

# NeRF-Pose: A First-Reconstruct-Then-Regress Approach for Weakly-supervised 6D Object Pose Estimation

Fu Li<sup>1,2\*</sup> Shishir Reddy Vutukur<sup>2,3\*</sup> Hao Yu<sup>2</sup> Ivan Shugurov<sup>2,3</sup> Benjamin Busam<sup>2</sup>  
 Shaowu Yang<sup>1</sup> Slobodan Ilic<sup>2,3</sup>  
<sup>1</sup> National University of Defense Technology <sup>2</sup> Technical University of Munich  
<sup>3</sup> Siemens AG, Munich.

## Abstract

Pose estimation of 3D objects in monocular images is a fundamental and long-standing problem in computer vision. Existing deep learning approaches for 6D pose estimation typically rely on the availability of 3D object models and 6D pose annotations. However, precise annotation of 6D poses in real data is intricate, time-consuming and not scalable, while synthetic data scales well but lacks realism. To avoid these problems, we present a weakly-supervised reconstruction-based pipeline, named **NeRF-Pose**, which needs only 2D bounding boxes and relative camera poses during training. Following the first-reconstruct-then-regress idea, we first reconstruct the objects from multiple views in the form of an implicit neural representation. Then, we train a pose regression network to predict pixel-wise 2D-3D correspondences between images and the reconstructed model. A NeRF-enabled PnP+RANSAC algorithm is used to estimate stable and accurate pose from the predicted correspondences. Experiments on LineMod and LineMod-Occlusion show that the proposed method has state-of-the-art accuracy in comparison to the best 6D pose estimation methods in spite of being trained only with weak labels. We extend the Homebrewed DB dataset with real training images to support the weakly supervised task and achieve compelling results. The extended dataset and code will be released soon.

## 1. Introduction

Several computer vision tasks, such as 2D object detection and semantic segmentation, have experienced tremendous progress in recent years thanks to the development of deep learning. However, 2D object detection [12, 13, 26, 32, 38, 40] alone is limited and insufficient for real-world applications such as Augmented Reality, Robotic Manipulation, Autonomous driving, etc., which often require the

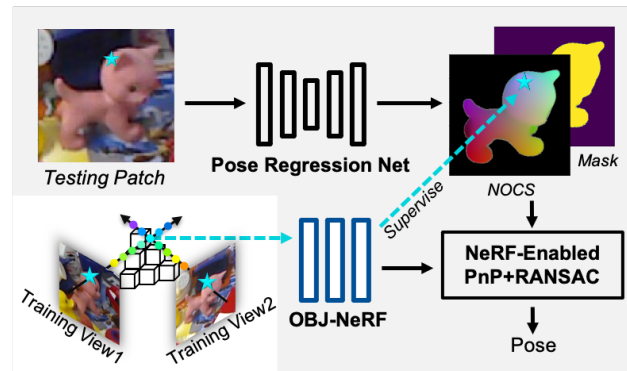


Figure 1. The inference of **NeRF-Pose**. 1) The pose regression network predicts a binary segmentation mask and 2D-3D correspondences represented as *NOCS* maps. The *g.t.* *NOCS* maps are generated from the OBJ-NeRF network, which encodes implicit 3D object representation recovered from the multi-view weak labels, *i.e.*, 2D object segmentation masks and relative camera poses. 2) The proposed NeRF-enabled PnP+RANSAC is performed on the regressed results to predict the object pose.

knowledge of the full 6 degrees of freedom (DoF) pose of the object in the scene. Therefore, the ability to recover the pose of the objects in 3D environments is essential for a better understanding of the 3D scene.

This fundamental 3D vision problem has been addressed by scholars for decades and is particularly difficult if the estimation is done only from a single RGB image. Recent research tackles the ill-posed nature of this problem in a data-driven way where the pose is usually computed with respect to a 3D CAD model. Most of the recent approaches [23, 29, 30, 36, 37, 47, 49, 59, 63, 25, 53, 7] require 6D pose labels as supervision signals. Moreover, most of the recent methods posed the pose estimation as a problem of correspondence estimation between a known 3D object model, such as a CAD model, and image pixels. However, it is hard and expensive to obtain accurate pose labels and fine-grained CAD models for all objects for real-world scenarios [22]. On the other hand, synthetically

\*The first two authors contributed equally to this work

generated images have the advantage of a potentially unlimited amount of labeled data. However, methods trained only on synthetic data [23, 48, 63, 4, 21] have worse performance than their real counterparts due to the lack of realism. Moreover, precise textured CAD models are required to render synthetic data. We argue that it is much easier to obtain 2D image labels, such as segmentation masks, and relative camera poses. As widely used in the recent methods [23, 29, 30, 36, 37, 47, 49, 59, 63, 25, 53, 7], segmentation mask can be obtained either manually or automatically with the off-the-shelf object segmentation approaches [24, 13, 6], few-shot segmentation[56, 34], depth driven segmentation or background subtraction[11], while relative camera poses can be obtained with Structure from Motion (SfM) [42], Simultaneous Localization And Mapping (SLAM) [1], Inertial Visual Odometry [28], or simply a marker board [22]. Different from fully-supervised methods trained on all labels, *e.g.* CAD models, segmentation masks and 6D pose labels, the methods, that don't use textured CAD models and 6D pose annotations, use weaker labels and, so, can be considered as weakly-supervised.

Our approach is to first recover implicit 3D object representation from training images containing weak labels: 2D segmentation masks and relative camera poses. Next, we use this implicit representation to supervise the regression of dense correspondences between training images and previously recovered object's implicit 3D object representation. Therefore, we propose a *first-reconstruct-then-regress* training pipeline, named **NeRF-Pose**, which builds atop of the success of Neural Radiance Fields (NeRF) [33] and its successors [31, 62, 39, 60]. We first reconstruct the object as a NeRF-based network trained with weak labels. Then, we train a pose regression network to regress the dense image pixel (2D)-object model (3D) correspondences.

