

Robust 6D Object Pose Estimation by Learning RGB-D Features

Meng Tian¹, Liang Pan¹, Marcelo H Ang Jr¹ and Gim Hee Lee²

Abstract—Accurate 6D object pose estimation is fundamental to robotic manipulation and grasping. Previous methods follow a local optimization approach which minimizes the distance between closest point pairs to handle the rotation ambiguity of symmetric objects. In this work, we propose a novel discrete-continuous formulation for rotation regression to resolve this local-optimum problem. We uniformly sample rotation anchors in $SO(3)$, and predict a constrained deviation from each anchor to the target, as well as uncertainty scores for selecting the best prediction. Additionally, the object location is detected by aggregating point-wise vectors pointing to the 3D center. Experiments on two benchmarks: LINEMOD and YCB-Video, show that the proposed method outperforms state-of-the-art approaches. Our code is available at <https://github.com/mentian/object-posenet>.

I. INTRODUCTION

6D object pose estimation is an important foundation for various robotic tasks, such as robotic manipulation and grasping [1], [2]. This problem is very challenging due to the varying illumination conditions, background clutters, and heavy occlusions between objects. Objects that are symmetrical or textureless further aggravate this problem. Additionally, a viable solution should be able to infer at real-time speed.

The recent explosion in 6D object pose estimation is arguably a result of the application of deep neural networks. Many proposed deep networks [3]–[10] only leverage RGB data, which are inherently sensitive to changing lighting conditions [11] and object appearance variations [6]. To mitigate these problems, researchers start to take advantage of 3D geometric features and use RGB-D images for object pose estimation. [12] proposes a unified scalable framework with CNN architectures to exploit features from both RGB and depth images. [13] introduces a PointNet-based [14] dense-fusion network that extracts geometric features from a point cloud and fuses with color features. However, both of them fall short of fully resolving pose ambiguities caused by symmetric objects. The former does not differentiate between symmetric and asymmetric objects and ignores the ambiguities. The latter adopts a naive version of ShapeMatch-Loss [5] which suffers the local-optimum problem (Fig. 1) caused by minimizing the distance between closest point pairs.

A symmetric object can give rise to ambiguous rotation estimations, while it does not influence translation. We argue that rotation and translation should be solved separately

¹Department of Mechanical Engineering, National University of Singapore tianmeng@u.nus.edu, pan.liang@u.nus.edu, mpeangh@nus.edu.sg

²Computer Vision and Robotic Perception Lab, Department of Computer Science, School of Computing, National University of Singapore gimhee.lee@nus.edu.sg



Fig. 1. Some incorrect predictions from PoseCNN [5] (left), DenseFusion [13] (middle), and our implementation of ShapeMatch-Loss (right). These errors are typical and commonly seen in the results of all three methods. It is strongly indicated that ShapeMatch-Loss suffers the local-optimum problem.

as opposed to jointly optimizing them. To this end, we propose a novel end-to-end deep network which densely extracts features from RGB-D images and estimate rotation and translation in two separate branches. For rotation, we initialize the estimation with uniformly sampled rotations (referred to as rotation anchors) in $SO(3)$. Given a target rotation, our network predicts a deviation from each rotation anchor to the target. All the rotations anchor together cover the whole $SO(3)$ space and each anchor is only responsible for a local region. By local prediction and local optimization, our method resolves the local-optimum problem. We also predict an uncertainty score for each anchor, which is used to select the best prediction during inference. Inspired by [15], [16], the score is learnt in a self-supervised way by maximizing the conditional probability of the rotation error given the uncertainty score. For translation, our network regresses a unit vector pointing towards the 3D object center at each point. We integrate a RANSAC-based voting layer to aggregate all the vectors and select the best hypothesis as the final estimation. Robustness of our method is twofold: first, by densely extracting RGB-D features our method is robust to changing lighting conditions and object appearance variations; and secondly, our regression strategies for both rotation and translation are robust. To be specific, the proposed rotation regression consists of local region classification and constrained residual regression, which is more robust than a single regression strategy. The RANSAC-based voting enforces local prediction such that it is robust to occlusions and clutters.

We evaluate our method on two widely-used benchmarks for 6D pose estimation: LINEMOD dataset [17] and YCB-Video [5] dataset. Our method achieves superior performance on both datasets over the state-of-the-art methods. We also demonstrate its robustness to heavy occlusions, background clutters and varying lighting conditions.

