

RDPN6D: Residual-based Dense Point-wise Network for 6DoF Object Pose Estimation Based on RGB-D Images

Zong-Wei Hong Yen-Yang Hung Chu-Song Chen
National Taiwan University

{r10922190, r11922a18}@g.ntu.edu.tw, chusong@csie.ntu.edu.tw

Abstract

In this work, we introduce a novel method for calculating the 6DoF pose of an object using a single RGB-D image. Unlike existing methods that either directly predict objects' poses or rely on sparse keypoints for pose recovery, our approach addresses this challenging task using dense correspondence, i.e., we regress the object coordinates for each visible pixel. Our method leverages existing object detection methods. We incorporate a re-projection mechanism to adjust the camera's intrinsic matrix to accommodate cropping in RGB-D images. Moreover, we transform the 3D object coordinates into a residual representation, which can effectively reduce the output space and yield superior performance. We conducted extensive experiments to validate the efficacy of our approach for 6D pose estimation. Our approach outperforms most previous methods, especially in occlusion scenarios, and demonstrates notable improvements over the state-of-the-art methods. Our code is available on <https://github.com/AI-Application-and-Integration-Lab/RDPN6D>.

1. Introduction

Estimating the pose of every object in an image is an essential task. This paper introduces a simple but effective 6DoF object pose estimation method based on RGB-D images. Object pose estimation is a crucial technique in many applications, such as augmented reality [28], autonomous driving [9, 15], and robotic manipulation [10, 39]. Traditional methods [10, 18, 27] use handcrafted features to establish 2D-3D correspondences between the image and 3D object meshes. However, they face challenges such as sensor noise, variable lighting conditions, and visual ambiguity. Recent methods leverage deep learning using convolutional neural networks (CNNs) or Transformer-based [3] architectures in RGB images to address these challenges. These methods typically provide pixel-level correspondences [11, 19, 25, 41] or predict 2D image locations

for predefined 3D keypoints [4, 31], resulting in more robust results. Nonetheless, they still face challenges such as handling textureless objects and severe occlusion.

To address the issues mentioned above, and because the cost of RGB-D sensors continues to decrease, recent studies such as PVN3D [17], FFB6D [16], RCVPose [42], and DFTr [46] have employed RGB-D images for 6DoF object pose estimation. Given the large output space in 6DoF, existing methods based on RGB-D images predict predefined keypoints in the 3D space (i.e., establishing 3D-3D sparse correspondences). The pose is then solved using least squares fitting methods. Although sparse keypoint methods exhibit essential clues that are explainable for object pose estimation, they still suffer from some drawbacks. First, they often involve time-consuming procedures such as keypoint voting. Second, predicting the keypoints that are heavily occluded remains challenging. Third, they might fail to find object landmarks under viewpoint changes [37].

Another approach in the RGB-D scheme aims to directly predict the object poses from the feature representations generated by the neural networks, such as DenseFusion [40] and ES6D [29]. Uni6D [23] and Uni6Dv2 [38] have also integrated RGB and depth feature extraction for direct pose prediction. The pose-deriving procedure is less explainable because the object pose is regressed directly from the feature embedding. The object poses estimated are often less accurate using the direct pose estimation approaches. This weakness is also observed in a related track of study, Camera Pose Estimation (aka Ego-motion Estimation) [7, 34]. Reasons why pose regression methods are less accurate are provided in [33].

To overcome the limitations of the above-mentioned RGB-D image-based methods, we propose to use dense correspondence in 3D space to mitigate the shortcomings of sparse correspondence and pose regression. This involves the utilization of both dense 2D-3D correspondence and 3D-3D correspondence, where the former employs the RGB-channels and the latter employs the D-channel inputs to establish the correspondences in association with the point clouds of 3D object models. To better com-

bine the RGB and depth information, we modify the approach of ES6D [29]. Our approach generates intermediate feature representations for the RGB images and the positionally encoded depth map (camera xyz map) at first. Then, it uses the concatenated feature maps to produce the final embedding jointly. Instead of directly regressing the pose from the embedding, we use the embedding to bring about the dense correspondence from the joint intermediate representation. This avoids the shortcomings of the pose regression approach, which is more related to pose approximation than to accurate pose estimation. Our dense-correspondence method maintains explainability and can achieve better accuracy than the sparse-correspondence method because complete 2D-3D and 3D-3D correspondence information is used. An overview is given in Fig. 1.

Furthermore, to avoid the need to predict the coordinates in a large and relatively unlimited space when finding the correspondence, we transform the coordinates of the 3D points in the object database into *residual-based* representations. A set of scattered anchor points is selected from the 3D model, and the coordinate of an object point is expressed as the residue vector for its nearest anchor point, as illustrated in Fig. 2. The representation ensures that the prediction range is more condensed and focused. Subsequently, the object pose is estimated using a neural network based on the dense correspondence, which is more efficient than the time-consuming process such as PnP-RANSAC.

To evaluate our proposed residual-based dense point network (RDPN) approach in a broad range of scenarios, we conducted experiments using four benchmark datasets: MP6D [6], YCB-Video [43], LineMOD [18], and Occlusion LineMOD [5]. Our approach generally outperforms most state-of-the-art methods on the four datasets, especially on heavily occluded datasets. Specifically, our approach achieves the best ADD-S AUC on the YCB-Video and MP6D datasets and the best ADD(-S) 0.1d on the LineMOD and Occlusion LineMOD datasets.

2. Related work

Two approaches are mainly used to infer the object poses related to the 3D models stored in the object database: approaches based on RGB images and RGB-D images.

In RGB-based approaches, SSD6D [24] and T6D-Direct [3] adopt object detection networks to infer an additional output for the object pose. YOLOPose [4] uses keypoint matching to derive the rotation and translation. DPOD [45] projects the 3D model points onto a 2-channel texture map and then predicts the coordinates on the corresponding texture map. Pix2Pose [30] uses an encoder-decoder architecture to predict the 3D coordinates and also uses a GAN to improve stability. CDPN [25] predicts correspondences in two steps and directly predicts translations from images instead of relying on the results of PnP-RANSAC. To address




Figure 1. **Overview of our approach.** Our method predicts the 3D coordinates of each pixel on the object’s surface, resulting in pixel-wise (or dense) correspondence. The object pose is then estimated based on the pixel-wise correspondence.

symmetric objects, EPOS [19] discretizes the object surface into fragments and predicts a probability distribution over fragments. GDR-Net [41] utilizes a simple yet effective 2D convolutional Patch-PnP to regress the 6D pose directly. It also introduces a symmetry-aware module from EPOS [19] to handle symmetric objects. ZebraPose [37] proposes a hierarchical surface encoding technique and then efficiently solves the pose through the one-to-one correspondence.