As depicted in Fig. 1, during inference, we first detect the objects in a 2D image using an off-the-shelf 2D object detection network and then predict segmentation masks and dense correspondences represented in terms of Normalized Object Coordinates (NOCS) [54]. With regressed correspondences and the NeRF object model, we propose a NeRF-enabled PnP+RANSAC method in order to compute the object pose in the end. Our key contributions can be summarized as follows:

- A weakly-supervised object pose estimation approach, which is trained only with 2D annotations and relative camera poses, instead of relying on an explicit CAD model and accurate 6D pose labels.
- OBJ-NeRF neural network encoding an implicit 3D object representation obtained from weak labels: segmentation masks and relative camera poses.
- Pose regression network, which relies on the above object's implicit NeRF-based representation.

- NeRF-enabled PnP+RANSAC approach enabling highly accurate 6D pose computation.
- An extension of the HomebrewedDB dataset containing real video sequences with weak labels (segmentation masks and relative camera poses).

We conduct experiments on LineMod (LM) [15] and LineMod Occlusion (LMO)[3] datasets, and extend the HomebrewedDB (HBD)[22] dataset. Though training in weakly supervised settings, we have about 15% (LM) and 20% (LMO) improvement on ADD(-S) metric compared to the methods trained without CAD models but with accurate pose labels. We also achieve comparable results on LM, LMO, and HBD datasets in comparison to fully supervised methods. The experiments show that our weakly-supervised NeRF-Pose approach achieves accurate and robust object pose estimation.

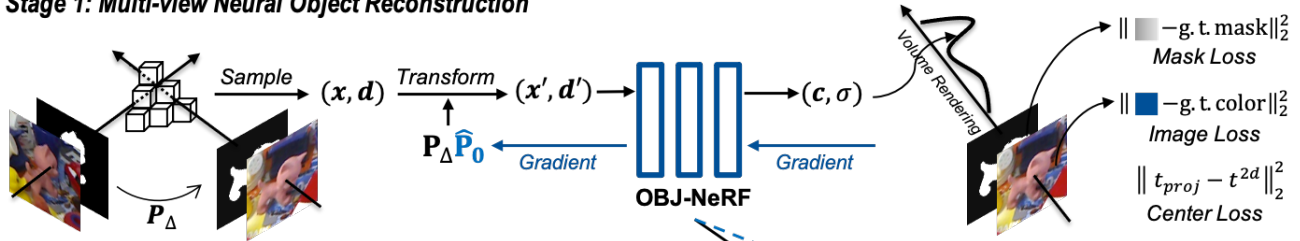
## 2. Related Work

The first type of pose estimation methods that gained popularity in recent years are dense correspondence-based methods [63, 16, 55, 30, 35, 45, 44]. While being different in implementation, their common denominator is the key idea to train a neural network to predict 2D-3D correspondences between each object pixel in the image and the 3D location of the corresponding point on the object's surface. Those correspondences are consecutively used either with PnP+RANSAC [27, 10] or the Umeyama algorithm [50] to compute the 6D object pose. DPOD [63, 45] is proposed to use discrete UV maps to uniquely parameterize the object surface. With this parameterization, the UNet-like network [41] predicts two discrete UV coordinates for each visible pixel occupied by the object. Pix2Pose [35] and CDPN [30] leverage two-stage detectors and use 3D normalized vertex coordinates to parameterize the correspondences. Correspondences in CDPN are only used to estimate rotation, while translation there is predicted directly. NOCS [55] uses Mask-RCNN [13] and a normalized object coordinate space to predict correspondences. However, NOCS mainly focuses on estimating the scale and 6D pose for unseen objects.

Furthermore, CosyPose [25] and Self6D [52] use a similar pose parameterization which allows for direct pose prediction. They predict the translation by estimating the 2D center location of the object and its distance to the camera center. Allocentric rotation parameterization is used for simpler rotation prediction.

The methods above are all fully supervised with the accurate pose labels and object CAD models available. The following methods relax the constraints and could be trained without object CAD models. LieNet [8] directly regresses the pose with the Mask-RCNN [13] as the backbone network. With known object pose labels and cam-

### Stage 1: Multi-view Neural Object Reconstruction



### Stage 2: Single view Pose Estimation

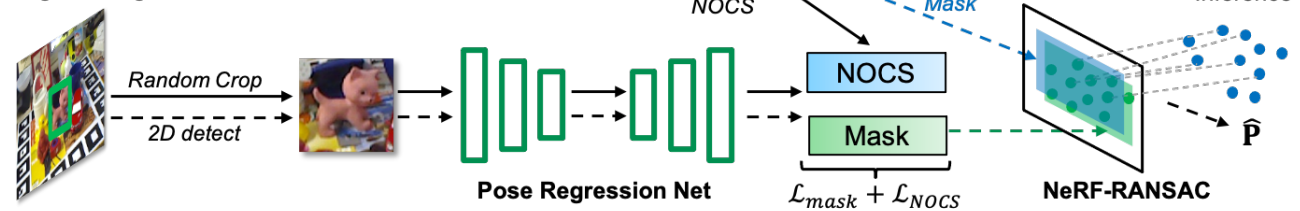


Figure 2. NeRF-Pose Pipeline. *Stage 1: Multi-view Neural Object Reconstruction.* During training, we first reconstruct the neural object model from multiple views. The rays generated from different viewpoints are sampled to produce 3D points along the rays (top left). The sampled points and directions  $(\mathbf{x}, \mathbf{d})$  are then transformed into an object-centric coordinate system according to the estimated object pose  $\hat{\mathbf{P}}_0$  and relative camera pose  $\mathbf{P}_\Delta$ . The transformed 3D points and directions  $(\mathbf{x}', \mathbf{d}')$  are passed through the learnable NeRF model. We can get the image RGB values and the mask with the NeRF rendering process. The image rendering loss and segmentation mask loss are used to guide the learning of the neural object model. *Stage 2: Single View Object Pose Estimation.* In stage two, we directly regress the object coordinate map and the segmentation mask supervised by the rendering results from the neural object model. At inference, we first detect and crop the objects from 2D images (lower left). The coordinates and segmentation mask are predicted from the pose regression network. In the end, the NeRF-enabled PnP+RANSAC is performed to recover the precise object pose in the canonical object space defined at the reconstruction stage.

era intrinsics, Cai et al. [5] supervise the coordinate prediction network with multi-view consistency, minimizing re-projection error across different views. Limited by the accuracy of the coordinate prediction network, the re-projection error is insufficient to guide reliable correspondence learning. Bundle-SDF [58] performs 6D tracking and reconstructs an object without assuming a CAD model from a video. WeLSA [51] generates pose labels for weakly labeled data by training on very few labeled data using shape alignment and feature alignment. Although WeLSA works with very few labeled data, they employ depth maps for training the pipeline and estimating pose labels for weakly labeled data.