In summary, our key contributions are as follows:

- We propose a novel deep network for 6D pose estimation, which robustly predicts rotation and translation from densely extracted RGB-D features.
- We propose a discrete-continuous formulation to resolve the local-optimum problem of ShapeMatch-Loss and handle rotation ambiguities.
- Our network outperforms the state-of-the-art results on both LINEMOD (ADD: 92.8% vs. 86.3%) and YCB-Video (ADD: 83.8% vs. 79.2%).

II. RELATED WORK

RGB methods. Traditionally, object pose estimation from single RGB image is tackled by matching local features [18], [19]. However, sufficient textures are required to compute the features. To handle textureless objects, [20] and [21] propose to directly regress 3D object coordinates. But the sampling of pose hypothesis and refinement are very time-consuming. Recent methods apply deep learning techniques to this task. [3] extends an 2D object detection network [22] to predict object’s identity, 2D bounding box and a discretized orientation. [4] and [23] first detect the keypoints of the object and then solve a Perspective-n-Point (PnP) problem for pose estimation. These methods predict from the global feature of the object. Hence, performance on occluded objects are not satisfactory. To address the occlusion problem, dense prediction methods [5], [7]–[10] are proposed. [9] and [10] estimate a pose at each visible local patch independently, and accumulate results over all patches. [5] and [7] predict a unit vector at each pixel pointing towards keypoints. [8] estimates a dense 2D-3D correspondence map between input image and object models. Among them, [7] and [8] achieve top performances. But they are still inferior to RGB-D based methods in terms of detection accuracy.

RGB-D methods. LINEMOD [17] is considered as a seminal work, after which all methods are benchmarked. They estimate object pose by matching templates consisting of color gradient and surface normal features. Their templates are holistic, thus cannot handle occluded objects. [20], [24], [25] regress a 3D object coordinate for each pixel on image, from which pose hypotheses are generated and further refined. These methods rely on hand-crafted features which limits their performances. Both [26] and [27] use locally-sampled RGB-D patches to cast votes for 6D pose. [26] proposes a Latent-Class Hough Forest which stores center position and rotation at its leaf nodes. [27] employs a convolutional auto-encoder to extract descriptors, and match to the nearest neighbor in a codebook. [12] and [13] are two recent works that predict 6D object pose from RGB-D images in an end-to-end way. [12] processes color image and depth image separately, then concatenates two feature maps along feature dimension. The global fusion does not make full use of geometric information. In contrast, we follow a dense-fusion way as in [13] and [28]. [13] is the most similar work to us. However, we propose a novel discrete-continuous formulation for rotation prediction to resolve the local-optimum issue. We also extend the 2D RANSAC-based

voting to 3D space for estimating translation, which is a much more robust method than single point prediction.

III. METHOD

The task of 6D object pose estimation is to detect each object instance and meanwhile estimate a rigid transformation from the object coordinate frame to the camera coordinate frame. Specifically, this rigid transformation is represented by a rotation $\mathbf{R} \in SO(3)$ and a translation $\mathbf{t} \in R^3$.

A. Overview

Fig. 2 overviews the proposed two-stage pipeline: instance segmentation and pose estimation. The first stage is to detect each object instance on the color image and obtain a foreground mask (e.g. instance segmentation). Considering that each object only appears at most once on the images we evaluate on, we adopt the well-developed SegNet [29], which is originally designed for semantic segmentation, as our pre-processing network. If there are multiple instances of same object category presented on an image, this part could be straightforwardly replaced by an instance segmentation network, such as [30] and [31]. With the obtained mask, we crop an image window which tightly encloses the object, and extract a point cloud from depth image as inputs to the second-stage pose estimation network. Our network concatenates densely extracted color and geometric features in a pixel-wise manner, from which rotation and translation are predicted in two separate branches. For rotation branch, we propose a robust discrete-continuous regression scheme to handle ambiguities caused by symmetrical objects. The translation branch regresses a unit vector pointing towards the object center in 3D space at each point, and aggregates these vectors through RANSAC-based voting.