Despite their excellent performance, approaches based on RGB images are still limited by texture-less objects and severe occlusions, which can result in wrong 2D-3D correspondences. With the help of depth-sensing information, approaches based on RGB-D images can generally perform better. Most RGB-D based methods assume pre-segmentation of object instances in the input image. After instance segmentation, the pose of each object is estimated individually. Existing approaches can be divided into two main categories: Direct Pose Prediction and Keypoint-based Prediction. The former regresses the poses of the segmented objects from the feature embedding layer directly in the image. It is computationally more efficient, but the pose obtained relies directly on the feature embedding, which would be less accurate and lacks explainability. The latter learns to select a set of keypoints (or proposals) for each object in the database. Given a segmented object in the input image, the corresponding points are found and matched at first. The object pose is then estimated based on the keypoint correspondence between the input image and the 3D models using closed-form solutions with outlier removal (e.g., RANSAC) or a network for pose estimation. The approach is more explainable since the keypoints acting as proposals are learnable and we know the correspondence between the input images and the database objects. However, the sparsity of keypoints may restrict the performance.

They are reviewed respectively in the following.

2.1. Direct Pose Prediction

Researches in [23, 29, 38, 40] opt for a direct prediction of rotation and translation. DenseFusion [40] introduces the densefusion module, which fuses RGB and depth fea-

tures at the pixel level and predicts the object pose accordingly. The final predicted object pose is determined by selecting the pixel with the highest confidence level. ES6D [29] presents XYZNet as a solution to mitigate the need for random memory accesses in the densefusion module to improve time efficiency. Additionally, they introduce the A(M)GPD loss, specifically designed to address the challenges posed by symmetric objects, as an improvement over the previous ADD-S loss. Experimental results demonstrate that the A(M)GPD loss yields greater efficacy in handling symmetric objects. In Unit6D [23] and Unit6Dv2 [38], the authors address the issue of projection breakdown by introducing supplementary UV data as an input, enabling a unified backbone to estimate the object pose accurately.

While direct pose prediction (or regression) methods are time-efficient, they usually have degraded performance compared to keypoint-based methods due to sensor noise.

2.2. Keypoint-based Prediction

Unlike direct pose prediction methods, keypoint-based approaches [16, 17, 42, 46] improve robustness using the projection equation. These methods define predetermined keypoints on the object’s surface and predict their positions in the image frame or camera coordinate system. The object’s pose is then computed based on these correspondences through PnP-RANSAC or least-square fitting [1, 13].

PVN3D [17] extends PVNet [31] to predict the keypoints in the 3D space because errors that may appear small in the projection can significantly impact the real world. FFB6D [16] enhances PVN3D [17] by incorporating bidirectional fusion modules to share information between the two modalities at an early stage. RCVPose [42] proposes a novel keypoint voting scheme that uses 1D radial voting to mitigate cumulative errors in each channel, which can significantly improve the accuracy of keypoint localization predictions. DFTr [46] fuses two cross-modal features using a transformer block to communicate global information between the two modalities. Moreover, instead of using the MeanShift [2] algorithm to perform keypoint voting, DFTr [46] proposes non-iterative weighted vector-wise voting scheme to reduce computational costs.

3. Method

Given an RGB-D image \mathcal{I} and the set of 3D CAD models $\mathcal{M} = \{\mathcal{M}_i | i = 1, \dots, N\}$, the primary objective of 6Dof object pose estimation is to accurately determine the rotation $\mathbf{R} \in SO(3)$ and the translation $\mathbf{t} \in \mathbb{R}^3$ that correspond to the camera coordinate system for each detected object $\mathcal{O} = \{\mathcal{O}_i | i = 1, \dots, O\}$ in the image \mathcal{I} .

3.1. Overview

Previous methods combining RGB and depth features at the pixel level require either an external instance segmen-




Figure 2. **Our proposed residual representation.** We use distributedly located anchor points and a fine-level residual vector (to the nearest anchor) to map each point on the object’s surface. This eliminates the need for the network to directly predict the exact coordinates where the range is extensively large, and makes the correspondence prediction more robust.

tation network [29, 40] or an integrated module designed for object segmentation [16, 17, 23, 38]. To address the demands of processing efficiency and the challenges posed by small objects, we adopt a two-step approach similar to the approaches proposed in [25, 41] that rely on object detection. However, unlike those that use RGB images, our approach uses RGB-D images. Processing the depth information (represented as the camera xyz map) is crucial to the success of object pose estimation.

Our proposed network, RDPN, is illustrated in Fig. 3. First, we use a readily available object detector to derive detection results. Afterward, for each detection, we crop the RGB-D image using the detected bounding box to obtain resized RGB-D images (\mathcal{I}_{rgb} and \mathcal{I}_{depth}). We first use the camera intrinsic matrix to derive the camera xyz map to utilize the depth information better. To obtain an accurate projected camera xyz map, we adjust \mathbf{K}_{org} , the original camera intrinsic matrix, to the adjusted intrinsic matrix \mathbf{K}_{crop} for the cropped window, and apply it to obtain the camera xyz map ($\mathcal{I}_{C_{xyz}}$). Then, we feed \mathcal{I}_{rgb} and $\mathcal{I}_{C_{xyz}}$ to our network. The network predicts both coarse anchors (\mathcal{A}_{coarse}) and residual vector ($\mathcal{F}_{residual}$) for each visible pixel of the object, as well as a mask (\mathcal{F}_{mask}) for removing background pixels. Leveraging these predictions, in conjunction with the 2D UV map (\mathcal{I}_{UV}) and down-sampled camera xyz map ($\mathcal{I}_{C_{xyz64}}$), our model predicts the object pose jointly using the 2D-3D and 3D-3D dense correspondences established.