Concerning weakly supervised methods, there are few works related to ours. On the basis of NeRF[33], i-NeRF [61] presents a differentiable camera pose refinement method, which treats the camera pose as network parameters and iteratively updates the pose by minimizing the discrepancy between the input images and the rendered outputs in particular views. However, because of the high computational burden, it takes about half a minute to process one image and is very sensitive to pose initialization. BARF [31] and NeRF-[57] introduce the methods to estimate camera poses and train NeRF concurrently. It inspires us to optimize the object pose, while simultaneously reconstructing

the object with known relative poses.

Different from the model-free approach [5], which optimizes the network via minimizing the re-projection error, we aim at generating 2D-3D correspondences between input images and an implicit NeRF object representation obtained from weak labels, yielding better performance.

## 3. Methods

In this section, we present NeRF-Pose for 6D pose estimation, which only requires weak supervision. We assume that real images with 2D segmentation masks and relative camera poses are available during training. We first present OBJ-NeRF, an implicit 3D model representation, learned under the above-defined constraints. This representation is then used to generate object correspondence maps, which are afterward used to train our proposed pose regression network. Finally, the regressed correspondences are used in a novel NeRF-enabled PnP+RANSAC algorithm for iterative pose estimation.

### 3.1. OBJ-NeRF

NeRF [33] and its follow-ups recover the 3D scene from multiple views with known camera poses. Since we deal with the problem of 6D object pose estimation, our aim is to compute object-centric NeRF and use it as an implicit 3D

model representation for object pose estimation. Thus, we propose to modify the original NeRF approach. We take as input the images, segmentation masks and relative camera poses, and output implicit object-specific NeRF representation, named OBJ-NeRF. As OBJ-NeRF reconstruction relies only on relative camera poses, it will produce a 3D model in some uncertain coordinate systems. In order to make use of OBJ-NeRF for the purpose of 6D pose estimation based on dense 2D-3D correspondences it is necessary to recover it with respect to some reference coordinate systems. To achieve this, simultaneously with the NeRF reconstruction, we propose to regress poses of the NeRF reconstructed object with respect to some chosen reference frames. These estimated poses will be used later for the generation of correspondence maps needed for training the correspondence estimation network.

NeRF encodes a 3D scene as a continuous 3D representation using an MLP function  $\mathcal{N}_\psi: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3 \times \mathbb{R}$  with  $(\mathbf{x}; \mathbf{d}) \mapsto (\mathbf{c}; \sigma)$ , which predicts the RGB color  $\mathbf{c} \in \mathbb{R}^3$  and volume density  $\sigma \in \mathbb{R}$  for each input 3D point  $\mathbf{x} \in \mathbb{R}^3$  and its view direction  $\mathbf{d} \in \mathbb{R}^3$ . It can be summarized as  $\mathcal{N}_\psi(\mathbf{x}, \mathbf{d}) = (\mathbf{c}, \sigma)$ , where  $\psi$  is the set of network parameters.

Based on the NeRF scene representation, OBJ-NeRF reconstructs the object from multiple images, and optimizes the reference object pose at the same time. It takes images  $I$ , object masks  $\mathbf{M}$ , camera intrinsics  $\mathbf{K}$ , relative poses  $\mathbf{P}_\Delta$  as training input and outputs the rendered object view  $\hat{I}$ , its mask  $\hat{\mathbf{M}}$ , correspondence map  $\hat{\mathbf{O}}$  as well as the reference object pose  $\mathbf{P}_0$  in respect to the reference image  $I_0$ . As shown in stage 1 in Fig. 2, the steps of the full procedure, from input images to rendering the outputs, can be named as *Sample points*, *Transform points*, *Calculate*, and *Render*.

**Sample points.** Given 2D pixel coordinates  $\mathbf{u} \in \mathbb{R}^2$  in an input image  $I$ , we can express a 3D point  $\mathbf{x}$  along the viewing ray at depth  $z$  as  $\mathbf{x} = \mathcal{K}^{-1}(\mathbf{u}, z, \mathbf{K})$ , where  $z \in [z_{near}, z_{far}]$  is the depth sampling interval and  $\mathcal{K}^{-1}$  is the inverse projection function from image plane to 3D space. Different from the scene reconstruction, to reconstruct the object-centric model, we limit the ray sampling interval to be close to the estimated object center  $\hat{t}^{3d}$ , denoted as  $(z_{near}, z_{far}) = (|\hat{t}^{3d}| - s, |\hat{t}^{3d}| + s)$ , where  $|\hat{t}^{3d}|$  is the distance from estimated object center to the camera center, and  $s \in \mathbb{R}$  is larger than the object scale. Uniform sampling or weight guided re-sampling can be applied here to generate  $N$  points along each ray.

**Transform points.** Different from scene-centric BARF [31] and iNeRF [61], which optimize absolute camera poses in some arbitrary coordinate systems, we perform object-centric reconstruction, which instead estimates the object poses in the camera coordinate system. Therefore, the sampled points  $(\mathbf{x}, \mathbf{d})$  from the camera coordinate system are transformed to the object coordinate system  $(\mathbf{x}', \mathbf{d}')$

according to the object pose  $\hat{\mathbf{P}}_i$  with respect to the image  $I_i$ . However, poses of image objects are unavailable and thus need to be estimated. Notably, owing to our weakly-supervised settings, where relative camera poses from  $I_0$  to  $I_i$ , denoted by  $\mathbf{P}_{\Delta i0}$ , are accessible, only  $\mathbf{P}_0$  needs to be optimized. Thus, arbitrary poses can then be computed by  $\mathbf{P}_i = \mathbf{P}_{\Delta i0} \mathbf{P}_0$ . We further define the transformation functions:  $\mathcal{W}_{\mathbf{P}_0}^{\mathbf{x}}(\mathbf{x}, \mathbf{P}_{\Delta i0}) \rightarrow \mathbf{x}'$  and  $\mathcal{W}_{\mathbf{P}_0}^{\mathbf{d}}(\mathbf{d}, \mathbf{P}_{\Delta i0}) \rightarrow \mathbf{d}'$ , which transform the coordinate  $\mathbf{x}$  and view direction  $\mathbf{d}$  into unified object-centric coordinates, with the parameters  $\mathbf{P}_0$  to be optimized.