B. Rotation

Symmetric objects will cause convergence problem during training, since there are multiple rotation labels corresponding to an identical appearance. To alleviate this problem, ShapeMatch-Loss [5] is proposed and defined as

$$L = \begin{cases} \frac{1}{M} \sum_{\mathbf{x}_1 \in \mathcal{M}} \min_{\mathbf{x}_2 \in \mathcal{M}} \|\mathbf{R}\mathbf{x}_1 - \tilde{\mathbf{R}}\mathbf{x}_2\|_2 & \text{if symmetric,} \\ \frac{1}{M} \sum_{\mathbf{x} \in \mathcal{M}} \|\mathbf{R}\mathbf{x} - \tilde{\mathbf{R}}\mathbf{x}\|_2 & \text{if asymmetric;} \end{cases} \quad (1)$$

where M is the number of points in object model \mathcal{M} , $\tilde{\mathbf{R}}$ and \mathbf{R} are the ground truth and prediction. Ambiguous labels of symmetric objects become equivalent under this loss function. However, just like ICP, ShapeMatch-Loss is only a local optimization for symmetric objects. It will slow down the training process and converge to sub-optimal solution in the worst-case scenario.

We resolve this local-optimum issue by initializing our predictions with rotation anchors covering the whole $SO(3)$ space. Specifically, instead of directly estimating $\tilde{\mathbf{R}}$, our network predicts a deviation $\Delta\mathbf{R}_i$ at rotation anchor $\hat{\mathbf{R}}_i$:

$$\mathbf{R}_i = \Delta\mathbf{R}_i \hat{\mathbf{R}}_i, \quad (2)$$

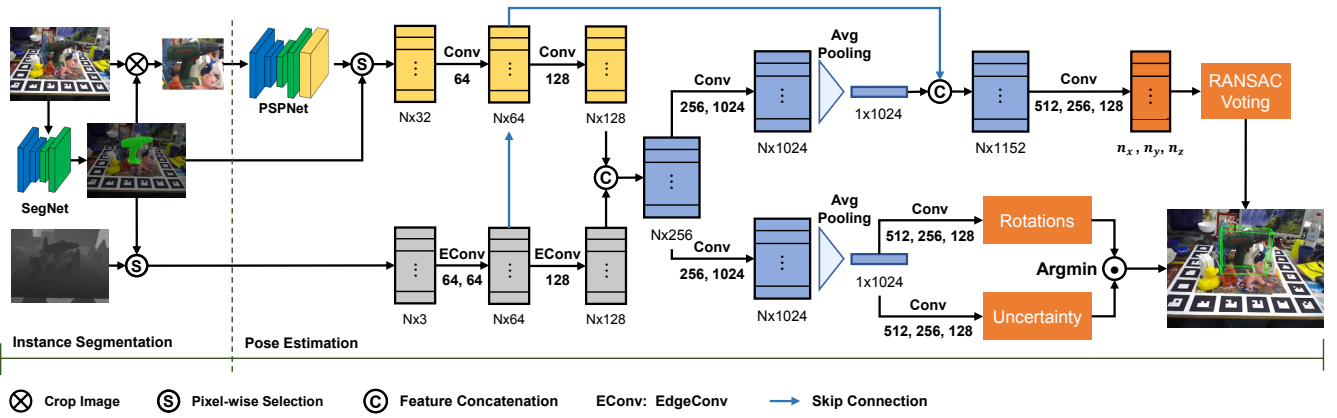


Fig. 2. **Overview of our pipeline.** Instance mask is first obtained through a semantic segmentation network (SegNet). With the obtained mask, our pose estimation network extracts a point cloud from depth image and crops a tight image window from color image as inputs. Color features and geometric features are densely concatenated, from which rotation and translation of the object are predicted in two separate branches.

where \mathbf{R}_i is the total prediction at i -th rotation anchor. $\Delta\mathbf{R}_i$ is constrained to be within the vicinity of $\hat{\mathbf{R}}_i$, such that each anchor responds for a local region. We investigate three discrete subgroups of $SO(3)$ as rotation anchors, namely tetrahedral group ($N=12$), octahedral group ($N=24$), and icosahedral group ($N=60$) [32], with N being the number of rotations in the corresponding group. Each of these three subgroups is a uniform sampling of $SO(3)$. In addition to $\Delta\mathbf{R}_i$, our network also predicts an uncertainty score $\sigma_i \in [0, 1]$ for each anchor. σ_i is proportional to the distance between \mathbf{R}_i and ground-truth rotation after the training is converged. During inference, prediction of the anchor with smallest uncertainty score is selected as the final output.

