after five viewpoints), the weight of the correct component (green) becomes higher than the false hypotheses.

VI. EXPERIMENTS

A. Datasets, Baselines and Evaluation Metrics

We evaluate our framework on the recently released ROBI dataset [33]. It provides multiple camera viewpoints and ground truth 6D poses for textureless reflective industrial parts. The objects were placed in challenging bin scenarios and captured using two sensors: a high-cost Ensenso camera and a low-cost RealSense camera. For network training purposes, we generate 80,000 synthetic images using Blender software [51] with Bullet physics engine [52] and train our object center localization network with only synthetic data. Figure 5 presents some examples of our generated synthetic images. We picked five objects that are textureless and evaluated them on both Ensenso and RealSense test sets.

Quantitatively, we compare our approach with CosyPose [28], a state-of-the-art multi-view pose fusion solution which takes the object pose estimates from individual viewpoints as the input and optimizes the overall scene consistency. Note that, CosyPose is an offline batch-based solution that is agnostic to any particular pose estimator. For a fair comparison, we use the same pose estimator (LINE-2D template matching with the same bounding boxes, object center, and segmentation mask from the object center localization network). Additionally, we provide the CosyPose with known camera poses. To feed the reliable single-view estimates to CosyPose, we use two strategies to obtain the scale information for the LINE-2D pose estimator. For the first strategy, we generate the templates at multiple distances in the training time (9 in our experiments) and perform the standard template

Objects	Ensenso						RealSense					
	4 Views			8 Views			4 Views			8 Views		
	CosyPose		Ours	CosyPose		Ours	CosyPose		Ours	CosyPose		Ours
Input Modality	RGB	RGBD	RGB	RGB	RGBD	RGB	RGB	RGBD	RGB	RGB	RGBD	RGB
Tube Fitting	51.6	80.8	86.1	66.2	86.7	88.7	42.6	64.7	77.9	72.1	76.5	85.2
Chrome Screw	39.7	66.1	64.9	58.6	75.3	67.8	60.0	62.9	67.1	72.9	84.3	81.4
Eye Bolt	36.5	77.0	78.4	62.1	87.8	83.8	26.5	58.8	79.4	55.9	91.1	91.1
Gear	33.3	83.9	81.5	45.6	85.2	86.4	41.7	77.8	75.0	61.1	83.3	86.1
Zigzag	51.7	75.8	82.7	60.3	89.7	94.8	53.6	71.4	78.6	64.3	78.6	89.3
ALL	42.6	76.7	78.7	58.6	84.9	84.3	44.9	67.1	75.6	65.3	82.8	86.6

TABLE I: 6D object pose estimation results on Ensenso test set from ROBI dataset, evaluated with the metrics of correct detection rate. There are a total of nine scenes for the Ensenso test set and four scenes for the RealSense test set.

matching in run-time. This strategy can significantly improve the single view pose estimation performance by sacrificing the run-time speed and is treated as the RGB version of CosyPose in our experiments. For the second strategy, we directly use the depth images at run-time to acquire the object scale and refer to it as the RGBD version for CosyPose. Note that, for our approach, we only use RGB images without any depth data.

We adopt the average distance (ADD) metric [53] for evaluation. We transform the object model points by the ground truth and the estimated 6D poses, respectively, and compute the mean of the pairwise distances between the two transformed point sets. A pose is claimed as correct if its ADD is smaller than 10% of the object diameter. A ground truth pose will be considered only if its visibility score is larger than 75%.

B. Results

We conduct the experiments on the ROBI dataset with a variable number of viewpoints (4 and 8). The object pose estimation results are presented in Table I of five highly reflective textureless objects. The results show our method outperforms the baseline CosyPose by a wide margin on RGB data, and is competitive with the RGBD approach. On the Ensenso test set, RGBD version CosyPose achieves an overall 76.7% detection with four views and 84.9% with eight views. It can be considered as an upper bound of CosyPose as it uses depth data at test time. In comparison, our approach outperforms the RGB version CosyPose by a large margin of 36.1% and 25.7% on the 4-view and 8-view test set, respectively, and achieves the upper bound of CosyPose. On the RealSense test set, RGBD version CosyPose performs slightly worse than Ensenso results, mainly due to the poorer depth data quality from the RealSense sensor. In comparison, our approach only relies on RGB images and outperforms both RGB and RGBD versions CosyPose by 26.0% and 6.1%, respectively.

C. Ablation Study on Decoupled Formulation

As discussed in Section III and V-A, the core idea of our method is the decoupling of 6D pose estimation into a sequential two-step process. This process first resolves the scale and depth ambiguities in the RGB images and greatly improves the rotation estimation performance. To justify

Method	Evaluation	Detection Rate (%)		Run-time ¹ (ms)
		4 Views	8 Views	
Simultaneous Process		74.4	81.5	104.5
Two-Step Process		77.2	85.5	25.6

TABLE II: Ablation studies on ROBI dataset with different configurations for object pose estimation.

its effectiveness, we consider an alternative version of our approach, one which simultaneously estimates the 3D translation and rotation. This version uses the same strategy to estimate the object translation. However, instead of using the provided scale from the translation estimates, it uses the multi-scale trained templates (similar to the RGB version of CosyPose) to acquire rotation measurements. Table II presents the result of our ablation study. Due to the large volume of the templates, the run-time for rotation estimation is generally slow for the simultaneous process version. In comparison, our two-step process not only operates with a much faster run-time speed but also has better overall performance.