**Calculate.** Until now, we have obtained the transformed ray points and their directions  $(\mathbf{x}', \mathbf{d}')$ . Then, we feed  $(\mathbf{x}', \mathbf{d}')$  into the original NeRF network  $\mathcal{N}_\psi$  to get color and volume density prediction  $(\mathbf{c}, \sigma)$ . In summary, our OBJ-NeRF network  $\mathcal{N}_{\psi, \mathbf{P}_0}$  can be formulated as:

$$\mathcal{N}_{\psi, \mathbf{P}_0}(\mathbf{x}', \mathbf{d}') = \mathcal{N}_{\psi, \mathbf{P}_0}(\mathcal{W}^{\mathbf{x}}(\mathbf{x}, \mathbf{P}_i), \mathcal{W}^{\mathbf{d}}(\mathbf{d}, \mathbf{P}_i)) = (\mathbf{c}, \sigma), \quad (1)$$

where  $\psi$  and  $\mathbf{P}_0$  are the network parameters.

**Render.** Akin to NeRF, we use volume rendering approach to map a set of calculated data  $(\mathbf{c}, \sigma)$  to the image plane along the rays. The volume rendering function sums up the product of transmittance  $T_i$  and alpha value  $\alpha_i$  of  $N$  sampled points along the ray, which is differentiable. Following [33, 60], the rendered color  $\hat{\mathbf{C}}$ , mask  $\hat{\mathbf{M}}$  and coordinates  $\hat{\mathbf{O}}$  can be formulated as:

$$\hat{\mathbf{C}}(\mathbf{u}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \hat{\mathbf{M}}(\mathbf{u}) = \sum_{i=1}^N T_i \alpha_i, \hat{\mathbf{O}}(\mathbf{u}) = \sum_{i=1}^N T_i \alpha_i \mathbf{x}_i, \quad (2)$$

with  $T_i = \exp(-\sum_{1 \leq j \leq i-1} \sigma_j \delta_j)$ , and  $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$ , where  $\delta_j$  is the sampling distance between sampled adjacent points.

**Constrain object pose.** With only a 2D bounding box and segmentation, where object center information is unavailable, the object canonical pose center has no constraint. Accordingly, we constrain the object center by projecting the object center  $t^{3d}$  to each view. More specifically, we minimize the reprojection error between the projected object center and 2D bounding box center  $t^{2d}$  on the image plane by minimizing  $e = \|\mathcal{K}(\mathcal{W}_{\mathbf{P}_0}^{\mathbf{x}}(t_0^{3d}, \mathbf{P}_\Delta), \mathbf{K}) - t^{2d}\|_2^2$ .

**Loss function.** We leverage the 2D segmentation mask and the segmented image as the supervision signals for the rendered pixel mask and color value. Denoting the segmentation mask at pixel  $\mathbf{u}$  as  $\mathbf{M}(\mathbf{u}) \in \{0, 1\}$ , and the color value as  $\mathbf{C}(\mathbf{u}) \in \mathbb{R}^3$ , our loss function for the reconstruction stage can be defined as:

$$\begin{aligned} \mathcal{L}_{rec} = & \sum_i \left( \sum_{\mathbf{u}} \mathbf{M}(\mathbf{u}) \|\mathbf{C}(\mathbf{u}) - \hat{\mathbf{C}}(\mathbf{u})\|_2^2 \right. \\ & + \sum_{\mathbf{u}} \|\mathbf{M}(\mathbf{u}) - \hat{\mathbf{M}}(\mathbf{u})\|_2^2 \\ & \left. + \|\mathcal{K}(\mathcal{W}_{\mathbf{P}_0}^{\mathbf{x}}(t_0^{3d}, \mathbf{P}_{\Delta i0}), \mathbf{K}) - t_i^{2d}\|_2^2 \right), \end{aligned} \quad (3)$$

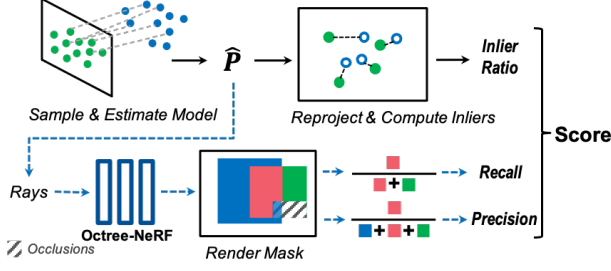


Figure 3. Illustration of a single iteration of the NeRF-enabled RANSAC. Similar to a standard PnP+RANSAC procedure (upper top), the pose hypothesis is first calculated by the PnP algorithm from sampled 2D-3D correspondences. Scored by Inlier Ratio, the pose hypothesis with the most inliers is selected as the estimated result. Different from standard PnP+RANSAC, we further assemble Recall and Precision into our criteria of pose selection (bottom). Detailly, we first render the object mask in this pose hypothesis using our well-trained octree-NeRF. Then, we calculate Precision by the Intersection (red)-over-Union (blue+red+green) between the rendered mask and the mask from our pose regression network. The recall is defined as the ratio of the intersection (red) over the area of the regressed mask (red+green). In the end, the selection score is calculated as the weighted sum of Inlier Ratio, Precision, and Recall.

where  $i$  iterates the training image views, and  $\mathbf{u}$  iterates the image pixels.

In the end, the newly defined object canonical pose is determined by optimizing the object pose in multi-view settings. With the learned canonical pose and neural model, the *NOCS* map ground truth can be rendered for the next pose estimation stage.

### 3.2. Pose Estimation

The overall three-step pipeline is shown in Stage 2 in Fig. 2. The first two steps depend on separately trained convolutional neural networks. The third step is purely optimization-based and does not require training. The first step represents the off-the-shelf 2D detector trained on ground truth crops of the objects of interest. In practice, we use YOLOv3 [9]. In the second step, a pose regression network is trained on images  $I$  to predict object coordinates  $\mathbf{O}$  and segmentation mask  $\mathbf{M}$ . The third step is our NeRF-enabled PnP+RANSAC algorithm, to further improve pose estimation by introducing the NeRF-Mask renderer.