For each prediction \mathbf{R}_i , we define the conditional probability of the normalized ShapeMatch-Loss $d_i = L_i/d$ given σ_i as

$$p(d_i|\sigma_i) = \frac{1}{\sigma_i} \exp\left(-\frac{d_i}{\sigma_i}\right), \quad (3)$$

where d is the object diameter, L_i is the ShapeMatch-Loss of \mathbf{R}_i . For a certain d_i , maximum of p is achieved when $\sigma_i = d_i$. Maximizing p equals to minimizing d_i and setting σ_i to d_i . Therefore, optimal parameters of the network can be obtained by minimizing the probabilistic loss

$$\begin{aligned} L_R &= -\sum_i \ln p(d_i|\sigma_i) \\ &= \sum_i \ln \sigma_i + \frac{d_i}{\sigma_i}. \end{aligned} \quad (4)$$

The regularization loss which constrains the range of $\Delta\mathbf{R}_i$ is

$$L_{reg} = \sum_i \max(0, \max_{j \neq i} \langle \mathbf{q}_i, \hat{\mathbf{q}}_j \rangle - \langle \mathbf{q}_i, \hat{\mathbf{q}}_i \rangle), \quad (5)$$

where \mathbf{q}_i and $\hat{\mathbf{q}}_i$ are the quaternion representations of \mathbf{R}_i and $\hat{\mathbf{R}}_i$.

The proposed regression strategy has another benefit of improving the regressing accuracy. It is actually a practice of reformulating a continuous regression problem as a joint classification and residual regression problem. This kind of

reformulation has been widely proven to be effective [33], [34] and indeed verified in our experiments.

Remarks: [12] also uses a discrete-continuous formulation for regressing rotation. However, they do not enforce local prediction for each anchor. Their classification over the anchors is mutually exclusive, such that ambiguous labels of symmetric objects will confuse the classifier. Our formulation can be viewed as independently conducting binary classification for each anchor, which is more suitable for symmetric objects, Loss function of [13] could be derived from our probabilistic loss (Eq.4) by a reparameterization trick. The confidence value c in their loss function is the inverse of our uncertainty score. Hence, optimizing c leads to unbounded prediction. Their loss equally penalizes the distance error of objects of different sizes, thus has a bias towards large objects. While our loss normalizes the error with the corresponding object size. [35] proposes to resolve the rotation ambiguity by generating multiple hypotheses (e.g. 5). Compared with their solution, our formulation dynamically selects the number of hypotheses (up to the total number of anchors) according to the uncertainty scores.

C. Translation

In order to make use of the complementary depth information, we extend a RANSAC-based voting method [5], [7] from 2D to 3D space. Point-wise prediction focuses more on local features, which makes our prediction robust to foreground occlusions, background clutters and segmentation noise in the first stage. Suppose that the object center is $\mathbf{c} \in R^3$. At each selected point \mathbf{p} , our network predicts a unit vector \mathbf{v} representing the direction from that point to the center:

$$\mathbf{v} = \frac{\mathbf{c} - \mathbf{p}}{\|\mathbf{c} - \mathbf{p}\|_2}. \quad (6)$$

Randomly sampled two points and their associated vectors define two 3D lines. However, lines do not necessarily intersect in 3D space. To mitigate this problem, a hypothesis \mathbf{h} is generated by taking the mid-point of the shortest line

segment between these two lines. Votes are accumulated over all points. A point is considered as an inlier if $(\mathbf{h} - \mathbf{p})^T \mathbf{v} / \|\mathbf{h} - \mathbf{p}\|_2 \geq \theta$, where θ is a threshold and set to 0.99 for all experiments. Final estimate is set to the point closest to all the lines defined by the inliers of the hypothesis with highest voting score. During training, we employ smooth L1 loss (L_t) for learning the vectors.

D. Network Architecture

Our network has two separate branches for color feature and geometric feature extraction. Features from two different modalities are densely concatenated, from which translation and rotation are predicted independently. The translation branch enforces dense prediction. Therefore, we concatenate high-level global feature with low-level color and geometric features (shown as skip connections in Fig. 2). While the rotation branch only leverages the global average-pooled feature.