VII. CONCLUSION

In this work, we have implemented a multi-view pose estimation framework for textureless objects using only RGB images. Our core idea of our method is to decouple the posterior distribution into the 3D translation and the 3D rotation of an object and integrate the per-frame measurements with a two-step multi-view optimization formulation. This process first resolves the scale and depth ambiguities in the RGB images and greatly simplifies the per-frame rotation estimation problem. Moreover, our rotation optimization module explicitly handles the object symmetries and counteracts the measurement uncertainties with a max-mixture-based formulation. Experiments on the real ROBI dataset demonstrate the effectiveness and accuracy compared to the state-of-the-art. Future work includes joint camera pose estimation and 6D object pose estimation, and the active perception to strategically select camera viewpoints for estimating the object poses.

¹We conduct the run-time analysis for **template matching only** and report with milliseconds per object. The analysis is conducted on a desktop with an Intel 3.40GHz CPU and an Nvidia RTX 2080 Ti GPU.

REFERENCES

- [1] X. Deng, Y. Xiang, A. Mousavian, C. Eppner, T. Bretl, and D. Fox, "Self-supervised 6d object pose estimation for robot manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3665–3671, IEEE, 2020.
- [2] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese, "Densefusion: 6d object pose estimation by iterative dense fusion," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3343–3352, 2019.
- [3] K. Wada, E. Sucar, S. James, D. Lenton, and A. J. Davison, "Morefusion: Multi-object reasoning for 6d pose estimation from volumetric fusion," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14540–14549, 2020.
- [4] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. Kelly, and A. J. Davison, "Slam++: Simultaneous localisation and mapping at the level of objects," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1352–1359, 2013.
- [5] S. Yang and S. Scherer, "Cubeslam: Monocular 3-d object slam," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 925–938, 2019.
- [6] J. Fu, Q. Huang, K. Doherty, Y. Wang, and J. J. Leonard, "A multi-hypothesis approach to pose ambiguity in object-based slam," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7639–7646, IEEE, 2021.
- [7] N. Merrill, Y. Guo, X. Zuo, X. Huang, S. Leutenegger, X. Peng, L. Ren, and G. Huang, "Symmetry and uncertainty-aware object slam for 6dof object pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14901–14910, 2022.
- [8] B. Drost, M. Ulrich, N. Navab, and S. Ilic, "Model globally, match locally: Efficient and robust 3d object recognition," in *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 998–1005, IEEE, 2010.
- [9] M. Bui, S. Zakharov, S. Albarqouni, S. Ilic, and N. Navab, "When regression meets manifold learning for object recognition and pose estimation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6140–6146, IEEE, 2018.
- [10] G. Gao, M. Lauri, Y. Wang, X. Hu, J. Zhang, and S. Frintrop, "6d object pose regression via supervised learning on point clouds," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3643–3649, IEEE, 2020.
- [11] G. Gao, M. Lauri, X. Hu, J. Zhang, and S. Frintrop, "Clouadae: Learning 6d object pose regression with on-line data synthesis on point clouds," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11081–11087, IEEE, 2021.
- [12] D. Cai, J. Heikkilä, and E. Rahtu, "Ove6d: Object viewpoint encoding for depth-based 6d object pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6803–6813, 2022.
- [13] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6d object pose and predicting next-best-view in the crowd," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3583–3592, 2016.
- [14] Y. He, W. Sun, H. Huang, J. Liu, H. Fan, and J. Sun, "Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11632–11641, 2020.
- [15] M. Tian, L. Pan, M. H. Ang, and G. H. Lee, "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.
- [16] L. Saadi, B. Besbes, S. Kramm, and A. Bensch, "Optimizing rgb-d fusion for accurate 6dof pose estimation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2413–2420, 2021.
- [17] C.-Y. Chai, Y.-P. Wu, and S.-L. Tsao, "Deep depth fusion for black, transparent, reflective and texture-less objects," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6766–6772, IEEE, 2020.
- [18] J. Yang and S. L. Waslander, "Next-best-view prediction for active stereo cameras and highly reflective objects," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 3684–3690, IEEE, 2022.
- [19] 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)*, pp. 3634–3642, IEEE, 2020.
- [20] 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*, pp. 11602–11610, 2020.
- [21] S. Hinterstoisser, C. Cagniart, S. Ilic, P. Sturm, N. Navab, P. Fua, and V. Lepetit, "Gradient response maps for real-time detection of textureless objects," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 5, pp. 876–888, 2011.
- [22] E. Brachmann, F. Michel, A. Krull, M. Y. Yang, S. Gumhold, et al., "Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3364–3372, 2016.
- [23] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again," in *Proceedings of the IEEE international conference on computer vision*, pp. 1521–1529, 2017.
- [24] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," in *Robotics: Science and Systems (RSS)*, 2018.
- [25] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in *Proceedings of the european conference on computer vision (ECCV)*, pp. 699–715, 2018.
- [26] 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*, pp. 4561–4570, 2019.
- [27] T. Hodan, D. Barath, and J. Matas, "Epos: Estimating 6d pose of objects with symmetries," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11703–11712, 2020.
- [28] Y. Labbé, J. Carpentier, M. Aubry, and J. Sivic, "Cosypose: Consistent multi-view multi-object 6d pose estimation," in *European Conference on Computer Vision*, pp. 574–591, Springer, 2020.
- [29] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "Poserbpf: A rao-blackwellized particle filter for 6-d object pose tracking," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1328–1342, 2021.
- [30] I. Shugurov, I. Pavlov, S. Zakharov, and S. Ilic, "Multi-view object pose refinement with differential renderer," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2579–2586, 2021.
- [31] K.-K. Maninis, S. Popov, M. Niesser, and V. Ferrari, "Vid2cad: Cad model alignment using multi-view constraints from videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [32] E. Olson and P. Agarwal, "Inference on networks of mixtures for robust robot mapping," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 826–840, 2013.
- [33] J. Yang, Y. Gao, D. Li, and S. L. Waslander, "Robi: A multi-view dataset for reflective objects in robotic bin-picking," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9788–9795, IEEE, 2021.
- [34] M. Imperoli and A. Pretto, "D co: Fast and robust registration of 3d textureless objects using the directional chamfer distance," in *Computer Vision Systems: 10th International Conference, ICVS 2015, Copenhagen, Denmark, July 6-9, 2015, Proceedings*, pp. 316–328, Springer, 2015.
- [35] T. Hodaň, X. Zabulis, M. Lourakis, Š. Obdržálek, and J. Matas, "Detection and fine 3d pose estimation of texture-less objects in rgb-d images," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4421–4428, IEEE, 2015.
- [36] M. Rad and V. Lepetit, "Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth," in *Proceedings of the IEEE international conference on computer vision*, pp. 3828–3836, 2017.
- [37] G. Pavlakos, X. Zhou, A. Chan, K. G. Derpanis, and K. Daniilidis, "6-dof object pose from semantic keypoints," in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2011–2018, IEEE, 2017.
- [38] S. Zakharov, I. Shugurov, and S. Ilic, "Dpod: 6d pose object detector and refiner," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1941–1950, 2019.
- [39] K. Park, T. Patten, and M. Vincze, "Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7668–7677, 2019.