3.2. Handling Cropping in RGB-D Images

To obtain the camera xyz map $\mathcal{I}_{C_{xyz}}$, we can utilize Eq. (1), which describes the projection of a point on the object’s surface (X, Y, Z) into the point of the image plane (u, v) through the camera intrinsic matrix \mathbf{K} after being rotated with rotation matrix \mathbf{R} and translated by \mathbf{t} :

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \frac{\mathbf{K}}{d} \cdot \left(\mathbf{R} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + \mathbf{t} \right), \quad (1)$$

where d is a scale factor.

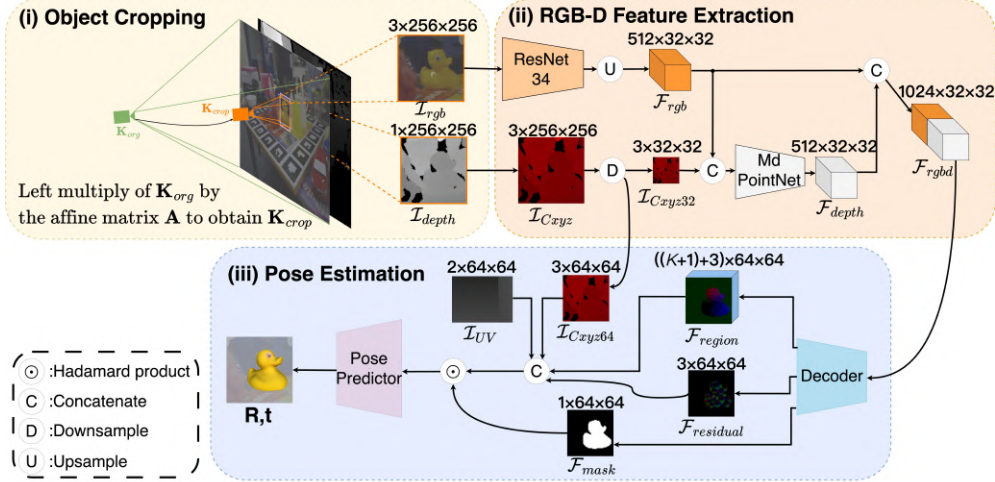


Figure 3. **Framework of RDPN.** (i) Starting with an RGB-D image, our initial step involves utilizing the outcomes of object detection to crop the region of interest (ROI), which results in a zoomed-in view ($\mathcal{I}_{rgb}, \mathcal{I}_{depth}$). In order to obtain the accurate projected camera xyz map \mathcal{I}_{Cxyz} , it is necessary to adjust the original camera intrinsic \mathbf{K}_{org} to \mathbf{K}_{crop} . (ii) Once we have prepared the \mathcal{I}_{rgb} and \mathcal{I}_{Cxyz} , the RGB-D feature extractor is responsible for capturing the RGB-D fusion features \mathcal{F}_{rgbd} , and feed them into a feature decoder to obtain both the mask (\mathcal{F}_{mask}) and per-pixel prediction to the point coordinates in the 3D model of the object. This includes a $(K + 1)$ -dimensional region probability (\mathcal{F}_{region}), 3-dimensional corresponding nearest anchors, and the residual vector ($\mathcal{F}_{residual}$). (iii) Finally, based on the mask and object coordinates, we utilize an image uv map (\mathcal{I}_{UV}) and a downsampled camera xyz map (\mathcal{I}_{Cxyz64}) to establish dense correspondences. These correspondences are then input into the pose predictor to regress the object pose \mathbf{R} and \mathbf{t} .

However, once we have acquired the object bounding boxes, we must crop and resize the images to obtain the Regions of Interest (RoI). This can be seen as a sub-image obtained by applying an affine transformation to the original image or modifying the intrinsic matrix of the original image. When transforming the depth image into a camera xyz map using the original intrinsic matrix, the projected location will be incorrect within the camera. Therefore, it is necessary to adjust the original camera intrinsic matrix \mathbf{K}_{org} by left multiplying it by an affine matrix \mathbf{A} (as illustrated in the upper left part of Fig. 3), which is a 3×3 matrix with the last row $[0 \ 0 \ 1]$. This gives us $\mathbf{K}_{crop} = \mathbf{A}\mathbf{K}_{org}$, which ensures accurate projection for the RoI.

In addition to achieving accurate projections, adjusting the intrinsic matrix addresses the diverse distribution of object positions between the training and testing datasets. Specifically, if we directly use \mathbf{K}_{org} , the object 3D points could appear in a wide range of locations in the image, such as the upper-left or lower-right. However, after adjusting the intrinsic matrix to \mathbf{K}_{crop} , the object will always be located in the center of the image, making it easier for the subsequent networks to robustly learn from uniformly distributed coordinates for correspondence finding.

3.3. Residual Representation

A possible approach to dense correspondence is to predict the 3D object coordinates of all object pixels in a single step. It is efficient in this way but requires the model to iden-

tify each mesh point within the image and store its coordinates in an extensively large range of the object coordinate system. Learning to predict them is challenging because of the large output space, which can lead to suboptimal performance for objects with complex shapes or symmetric features [17]. Additionally, symmetric features could cause ambiguities in the mapping, as there may be multiple points on the object similar to the given pixel in an image.