**Coordinates regression.** Our pose regression network is inspired by DPoD [63] and CDPN [30], the state-of-the-art dense correspondence-based methods for indirect pose estimation. We use ResNet[14] as the encoder backbone, and the decoder contains four upsampling layers. Our pose regression network outputs the predicted segmentation  $\hat{\mathbf{M}}$  and the object coordinates  $\hat{\mathbf{O}}$ , which encodes correspondences between input image pixels and the OBJ-NeRF 3D model representation.

The loss function is defined similarly to other depth regression works [2, 19]. Our loss function is composed of the segmentation loss and the normalized coordinate regression loss:

$$\mathcal{L} = w_1 \mathcal{L}_M + w_2 \mathcal{L}_{NOCS}. \quad (4)$$

The  $\mathcal{L}_M$  loss is defined by the mean squared error (MSE) between predicted mask  $\hat{\mathbf{M}}$  and ground truth mask  $\mathbf{M}$ .  $\mathcal{L}_{NOCS}$  is defined between the predicted object *NOCS* map  $\hat{\mathbf{O}}$  and their ground truth  $\mathbf{O}$  rendered from the implicit neural network OBJ-NeRF as given below:

$$\mathcal{L}_{NOCS} = \lambda_1 \mathcal{L}_{mse} + \lambda_2 \mathcal{L}_{grad} + \lambda_3 \mathcal{L}_{normal}, \quad (5)$$

The first term  $\mathcal{L}_{mse}$  is the element-wise loss that measures the average distance between these two coordinates written as:

$$\mathcal{L}_d = \|\mathbf{M} \odot (\mathbf{O} - \hat{\mathbf{O}})\|_2^2. \quad (6)$$

$\mathcal{L}_{grad}$  and  $\mathcal{L}_{normal}$  are used to penalize the coordinate errors in first and second order. Similar to the loss of [19], they are defined as:

$$\mathcal{L}_{grad} = \|\mathbf{M} \odot (\nabla_x \hat{\mathbf{O}} - \nabla_x \mathbf{O})\|_2^2 + \|\mathbf{M} \odot (\nabla_y \hat{\mathbf{O}} - \nabla_y \mathbf{O})\|_2^2, \quad (7)$$

$$\mathcal{L}_{normal} = \|\mathbf{M} \odot (1 - \cos(n(\hat{\mathbf{O}}), n(\mathbf{O}))\|_2^2, \quad (8)$$

where  $\nabla_x$  and  $\nabla_y$  are the gradients for the coordinates along  $x$  and  $y$  axis,  $n(o) = [\nabla_x(o), \nabla_y(o), -1]$  defines the surface normal vector and  $\cos(n_1, n_2)$  measures the cosine similarity of the two vectors.

**NeRF-enabled PnP+RANSAC.** With the predicted dense 2D-3D correspondences, PnP+RANSAC-based method [63, 30, 16] are typically used for estimating object pose  $\hat{\mathbf{P}}$ . As shown in Fig. 3, the PnP+RANSAC [10] algorithm iteratively selects a minimum number of correspondences for pose estimation and calculates object pose using PnP method. The pose hypothesis supporting the most inliers is selected as the calculated pose results.

The criteria of pose selection in PnP+RANSAC is inlier ratio  $S_{IR}$ , which highly depends on the correspondence quality. To make PnP+RANSAC more stable, we add two extra pose selection criteria which are independent to the correspondences. As shown in Fig. 3, with the calculated pose hypothesis and our implicit object representation, we can render the object mask  $\hat{\mathbf{M}}_{nerf}$  on the image plane without occlusion. With the mask  $\hat{\mathbf{M}}_{pose}$  predicted from our pose regression network, we add the *Overlap Recall*  $S_{Recall}$  and *Overlap Precision*  $S_{Prec}$  to the score which is defined as:  $S_{Recall} = |\hat{\mathbf{M}}_{nerf} \cap \hat{\mathbf{M}}_{pose}| / |\hat{\mathbf{M}}_{pose}|$  and  $S_{Prec} = |\hat{\mathbf{M}}_{nerf} \cap \hat{\mathbf{M}}_{pose}| / |\hat{\mathbf{M}}_{pose} \cup \hat{\mathbf{M}}_{nerf}|$ . The final score is defined as  $S = v_1 S_{IR} + v_2 S_{Prec} + v_3 S_{Recall}$ .

Following Plenotree [62], we convert the well-trained OBJ-NeRF model to an octree-based data structure, namely *octree-NeRF*, which stores the average value of  $\mathbf{c}, \sigma$  sampled in the leaf voxel of the octree. The NeRF calculation is simplified by the octree indexing, which significantly accelerates the mask rendering.




Figure 4. Qualitative results on **LM**, **HBD**. Column (a) contains reference ground truth CAD models. Column (b) and (f) illustrate the predicted object poses (blue) and ground truth poses (green). Column (c) and (g) are the rendered images from our well-trained OBJ-NeRF model, Column (d) and (h) are the produced *NOCS* map from our OBJ-NeRF, which are also *g.t.* of the regression network. Column (e) and (i) are the predicted *NOCS* from our regression network. Black dash circles show the *NOCS* prediction of occlusion regions.

## 4. Experiments

### 4.1. Datasets

We conduct extensive experiments to demonstrate that our proposed weakly-supervised pose estimation method, **NeRF-Pose**, produces highly-accurate 6D object pose.

**Datasets.** We conduct our experiments on three publicly available datasets: Linemod (LM)[15], Linemod-Occlusion (LMO)[3] and HomebrewedDB (HBD)[22]. The LM dataset is a standard benchmark for 6D object pose estimation of textureless objects. It offers 13 objects of various sizes in the scenes with large background clutter but with almost no occlusion. The LMO dataset consists of 8 objects from LM dataset but provides more challenging test data with more occlusion. According to our problem settings, we train our model on the real data in LM and strictly follow the training/testing split proposed in [3]. On the LMO dataset, to make a fair comparison to Cai et al. [5], we train our method with the real training data of LM. Since the HBD dataset has no real training data, we provide an **Real data EXTension** of the HBD dataset, named **HBD-REXT**. It consists of real images of each HBD object presented in BOP [18] challenge. These images are captured by the Azure Kinect camera with only relative poses and segmentation mask provided. We will make the HBD-

REXT dataset publicly available soon.