- [40] 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*, pp. 6749–6758, 2022.
- [41] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnnp: An accurate o (n) solution to the pnp problem," *International journal of computer vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [42] R. Eidenberger and J. Scharinger, "Active perception and scene modeling by planning with probabilistic 6d object poses," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1036–1043, IEEE, 2010.
- [43] A. Collet and S. S. Srinivasa, "Efficient multi-view object recognition and full pose estimation," in *2010 IEEE International Conference on Robotics and Automation*, pp. 2050–2055, IEEE, 2010.
- [44] J. N. Kundu, M. Rahul, A. Ganeshan, and R. V. Babu, "Object pose estimation from monocular image using multi-view keypoint correspondence," in *European Conference on Computer Vision*, pp. 298–313, Springer, 2018.
- [45] Y. Wu, Y. Zhang, D. Zhu, Y. Feng, S. Coleman, and D. Kerr, "Eao-slam: Monocular semi-dense object slam based on ensemble data association," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4966–4973, IEEE, 2020.
- [46] R. Y. Tsai, R. K. Lenz, *et al.*, "A new technique for fully autonomous and efficient 3 d robotics hand/eye calibration," *IEEE Transactions on robotics and automation*, vol. 5, no. 3, pp. 345–358, 1989.
- [47] G. Klein and D. Murray, "Parallel tracking and mapping for small ar workspaces," in *2007 6th IEEE and ACM international symposium on mixed and augmented reality*, pp. 225–234, IEEE, 2007.
- [48] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: a versatile and accurate monocular slam system," *IEEE transactions on robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [49] G. Jocher, "ultralytics/yolov5: v3.1 - Bug Fixes and Performance Improvements." <https://github.com/ultralytics/yolov5>, Oct. 2020.
- [50] M.-Y. Liu, O. Tuzel, A. Veeraraghavan, Y. Taguchi, T. K. Marks, and R. Chellappa, "Fast object localization and pose estimation in heavy clutter for robotic bin picking," *The International Journal of Robotics Research*, vol. 31, no. 8, pp. 951–973, 2012.
- [51] B. O. Community, *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.
- [52] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [53] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes," in *Asian conference on computer vision*, pp. 548–562, Springer, 2012.