To address this concern, we propose a novel residual-based strategy that decouples the prediction of the object coordinates into the coarse and fine parts. For the coarse part, we first establish multiple anchors for the object \mathcal{M} using Farthest Point Sampling (FPS) [12]. FPS effectively subdivides the object’s surface into distinct subregions. Specifically, we can establish the anchor set \mathcal{A}_i for the given object \mathcal{M}_i as follows:

$$\mathcal{A}_i = \{\mathcal{A}_i^k | \mathcal{A}_i^k \in FPS(\mathcal{M}_i, K), \forall k = 1, \dots, K\}, \quad (2)$$

where K is the number of anchors. Each object in the 3D model database can be divided into separate regions based on its anchor points, determined by nearest-neighbor grouping. For each point on the object’s surface, we then use these anchors as reference points and compute the residual vector (i.e., fine part) to its nearest anchor, as illustrated in Fig. 2. By utilizing a residual representation, we can reformulate the projection equation as Eq. (3) below:

$$\begin{aligned}
\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} &= \frac{\mathbf{K}}{d} \cdot \left(\mathbf{R} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} + \mathbf{t} \right) \\
&= \frac{\mathbf{K}}{d} \cdot \left(\underbrace{\mathbf{R} \begin{bmatrix} R_x \\ R_y \\ R_z \end{bmatrix}}_{\text{coarse}} + \underbrace{\mathbf{R} \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix}}_{\text{fine}} + \mathbf{t} \right)
\end{aligned} \tag{3}$$

where $(R_x, R_y, R_z) \in \mathcal{A}_i$ represents the closest anchor point of the given point (X, Y, Z) on the object’s surface and $(r_x, r_y, r_z) = (X - R_x, Y - R_y, Z - R_z)$ denotes the residual of the point. This residual projection enables a condensed and uniform range to represent the correspondence relationship between the camera frame and the object coordinate system.

3.4. RDPN

Our proposed RDPN takes as input a cropped RGB image \mathcal{I}_{rgb} , a camera xyz map $\mathcal{I}_{C_{xyz}}$, and the predicted goal coarse part and fine part for the visible object pixels. The network is designed using an encoder-decoder architecture to predict the object coordinates for each pixel. The predicted object coordinates, the 2D UV map (\mathcal{I}_{UV}), and the downsampled camera xyz map ($\mathcal{I}_{C_{xyz64}}$) are then used to compute the object pose via a simple but effective network. **RGB-D Feature Encoder.** Our approach to fusing the information from the two modalities draws inspiration from ES6D [29]. Nevertheless, a critical distinction between their method and ours is that we do not concatenate \mathcal{I}_{rgb} and $\mathcal{I}_{C_{xyz}}$ initially and then feed them into CNNs to extract local features. We avoid this approach because background noise may hinder our detection-based approach’s effectiveness for fusion of the two modalities.

Instead, we extract local texture features (\mathcal{F}_{rgb}) from the RGB image \mathcal{I}_{rgb} using CNNs. We then combine these features with the downsampled camera xyz map ($\mathcal{I}_{C_{xyz32}}$) and input them into a PointNet-like CNNs architecture, using 1×1 convolutions to compress both position and texture information for individual pixels to get the spatial features (\mathcal{F}_{depth}). We empirically concatenate the two features at a 32×32 resolution to leverage the depth information efficiently. Finally, we obtain the global RGB-D fusion feature (\mathcal{F}_{rgb-d}) by concatenating the \mathcal{F}_{rgb} and \mathcal{F}_{depth} , as shown in the upper right part of Fig. 3.

RGB-D Feature Decoder. The residual representation of the 3D model object coordinates is predicted by the RGB-D feature decoder, as illustrated in the lower part of Fig. 3. To utilize the dense correspondences accurately, we also let the decoder network predict the visible object mask \mathcal{F}_{mask} and guide it with ground truth $\hat{\mathcal{F}}_{mask}$ by applying \mathcal{L}_1 loss:

$$\mathcal{L}_{mask} = \|\mathcal{F}_{mask} - \hat{\mathcal{F}}_{mask}\|_1. \tag{4}$$

For the coarse part in the residual representation, we approach it as a **classification problem**, opting to choose the anchor with the highest probability for each pixel. Consider the anchor set of the object i , $\mathcal{A}_i = \{\mathcal{A}_i^k | k = 1, \dots, K\}$. The object points can be divided into regions by these K anchors according to nearest-neighbor grouping. This defines the ground-truth labels of the point in a region: if the region is associated with anchor k , every point inside this region is labeled by a one-hot vector with the k -th element 1 and the others 0. The one-hot vector is of length $K + 1$, where the $(K + 1)$ -th element denotes the background. Hence, every spatial site of the RGB-D fusion feature \mathcal{F}_{rgb-d} can be classified as one of the classes in $\{1, \dots, (K + 1)\}$ after decoding. We use $\hat{\mathcal{F}}_{region}$ to represent the classification labels thus obtained. Hence, we can supervise the predicted \mathcal{F}_{region} with cross-entropy loss, formulate as:

$$\mathcal{L}_{coarse} = CrossEntropy(\mathcal{F}_{mask} \odot \mathcal{F}_{region}, \hat{\mathcal{F}}_{region}), \tag{5}$$

where \odot denotes the Hadamard product, \mathcal{F}_{region} records the probability of being the k -th class ($k = 1 \dots (K + 1)$) for every spatial site. To infer coarse correspondence locations, we need anchor positions in addition to class (or anchor label) probabilities. We employ the anchor coordinate $\mathcal{A}_i^k = (R_x, R_y, R_z)$ as feature representation with k the most probable class. This yields an additional 3 channels. An illustration can be found in the lower part of Fig. 3.

For the fine part in the residual representation, we approach it as a **regression problem**. To generate the residual \mathcal{F}_{fine} , we employ a straightforward approach by applying an \mathcal{L}_1 loss to the ground truth residual $\hat{\mathcal{F}}_{fine}$:

$$\mathcal{L}_{fine} = \|\mathcal{F}_{mask} \odot (\mathcal{F}_{fine} - \hat{\mathcal{F}}_{fine})\|_1. \tag{6}$$

Pose Predictor. To effectively predict an object’s pose using dense correspondences, as opposed to keypoint-based methods [16, 17, 46] relying on least-squares fitting, we adopt the methodology presented in [41]. The pose predictor (PP) is a simple 2D CNN based on the dense visible correspondences.

The overall loss function can be summarized as

$$\mathcal{L}_{RDPN} = \mathcal{L}_{coarse} + \mathcal{L}_{fine} + \mathcal{L}_{mask} + \mathcal{L}_{Pose}, \tag{7}$$

where \mathcal{L}_{Pose} is a disentangled 6D pose loss detailed in [41], employed to monitor the estimation of the pose \mathbf{R} and \mathbf{t} .

4. Experiments

In this section, we present the experimental results of our RDPN method and comparisons to the other approaches.