**Evaluation Metrics.** On the LM and LMO datasets, we report the standard ADD(-S) metrics [15] with the 10% diameter threshold, as it is the most prevalent pose quality metric for these two datasets. The  $n^\circ, n\text{ cm}$  metric [43] measures whether the rotation error is less than  $n^\circ$  and the translation error is below  $n\text{ cm}$ . Besides, we also use  $n^\circ$  and  $n\text{ cm}$  [43], which measure whether the rotation error is less than  $n^\circ$  and the translation error is below  $n\text{ cm}$ , respectively. Moreover, following the setup of the BOP challenge, which aims at unifying the evaluation of 6D pose estimation methods, we report the Average Recall, AR, for the HBD dataset

Since we predict the poses in a canonical orientation, it is necessary to transform them to absolute poses. This is done by transforming them with the offset poses (obtained from the *g.t.* poses), which represent the transformation from absolute poses to the canonical poses. Note that this transformation is only leveraged for evaluation.

### 4.2. Comparisons to the State-of-the-art

**LM.** We train our method on real images, strictly following the training/testing split from [3]. In order to make a fair comparison to Cai et al. [5], apart from training using relative poses(**Our-weak**), we also train our method using pose labels(**Our-pose**). In Tab. 1, our method outperforms

Table 1. **LM** results in on ADD-10 metric. \*denotes that the objects is symmetric and is evaluated in ADD-S. **Our-pose** denotes our results trained on 6D pose labels, and **Our-weak** denotes training on camera relative pose labels. **Ours-sam** denotes our results trained on 6D pose labels and segmentation masks extracted using SegmentAnything. *w/o NeRF* denotes our results using original PnP+RANSAC and *w/NeRF* is our method with our NeRF-enabled PnP+RANSAC

Object	Tekin	PVNet	CDPN	GDR	SO-Pose	LieNet	Cai.	<b>Ours-sam</b>	<b>Ours-pose</b>	<b>Ours-pose</b>	<b>Ours-weak</b>
	[49]	[36]	[30]	[53]	[7]	[8]	[5]	<i>w/o NeRF</i>	<i>w/o NeRF</i>	<i>w/ NeRF</i>	<i>w/ NeRF</i>
<i>CAD</i>	<i>w/ CAD</i>						<i>w/o CAD</i>				
Mean	56.0	86.3	89.9	93.7	96.0	65.2	82.9	88.3(↑ 5.4)	91.8(↑ 8.9)	96.6(↑ 13.7)	<b>97.8(↑ 14.9)</b>

Table 2. Comparisons with state-of-the-art methods on LMO. We report the Average Recall(%) of ADD(-S) without refinement. *real* denotes the same real data as LM. *syn* denotes self-generated synthetic data, and *pbr* denotes blender rendered synthetic data from BOP[17]. \* denotes the symmetric objects. **Our-pose** denotes our results with accurate pose labels and **Our-weak** is with relative pose labels. *w/o NeRF* denotes our results using original PnP+RANSAC and *w/ NeRF* is our method with our NeRF-enabled PnP+RANSAC.

	PoseCNN	PVNet	Single-Stage	HybridPose	GDR	SO-Net	GDR	SO-Net	Cai.	<b>Our-pose</b>	<b>Our-pose</b>	<b>Our-weak</b>	<b>Our-weak</b>
	[59]	[36]	[20]	[46]	[53]	[7]	[53]	[7]	[5]	<i>w/o NeRF</i>	<i>w/ NeRF</i>	<i>w/o NeRF</i>	<i>w/ NeRF</i>
<i>CAD</i>	<i>w/ CAD</i>						<i>w/ CAD</i>			<i>w/o CAD</i>			
training	<i>real+syn</i>						<i>real+pbr</i>			<i>real</i>			
Mean	24.9	40.8	43.3	47.5	53.0	54.3	62.2	<b>62.3</b>	30.3	44.1	49.2	48.2	51.4(↑ 21.1)

the method of Cai at al. [5] by a large margin ( $\approx 14pp$ ) on a ADD(-S) metric. Moreover, our method with weakly-supervised settings performs on-par with the SoTA fully-supervised methods. We also perform an ablation, (**Ours-sam**) by using SegmentAnything[24] to generate segmentation masks using ground truth bounding boxes as input instead of using ground truth segmentation masks. The mild accuracy drop(3.5%) indicates that our approach can be applied in the real world much more easily compared to other approaches using relative poses from the sensor or SFM and segmentation masks obtained using SAM.

**LMO.** We compare our method with SoTA methods in terms of ADD(-S) on LMO. Compared to Cai et al. [5] trained with pose labels but without CAD models on real images, our method on the same settings achieves 49.2% on mean ADD score, surpassing it with a large margin of 20pp. In our setting, the CAD models are not accessible, so we do not train our method in self-generated synthetic images or *pbr* images from BOP. It is proved that training with more synthetic images can improve the model performance and *pbr* images can improve more. So, it is reasonable to claim that we stay comparable with the SoTA fully-supervised methods (GDR-Net [53]: 53.0% and SO-Pose [7]: 54.3%) trained on real and author-generated synthetic images.

In Tab.1 and Tab.2, **Our-weak**(uses relative poses) performs better than **Our-pose** (uses absolute poses). This is caused by noisy labels in LM, which harm more **Our-pose** than **Our-weak**. This is because in **Our-weak** we optimize all pose labels by minimizing the reprojection and rendering errors constrained with relative poses. In **Our-pose** with absolute poses we cannot refine them, since we cannot guar-

antee to obtain the correct object’s scale.