4.1. Implementation Details

We implemented our approach with PyTorch framework on a single RTX 3090 GPU. We employ YOLOX-x [14]

Method	RGB Approaches			RGB-D Approaches							
	PVNet [31] (CVPR' 19)	CDPN [25] (ICCV' 19)	DPODv2 [35] (TPAMI' 21)	PointFusion [44] (CVPR' 18)	DenseFusion [40] (CVPR' 19)	G2L-Net [8] (CVPR' 20)	PVN3D [17] (CVPR' 20)	FFB6D [16] (CVPR' 21)	RCVPose [42] (ECCV' 22)	DFTr [46] (ICCV' 23)	RDPN (Ours)
ape	43.6	64.4	100.0	70.4	92.3	96.8	97.3	98.4	99.2	98.6	99.7
benchvise	99.9	97.8	100.0	80.7	93.2	96.1	99.7	100.0	99.6	100.0	100.0
camera	86.9	91.7	100.0	60.8	94.4	98.2	99.6	99.9	99.7	100.0	99.9
can	95.5	95.9	100.0	61.1	93.1	98.0	99.5	99.8	99.0	100.0	100.0
cat	79.3	83.8	100.0	79.1	96.5	99.2	99.8	99.9	99.4	100.0	100.0
driller	96.4	96.2	100.0	47.3	87.0	99.8	99.3	100.0	99.7	100.0	100.0
duck	52.6	66.8	100.0	63.0	92.3	97.7	98.2	98.4	99.4	99.1	100.0
eggbox	99.2	99.7	99.04	99.9	99.8	100.0	99.8	100.0	98.7	100.0	100.0
glue	95.7	99.6	98.03	99.3	100.0	100.0	100.0	100.0	99.7	100.0	100.0
holepuncher	82.0	85.8	99.03	71.8	92.1	99.0	99.9	99.8	99.8	100.0	100.0
iron	98.9	97.9	100.0	83.2	97.0	99.3	99.7	99.9	99.9	99.9	100.0
lamp	99.3	97.9	100.0	62.3	95.3	99.5	99.8	99.9	99.2	100.0	100.0
phone	92.4	90.8	100.0	78.8	92.8	98.9	99.5	99.7	99.1	99.6	100.0
Avg (13)	86.3	89.9	99.70	73.7	94.3	98.7	99.4	99.7	99.43	99.8	99.97

Table 1. Quantitative evaluation of 6D Pose ADD(-S) 0.1d on the LineMOD dataset. Symmetric objects are in bold.

Method	PoseCNN [43] (RSS' 18)	Hybridpose [36] (CVPR' 18)	PVN3D [17] (CVPR' 20)	FFB6D [16] (CVPR' 21)	RCVPose [42] (ECCV' 22)	Uni6D [23] (CVPR' 22)	Uni6Dv2 [38] (AISTATS' 23)	DFTr [46] (ICCV' 23)	RDPN (Ours)
ape	9.6	20.9	33.9	47.2	60.3	33.0	44.3	<u>64.1</u>	64.6
can	45.2	75.3	88.6	85.2	92.5	51.0	53.3	<u>96.1</u>	97.0
cat	0.9	24.9	39.1	45.7	50.2	4.6	16.7	<u>52.2</u>	54.8
driller	41.4	70.2	78.4	81.4	78.2	58.4	63.0	<u>95.8</u>	<u>93.1</u>
duck	19.6	27.9	41.9	53.9	52.1	34.8	38.1	<u>72.3</u>	<u>68.8</u>
eggbox	22.0	52.4	<u>80.9</u>	70.2	81.2	1.7	4.6	75.3	78.1
glue	38.5	53.8	68.1	60.1	72.1	30.2	40.3	<u>79.3</u>	83.5
holepuncher	22.1	54.2	74.7	85.9	75.2	32.1	50.9	<u>86.8</u>	96.1
Avg (8)	24.9	47.5	63.2	66.2	70.2	30.7	40.2	<u>77.7</u>	79.5

Table 2. Quantitative evaluation of 6D Pose ADD(-S) 0.1d on the Occlusion-LineMOD dataset. Symmetric objects are in bold.

for object detection. The number of anchors is selected as $K = 32$. To address the issue of varying object sizes, we adopt Dynamic Zoom-In (DZI) to scale ground truth bounding boxes, as suggested in [25, 41]. More implementation details can be found in the supplementary material.

4.2. Benchmark Datasets

LineMOD [18] contains sequences of 13 objects with mild occlusion, cluttered scenes, texture-less objects, and variations in lighting conditions, presenting challenges for accurate pose estimation in real-world scenarios. Our approach follows the protocol established in [11, 31, 41–43], which employs the standard 15%/85% training/testing split. During training, we leverage further 1k rendered images for each object, as described in the settings in [11, 25, 41].

Occlusion LineMOD [5] is an extension of LineMOD that contains challenging test images of partially occluded objects. It involves training on LineMOD images and testing Occlusion LineMOD to assess the robustness of handling heavily occluded objects.

YCB-Video [43] is an extensive dataset including 130K key frames from 92 videos featuring 21 objects with varying lighting conditions and occlusions. Following [16, 17, 43], we split the dataset into 80 training videos and select 2,949 keyframes from the remaining 12 videos for testing.

MP6D [6] is a challenging dataset for the 6D pose estimation of metal parts in industrial environments. It contains 77 RGB-D video segments with simultaneous multi-target, occluded, and illumination changes. All objects are *symmet-*

ric, textureless, and of complex shape, high reflectivity, and uniform color, which makes them difficult to distinguish. We follow [6, 46] to split the training and testing sets.

For Occlusion LineMOD and YCB-V, we also follow [11, 41] that utilizes synthetic data with physically-based rendering [21] for training.

4.3. Metrics

We follow previous works [16, 20, 23, 43] and use the common average distance metrics ADD-S and ADD(-S). First, let us consider the ADD metric that evaluates the average of pairwise points transformed by the predicted pose $[\mathbf{R}, \mathbf{t}]$ and the ground truth pose $[\mathbf{R}^*, \mathbf{t}^*]$

$$\text{ADD} = \frac{1}{m} \sum_{x \in \mathcal{O}} \|(\mathbf{R}x + \mathbf{t}) - (\mathbf{R}^*x + \mathbf{t}^*)\|, \quad (8)$$

where x denotes the vertex sampled from the given 3D model \mathcal{O} , and m denotes the total number of vertices.