**HBD-REXT.** To support weakly-supervised training on real images, we extend the HomebrewedDB training set by capturing real sequences for all objects used in BOP [18] challenge. We provide 300 real images for each object with segmentation masks and relative camera poses generated from the markerboard. We train our model on this newly captured data. Since no other method was trained only with relative poses and 2D segmentation, we cannot run other methods using this new weakly-supervised data. In Tab. 3 we report the AR metric on the BOP challenge test set. It illustrates that we stay on-par with the methods trained on synthetic images [63, 20, 30] and fall behind the methods trained on *pbr* images [16, 25, 30]. Notably, neither CAD models nor ground truth poses are used in the **our-weak** case, whose results are shown in the last column of Tab. 3.

Table 3. Comparisons with state-of-the-art methods on **HBD**.

	DPoD	Single	CDPN	Pixel2Pose	EPOS	CosyPose	CDPN	<b>Our-weak</b>
	[63]	[20]	[30]	[35]	[16]	[25]	[30]	<i>w/ NeRF</i>
training	<i>syn</i>	<i>syn</i>	<i>syn</i>	<i>pbr</i>	<i>pbr</i>	<i>pbr</i>	<i>pbr</i>	<i>real-ext</i>
AR	28.6	29.3	47.0	44.6	58.0	65.6	72.2	51.8
Time(s)	0.18	0.19	0.31	0.98	0.66	0.42	0.27	0.25

### 4.3. Ablations

**Shape Analysis.** The visual results of the objects represented in OBJ-NeRF can be found in Fig. 4(c, g). Compared to their reference CAD models (Fig 4(a)), our OBJ-NeRF keeps both the detailed shapes and the texture information. Accurate shapes obtained in our model promises the accuracy of our consecutive pose estimation. We observe minor

defects on our predicted shapes, e.g. details of cat eyes and cow legs, which is mainly caused by the imprecise relative pose annotations and imperfect segmentation masks.

**View Numbers.** To figure out how the number of views used in NeRF reconstruction affects the pose estimation, we train OBJ-NeRF with 4 different numbers of views—32(B1), 64(B2), 128(B3) and 156(Baseline). From the results in Tab. 4, training using 156 views is slightly better than the others. It actually shows small difference (around 0.5%) on ADD metric and almost no influence of the number of views to the results. Thus, we use 156 views for training OBJ-NeRF. Moreover, the results also illustrate the number of views does not bring significant influence on the performance, i.e., the OBJ-NeRF can recover the highly precise object model, though trained with fewer views. It is an interesting conclusion that can inspire us to use few-shot labels for reconstruction and pose estimation.

**Regression Loss.** We study the contribution of our pose regression loss. As shown in Tab. 4 C1-C3, the loss with gradient and normal loss has the best performance, which indicates the efficiency of our utilized loss functions.

**Training with occlusions.** To further evaluate the influence of occlusions on the performance of OBJ-NeRF, we add random rectangles onto RGB images to mask out certain areas and simulate occlusions and errors in segmentation masks. We test the performance on the same objects as in other ablation studies. The reconstruction computed from the distorted images (Fig.5(d)) has no clear difference to the non-distorted one (Fig.5(e)). However, in Tab 4, training with occlusion(D1) has a slight drop (about 5 *pp* on ADD score) on pose estimation performance. We consider it acceptable, as the occlusion increases the difficulty of the reconstruction. Moreover, as indicated by the green ellipses in Fig.5, the *g.t.* masks and pose labels from the LM and HBD are actually noisy. Due to the imprecisions, our reconstructions are missing some details. However, all our experiments demonstrate that such reconstructed implicit object models and small pose errors can be tolerated.

**NeRF-enabled PnP+RANSAC.** Furthermore, we evaluate the effectiveness of our proposed correspondence solver. We present the results in Tab. 4, which shows about 2% improvement on LM dataset on ADD score (E1). Though on 2° metric, the original PnP+RANSAC is slightly better than ours, our method outperforms it on the 2 cm metric. On ADD score, our method shows superiority over the original PnP+RANSAC. The main reason is that our NeRF-enabled RANSAC incorporates mask-based Recall and Precision in scoring pose hypothesis, as illustrated in Fig. 3. Since Recall and Precision are more sensitive to translation, our NeRF enabled RANSAC prefers pose hypothesis with less translation error, leading to better performance on both ADD and 2 cm metrics that reflect translation performance. In Tab. 4 E2-E3, we also report the ADD(-S) score on LMO

Table 4. Ablation study on LM subset. We train the pipeline for the Ape, Cat, and Duck. A0 is the baseline setting in our paper, and B1-B3 shows the results on different numbers of training views in OBJ-NeRF. C1-C3 are the results using different loss combinations. D1 is the result of training with generated occlusions. E1-E3 consists of the results from different PnP+RANSAC strategies.

Case	Method	ADD(0.01d)	2cm	2°,2cm	2°
A0	NeRF-Pose	95.1	97.5	82.3	83.2
B1	rec. 32 views	94.6	97.0	79.6	79.9
B2	rec. 64 views	94.0	96.6	78.4	78.9
B3	rec. 128 views	93.6	96.2	78.5	79
C1	w/o $\mathcal{L}_{grad}$	91.5	95.2	77.2	78.4
C2	w/o $\mathcal{L}_{norm}$	91.0	94.4	76.9	78.0
C3	w/o $\mathcal{L}_{grad}, \mathcal{L}_{grad}$	92.9	96.1	77.7	78.3
D1	rec. w/ occlusion	89.8(↓ 5.3 <i>pp</i> )	93.8	74.4	76.0
E1	PnP+RANSAC	93.6(↓ 1.5 <i>pp</i> )	96.8	83.0	83.8
E2(LMO)	NeRF-PnP+RANSAC	51.4	48.4	9.6	12.1
E3(LMO)	PnP+RANSAC	48.2(↓ 3.2 <i>pp</i> )	45.7	8.7	11.6

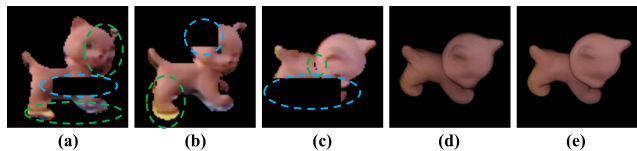


Figure 5. (a),(b) and (c) are masked RGB images (green ellipses are noises and blue ellipses are our generated occlusions), and (d) is rendered with occluded images training, (e) without.

dataset. We observe an average 3% improvement using our proposed NeRF-enabled method. The results show the superiority of our proposed NeRF-enabled PnP-RANSAC algorithm over the original one. The LMO dataset contains more occlusion data than the LM dataset. The occlusions can make mask scores used in NeRF-PnP-Ransac imprecise. In that case, overlap recall and precision can deteriorate. However, since we weight the inlier score the most with 0.6 vs. 0.2 for precision and recall scores, and since our inliers are quite accurate we still get better performance than standard PnP-RANSAC as shown in the Tab.4 for the LM-O benchmark.