The ADD-S metric measures the closest point distance between two point clouds, making it more suitable for evaluating symmetric objects.

$$\text{ADD-S} = \frac{1}{m} \sum_{x_1 \in \mathcal{O}} \min_{x_2 \in \mathcal{O}} \|(\mathbf{R}x_1 + \mathbf{t}) - (\mathbf{R}^*x_2 + \mathbf{t}^*)\|. \quad (9)$$

ADD(-S) [18] calculates ADD for non-symmetric objects and ADD-S for symmetric objects, respectively.

For Linemod and Occlusion LineMOD datasets, we follow [23, 31] to report ADD(-S) 0.1d. This metric measures

	PVN3D [17]		FFB6D [16]		RCVPose [42]		ES6D [29]		Uni6D [23]		Uni6Dv2 [38]		DFTr [46]		RDPN (Ours)	
	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)	ADD-S	ADD(-S)
002_master_chef_can	96.0	80.5	96.3	80.6	95.7	93.6	96.5	73.0	95.4	70.2	96.0	74.2	97.0	92.3	99.8	78.7
003_cracker_box	96.1	94.8	96.3	94.6	97.2	95.7	95.3	94.0	91.8	85.2	96.0	94.2	95.9	93.9	98.3	88.3
004_sugar_box	97.4	96.3	97.6	96.6	97.6	97.2	97.9	97.3	96.4	94.5	97.6	96.6	97.1	95.5	100.0	99.9
005_tomato_soup_can	96.2	88.5	95.6	89.6	98.2	94.7	97.3	90.4	95.8	85.4	96.1	86.6	95.6	92.6	98.4	97.1
006_mustard_bottle	97.5	96.2	97.8	97.0	97.9	97.2	98.2	97.9	95.4	91.7	97.8	96.7	97.6	96.3	99.7	97.0
007_tuna_fish_can	96.0	89.3	96.8	88.9	98.2	96.4	97.4	93.7	95.2	79.0	96.3	76.0	97.3	94.5	100.0	99.2
008_pudding_box	97.1	95.7	97.1	94.6	97.7	97.1	96.5	93.4	94.1	89.8	96.6	94.7	97.4	95.7	95.0	89.3
009_gelatin_box	97.7	96.1	98.1	96.9	97.7	96.5	97.7	96.5	97.4	96.2	98.0	97.0	97.6	96.3	100.0	100.0
010_potted_meat_can	93.3	88.6	94.7	88.1	97.9	90.2	92.5	84.6	93.0	89.6	95.7	91.9	95.9	92.1	97.9	90.1
011_banana	96.6	93.7	97.2	94.9	97.9	96.7	97.9	85.8	96.4	93.0	98.0	96.9	97.1	95.0	100.0	96.6
019_pitcher_base	97.4	96.5	97.6	96.9	96.2	95.7	97.8	97.7	96.2	94.2	97.5	96.9	96.0	93.1	100.0	98.7
021_bleach_cleanser	96.0	93.2	96.8	94.8	99.2	97.8	96.3	92.8	95.2	91.1	97.0	95.3	96.8	94.9	98.3	89.8
024_bowl	90.2	90.2	96.3	96.3	95.2	94.9	96.4	96.4	95.5	95.5	96.8	96.8	96.9	96.9	87.8	87.8
025_mug	97.6	95.4	97.3	94.2	98.4	96.3	97.3	95.0	96.6	93.0	97.7	96.3	97.6	94.9	100.0	97.9
035_power_drill	96.7	95.1	97.2	95.9	96.2	95.4	97.2	96.3	94.7	91.1	97.6	96.8	96.9	95.2	99.9	97.8
036_wood_block	90.4	90.4	92.6	92.6	89.1	89.3	94.4	94.4	94.3	94.3	96.1	96.1	96.2	96.2	96.7	96.7
037_scissors	96.7	92.7	97.7	95.7	96.2	94.7	87.1	61.5	87.6	79.6	95.0	90.3	97.2	93.3	96.6	89.3
040_large_marker	96.7	91.8	96.6	89.1	95.9	92.4	97.8	90.6	96.7	92.8	97.0	93.1	96.9	92.7	99.9	93.9
051_large_clamp	93.6	93.6	96.8	96.8	95.2	96.4	61.0	61.0	95.9	95.9	97.0	97.0	96.3	96.3	98.9	98.9
052_extra_large_clamp	88.4	88.4	96.0	96.0	94.7	94.7	59.6	59.6	95.8	95.8	96.5	96.5	96.4	96.4	98.6	98.6
061_foam_brick	96.8	96.8	97.3	97.3	95.7	95.7	96.6	96.6	96.1	96.1	97.4	97.4	97.3	97.3	100.0	100.0
Avg (21)	95.5	91.8	96.6	92.7	96.6	95.2	93.6	89.0	95.2	88.8	96.8	91.5	96.7	94.4	98.4	94.6

Table 3. Quantitative results of 6D Pose (ADD-S AUC and ADD(-S) AUC) on the YCB-Video dataset. Symmetric objects are in bold.

	Hodan [22]	PointFusion [44]	DCF [26]	DF (per-pixel) [40]	MaskedFusion [32]	G2L-Net [8]	PVN3D [17]	FFB6D [16]	DFTr [46]	RDPN (Ours)
Avg (20)	74.20	75.54	75.93	79.84	79.38	83.51	85.42	86.29	93.01	95.90

Table 4. Quantitative evaluation of 6D Pose (ADD-S AUC) on the MP6D dataset. DF (per-pixel) denotes DenseFusion (per-pixel). Note that all objects are symmetric.

the accuracy of distances less than 10% of the objects’ diameter. For the YCB-Video dataset, we follow [16, 46] to report the area under the accuracy-threshold curve computed by varying the distance threshold ADD-S AUC and ADD(-S) AUC. For MP6D, we report ADD-S AUC, following [6, 16, 23, 43] since all objects are symmetric.