## 5. Conclusion

In this paper, we propose **NeRF-Pose**, a first-reconstruct-then-regress approach for weakly-supervised object pose estimation. **NeRF-Pose** first implicitly reconstructs the object as the proposed neural network, namely OBJ-NeRF, from the weak labels and generates the signals to supervise the correspondences predicted from our pose regression network. At inference, a NeRF-enabled PnP+RANSAC algorithm is used to estimate the pose from the predicted correspondences. Finally, A thorough evaluation on LM, LMO, and HDB datasets shows our leading performance on the task of weakly-supervised object pose estimation.



## References

- [1] Tim Bailey and Hugh Durrant-Whyte. Simultaneous localization and mapping (slam): Part ii. *IEEE robotics & automation magazine*, 13(3):108–117, 2006.
- [2] Jonathan T Barron. A general and adaptive robust loss function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4331–4339, 2019.
- [3] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *European conference on computer vision*, pages 536–551. Springer, 2014.
- [4] Benjamin Busam, Hyun Jun Jung, and Nassir Navab. I like to move it: 6d pose estimation as an action decision process. *arXiv preprint arXiv:2009.12678*, 2020.
- [5] Ming Cai and Ian Reid. Reconstruct locally, localize globally: A model free method for object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3153–3163, 2020.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [7] Yan Di, Fabian Manhardt, Gu Wang, Xiangyang Ji, Nassir Navab, and Federico Tombari. So-pose: Exploiting self-occlusion for direct 6d pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12396–12405, 2021.
- [8] T Do, Trung Pham, Ming Cai, and Ian Reid. Real-time monocular object instance 6d pose estimation. 2019.
- [9] Joseph Redmon Ali Farhadi and Joseph Redmon. Yolov3: An incremental improvement. Retrieved September, 17:2018, 2018.
- [10] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [11] Belmar Garcia-Garcia, Thierry Bouwmans, and Alberto Jorge Rosales Silva. Background subtraction in real applications: Challenges, current models and future directions. *Computer Science Review*, 35:100204, 2020.
- [12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Asian conference on computer vision*, pages 548–562. Springer, 2012.
- [16] Tomas Hodan, Daniel Barath, and Jiri Matas. Epos: Estimating 6d pose of objects with symmetries. *arXiv preprint arXiv:2004.00605*, 2020.
- [17] Tomas Hodan and Antonin Melenovsky. Bop: Benchmark for 6d object pose estimation: <https://bop.felk.cvut.cz/home/>, 2019.
- [18] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: Benchmark for 6d object pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018.
- [19] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1043–1051. IEEE, 2019.
- [20] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2930–2939, 2020.
- [21] Roman Kaskman, Ivan Shugurov, Sergey Zakharov, and Slobodan Ilic. 6 dof pose estimation of textureless objects from multiple rgb frames. In *European Conference on Computer Vision*, pages 612–630. Springer, 2020.
- [22] Roman Kaskman, Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Homebreweddb: Rgb-d dataset for 6d pose estimation of 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [23] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1521–1529, 2017.
- [24] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023.
- [25] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [26] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018.
- [27] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009.
- [28] Stefan Leutenegger, Simon Lynen, Michael Bosse, Roland Siegwart, and Paul Furgale. Keyframe-based visual-inertial odometry using nonlinear optimization. *The International Journal of Robotics Research*, 34(3):314–334, 2015.
- [29] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In

- Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–698, 2018.
- [30] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7678–7687, 2019.
- [31] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. *arXiv preprint arXiv:2104.06405*, 2021.
- [32] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [33] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [34] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 622–631, 2019.
- [35] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7668–7677, 2019.
- [36] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019.
- [37] Mahdi Rad and Vincent 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*, pages 3828–3836, 2017.
- [38] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [39] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. *arXiv preprint arXiv:2103.13744*, 2021.
- [40] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [42] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016.
- [43] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2930–2937, 2013.
- [44] Ivan Shugurov, Ivan Pavlov, Sergey Zakharov, and Slobodan Ilic. Multi-view object pose refinement with differentiable renderer. *IEEE Robotics and Automation Letters*, 2021.
- [45] Ivan Shugurov, Sergey Zakharov, and Slobodan Ilic. Dpodv2: Dense correspondence-based 6 dof pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [46] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6d object pose estimation under hybrid representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 431–440, 2020.
- [47] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018.
- [48] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, and Rudolph Triebel. Augmented autoencoders: Implicit 3d orientation learning for 6d object detection. *International Journal of Computer Vision*, 128(3):714–729, 2020.
- [49] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 292–301, 2018.
- [50] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, pages 376–380, 1991.
- [51] Shishir Reddy Vutukur, Ivan Shugurov, Benjamin Busam, Andreas Hutter, and Slobodan Ilic. Welsa: Learning to predict 6d pose from weakly labeled data using shape alignment. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 645–661, Cham, 2022. Springer Nature Switzerland.
- [52] Gu Wang, Fabian Manhardt, Jianzhun Shao, Xiangyang Ji, Nassir Navab, and Federico Tombari. Self6d: Self-supervised monocular 6d object pose estimation. In *The European Conference on Computer Vision (ECCV)*, August 2020.
- [53] Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16611–16621, 2021.
- [54] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J. Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [55] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object co-

- ordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019.
- [56] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9197–9206, 2019.
- [57] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.
- [58] Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Muller, Alex Evans, Dieter Fox, Jan Kautz, and Stan Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. *CVPR*, 2023.
- [59] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017.
- [60] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021.
- [61] Lin Yen-Chen, Pete Florence, Jonathan T Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. inerf: Inverting neural radiance fields for pose estimation. *arXiv preprint arXiv:2012.05877*, 2020.
- [62] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. *arXiv preprint arXiv:2103.14024*, 2021.
- [63] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1941–1950, 2019.