4.4. Comparison with State-of-the-Art Methods

Results on LineMOD & Occlusion LineMOD. The quantitative results of our proposed RDPN are presented in Tab. 1 and Tab. 2 for the LineMOD and Occlusion LineMOD datasets, respectively. Our method achieves state-of-the-art performance on both datasets. Particularly, in occlusion scenes, our approach outperforms most previous methods extensively, demonstrating its effectiveness.

Results on YCB-V. Results of our proposed RDPN on the YCB-Video dataset are shown in Tab. 3. Our approach outperforms the previous state-of-the-art approaches, achieving a superior ADD-S AUC metric by 1.7%. It also demonstrates competitive performance in the ADD(-S) AUC without time-consuming post-processing.

Results on MP6D. The results of our RDPN on the MP6D dataset are shown in Tab. 4. Again, our method outperforms the previous state-of-the-art methods in the ADD-S AUC metric, demonstrating its generalizability across different datasets. Compared to other RGB-D based solutions, RDPN is more robust towards objects with texture-less or heavy reflective surfaces and severe occlusions.

Qualitative Results. Furthermore, to illustrate the robustness of our RDPN in handling occlusion, we present some

estimation results on the Occlusion LineMOD dataset. As shown in Fig. 4, RDPN demonstrates superior performance and robustness across various occlusion scenarios.

4.5. Ablation Study on Occlusion LineMOD

Effectiveness of the residual representation is shown in Tab. 5. We compare the approaches using the original and our anchor+residual representations for predicting the object coordinates. As the table shows, it is clear that residual representations outperform the original representation, highlighting the advantage of reducing the output space.

Object Coordinate	Average (8)		
	ADD(-S) 0.02d	ADD(-S) 0.05d	ADD(-S) 0.1d
Original	7.7	50.2	78.5
Residual Representation	9.1	51.3	79.5

Table 5. Ablation on the effectiveness of residual representation.

2D-3D Correspondences	3D-3D Correspondences	Average (8)		
		ADD(-S) 0.02d	ADD(-S) 0.05d	ADD(-S) 0.1d
v	-	4.6	35.3	67.3
-	v	9.4	50.9	78.4
v	v	9.1	51.3	79.5

Table 6. Ablation study on the effectiveness of different component of dense correspondences.

Effectiveness of the Dense Correspondence components is compared and presented in Tab. 6. We conducted experiments under three different configurations.

Initially, when the pose predictor relies solely on dense correspondence with the 2D image UV map \mathcal{I}_{UV} to predict the object coordinates (i.e., 2D-3D correspondences),

Intrinsic Parameter	Average (8)		
	ADD(-S) 0.02d	ADD(-S) 0.05d	ADD(-S) 0.1d
\mathbf{K}_{org}	7.1	41.2	72.4
\mathbf{K}_{crop}	9.1	51.3	79.5

Table 7. Ablation study on the effectiveness of adjusting the intrinsic parameter.

the performance is inferior to the previously employed direct regression methods. However, replacing \mathcal{I}_{UV} with the camera xyz map $\mathcal{I}_{C_{xyz64}}$ (i.e., 3D-3D correspondences) significantly improved all ADD(-S) metrics. This observation is intuitive, as the pose predictor can effectively predict the object’s pose without considering the camera intrinsic matrix \mathbf{K} through Eq. (1). Finally, utilizing dense correspondence with both \mathcal{I}_{UV} and $\mathcal{I}_{C_{xyz}}$, along with the predicted object coordinates (2D/3D to 3D correspondences), further enhances the performance. We attribute that the sensor noise in collected depth images can be rectified by incorporating \mathcal{I}_{UV} , when both \mathcal{I}_{UV} and $\mathcal{I}_{C_{xyz64}}$ are available.

Effectiveness of adjusting intrinsic \mathbf{K}_{org} is shown in Tab. 7, which reveals better performance when using \mathbf{K}_{crop} to achieve accurate projection. This suggests that the intrinsic \mathbf{K} adjustment step is necessary in our method.

Effectiveness of different number of anchors is shown in Fig. 5. We observe a noticeable improvement when increasing the numbers of anchors from 4 to 16, with the best performance achieved at 32 anchors. However, the performance starts to decline after using 32 anchors. We believe that as the number of anchors increases, the output space of the fine part is reduced, while the output space of the coarse part inevitably increases. Therefore, there is a trade-off between the number of anchors and overall performance.

5. Limitation

Our method demonstrates remarkable speed, achieving pose estimation for 8 objects in a 640x480 image within a mere 29ms/frame. However, the overall processing time is currently constrained by the object detection stage, which requires 66ms/frame using the YOLOX-x [14] detector. To further enhance the efficiency of our method, we plan to explore the development of a unified network for object detection and pose estimation in future work.

6. Conclusion

We introduce a novel Residual-based Dense Point-wise Network (RDPN) for precise and efficient object pose estimation from RGB-D images. By employing an intrinsic adjustment and a residual representation for object coordinates, RDPN effectively condenses and concentrates the output range, eliminating the need to predict coordinates within a vast and relatively unbounded space. Leveraging both 2D-3D and 3D-3D dense correspondences, our method achieves state-of-the-art performance on public benchmarks



Figure 4. **Qualitative results on Occlusion LineMOD.** The images are rendered by projecting the 3D object model onto the image plane using the estimated pose. Our two-step dense correspondence method can accurately capture the object and predict its pose, even under heavy occlusion. This contrasts previous keypoint-based methods, which often struggle in such scenarios.




Figure 5. Ablation study on the number of anchors. The 10°, 10 cm metric measures whether the rotation and the translation error is less than 10° and 10 cm, respectively.

with high efficacy.

Acknowledgement This work was supported in part by the National Science and Technology Council, Taiwan under Grant NSTC 112-2221-E-002-182-MY3 and 112-2634-F-002-005, and also under Grant UR2205 of Delta-NTU joint R&D center. We thank to National Center for High-performance Computing (NCHC) of National Applied Research Laboratories (NARLabs) in Taiwan for providing computational and storage resources.

References

- [1] Least-squares fitting of two 3-d point sets. *IEEE Transactions on pattern analysis and machine intelligence*, (5):698–700, 1987. [3](#)
- [2] Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002. [3](#)
- [3] Arash Amini, Arul Selvam Periyasamy, and Sven Behnke. T6d-direct: Transformers for multi-object 6d pose direct regression. In *DAGM German Conference on Pattern Recognition*, pages 530–544. Springer, 2021. [1](#), [2](#)
- [4] Arash Amini, Arul Selvam Periyasamy, and Sven Behnke. Yolopose: Transformer-based multi-object 6d pose estimation using keypoint regression. In *International Conference on Intelligent Autonomous Systems*, pages 392–406. Springer, 2022. [1](#), [2](#)
- [5] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part II 13*, pages 536–551. Springer, 2014. [2](#), [6](#)
- [6] Long Chen, Han Yang, Chenrui Wu, and Shiqing Wu. Mp6d: An rgb-d dataset for metal parts’ 6d pose estimation. *IEEE Robotics and Automation Letters*, 7(3):5912–5919, 2022. [2](#), [6](#), [7](#)
- [7] Shuai Chen, Xinghui Li, Zirui Wang, and Victor Prisacariu. Dfnet: Enhance absolute pose regression with direct feature matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [1](#)
- [8] Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, and Ales Leonardis. G2l-net: Global to local network for real-time 6d pose estimation with embedding vector features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4233–4242, 2020. [6](#), [7](#)
- [9] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. [1](#)
- [10] Alvaro Collet, Manuel Martinez, and Siddhartha S Srinivasa. The moped framework: Object recognition and pose estimation for manipulation. *The international journal of robotics research*, 30(10):1284–1306, 2011. [1](#)
- [11] 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. [1](#), [6](#)
- [12] Y. Eldar, M. Lindenbaum, M. Porat, and Y.Y. Zeevi. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9):1305–1315, 1997. [4](#)
- [13] 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. [3](#)
- [14] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. [5](#), [8](#)
- [15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. [1](#)
- [16] Yisheng He, Haibin Huang, Haoqiang Fan, Qifeng Chen, and Jian Sun. Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3003–3013, 2021. [1](#), [3](#), [5](#), [6](#), [7](#)
- [17] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian 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*, pages 11632–11641, 2020. [1](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [18] 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 *Computer Vision—ACCV 2012: 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5–9, 2012, Revised Selected Papers, Part I 11*, pages 548–562. Springer, 2013. [1](#), [2](#), [6](#)
- [19] Tomas Hodan, Daniel Barath, and Jiri Matas. Epos: Estimating 6d pose of objects with symmetries. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11703–11712, 2020. [1](#), [2](#)
- [20] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6d object pose estimation. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pages 606–619. Springer, 2016. [6](#)
- [21] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. Bop challenge 2020 on 6d object localization. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 577–594. Springer, 2020. [6](#)
- [22] Tomáš Hodaň, Xenophon Zabulis, Manolis Lourakis, Štěpán Obdržálek, and Jiří 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)*, pages 4421–4428. IEEE, 2015. [7](#)
- [23] Xiaoke Jiang, Donghai Li, Hao Chen, Ye Zheng, Rui Zhao, and Liwei Wu. Uni6d: A unified cnn framework without projection breakdown for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11174–11184, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [24] 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. [2](#)
- [25] 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/CVF International Conference on Computer Vision*, pages 7678–7687, 2019. [1](#), [2](#), [3](#), [6](#)
- [26] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *Proceedings of the European conference on computer vision (ECCV)*, pages 641–656, 2018. [7](#)
- [27] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1150–1157. Ieee, 1999. [1](#)
- [28] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE transactions on visualization and computer graphics*, 22(12):2633–2651, 2015. [1](#)
- [29] Ningkai Mo, Wanshui Gan, Naoto Yokoya, and Shifeng Chen. Es6d: A computation efficient and symmetry-aware 6d pose regression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6718–6727, 2022. [1](#), [2](#), [3](#), [5](#), [7](#)
- [30] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7668–7677, 2019. [2](#)
- [31] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pynet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019. [1](#), [3](#), [6](#)
- [32] Nuno Pereira and Luís A Alexandre. Maskedfusion: Mask-based 6d object pose estimation. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 71–78. IEEE, 2020. [7](#)
- [33] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#)
- [34] Yoli Shavit, Ron Ferens, and Yosi Keller. Coarse-to-fine multi-scene pose regression with transformers. *IEEE transactions on pattern analysis and machine intelligence*, PP, 08 2023. [1](#)
- [35] Ivan Shugurov, Sergey Zakharov, and Slobodan Ilic. Dpodv2: Dense correspondence-based 6 dof pose estimation. *IEEE transactions on pattern analysis and machine intelligence*, 44(11):7417–7435, 2021. [6](#)
- [36] 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. [6](#)
- [37] Yongzhi Su, Mahdi Saleh, Torben Fetzer, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6738–6748, 2022. [1](#), [2](#)
- [38] Mingshan Sun, Ye Zheng, Tianpeng Bao, Jianqiu Chen, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xiaoke Jiang. Uni6dv2: Noise elimination for 6d pose estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 1832–1844. PMLR, 2023. [1](#), [2](#), [3](#), [6](#), [7](#)
- [39] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790*, 2018. [1](#)
- [40] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3343–3352, 2019. [1](#), [2](#), [3](#), [6](#), [7](#)
- [41] 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. [1](#), [2](#), [3](#), [5](#), [6](#)
- [42] Yangzheng Wu, Mohsen Zand, Ali Etemad, and Michael Greenspan. Vote from the center: 6 dof pose estimation in rgb-d images by radial keypoint voting. In *European Conference on Computer Vision*, pages 335–352. Springer, 2022. [1](#), [3](#), [6](#), [7](#)
- [43] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. 2018. [2](#), [6](#), [7](#)
- [44] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 244–253, 2018. [6](#), [7](#)
- [45] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1941–1950, 2019. [2](#)
- [46] Jun Zhou, Kai Chen, Linlin Xu, Qi Dou, and Jing Qin. Deep fusion transformer network with weighted vector-wise keypoints voting for robust 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13967–13977, October 2023. [1](#), [3](#), [5](#), [6](#), [7](#)