We use a fully convolutional network modified from PSPNet [36] to densely extract color features. With ResNet-18 [37] as backend, up-sampling layers are applied after pyramid pooling to restore the size of feature map. For arbitrary-sized image crop, network will output a feature map of same size. Then, those features belong to foreground pixels are selected, and lifted to the same dimension with geometric features by two 1x1 convolution layers. Local geometric structure has been proven useful for reasoning on point cloud [38]. We recurrently apply edge convolution operation (EdgeConv) proposed in [38] to extract geometric features. The obtained dense color and geometric features are concatenated based on pixel-to-point correspondences. Observing that color feature and geometric feature play different roles in estimating translation and rotation, it is better to separate two prediction branches and let the network learn to balance two-modal information, rather than constraining them to share one global feature. Our experiment validates this particular design of network architecture.

IV. EXPERIMENTS

A. Datasets

LINEMOD [17] is the standard and most popular benchmark for 6D object pose estimation. This dataset has 13 texture-less objects placed in heavily cluttered scenes, of which two are symmetric. Following [21], [23] and [13], training images are selected such that relative orientation between two images are larger than a threshold (e.g. 15°), and viewpoints of the upper hemisphere are regularly covered. About 15% of all images are selected for training. The rest are used for testing.

YCB-Video [5] consists of 92 real video sequences collected for 21 objects from YCB objects [39] and 80k synthetically rendered images. 2949 key frames are extracted from 12 video sequences for testing, while the remaining sequences and synthetic images can be used for training. This dataset is extremely challenging due to varying light conditions, presence of heavy occlusions, and significant image noise.

TABLE I

ABLATION STUDIES: ACCURACIES ON LINEMOD DATASET IN TERMS OF ADD(-S) METRIC ARE REPORTED.

Architecture	Accuracy
DenseFusion [13]	86.20
RANSAC voting	89.42
local geometry	90.36
60 anchors	92.77
shared global feature	91.89
w/o regularization	84.71
24 anchors	89.89
12 anchors	87.86

B. Evaluation Metrics

We report the most widely used pose error in 3D space, which is referred to as ADD metric [17]. ADD is defined as the average distance between model vertices transformed with estimated pose (\mathbf{R}, \mathbf{t}) and ground-truth pose $(\tilde{\mathbf{R}}, \tilde{\mathbf{t}})$:

$$\text{ADD} = \frac{1}{M} \sum_{\mathbf{x} \in \mathcal{M}} \|(\mathbf{R}\mathbf{x} + \mathbf{t}) - (\tilde{\mathbf{R}}\mathbf{x} + \tilde{\mathbf{t}})\|_2, \quad (7)$$

where $\mathbf{x} \in \mathcal{M}$ denotes the 3D model point and M is the total number of points. A predicted 6D pose is considered to be correct if ADD is less than a threshold (e.g. 10% of object diameter). For symmetric objects, the error is calculated as the average distance to the closest model point [40]:

$$\text{ADD-S} = \frac{1}{M} \sum_{\mathbf{x}_1 \in \mathcal{M}} \min_{\mathbf{x}_2 \in \mathcal{M}} \|(\mathbf{R}\mathbf{x}_1 + \mathbf{t}) - (\tilde{\mathbf{R}}\mathbf{x}_2 + \tilde{\mathbf{t}})\|_2. \quad (8)$$

ADD(-S) combines two metrics and means we use ADD metric for asymmetric objects and ADD-S metric for symmetric objects during evaluation. AUC (area under accuracy-threshold curve) is a commonly used metric on YCB-Video dataset [5], [9], [12], [13]. We also report AUC in terms of ADD-S metric within a range of [0, 0.1].

C. Implementation Details

The total loss function is $L = L_R + \lambda_1 L_{reg} + \lambda_2 L_t$. We empirically find that $\lambda_1 = 2$ and $\lambda_2 = 5$ are good choices. Initial learning rate is set to 0.0001, and decayed by a rate of 0.6 and 0.5 when ADD(-S) metric on a validation set is less than 0.016 and 0.013 respectively. We randomly sample 500 and 1,000 foreground points for each object in LINEMOD and YCB-Video datasets as inputs. 3D RANSAC-voting layer is implemented using CUDA. Each iteration parallelly generates 128 hypotheses. The maximum number of iterations is 20, resulting in 2,560 hypotheses altogether at most. Early stop requires the probability of success above 0.99. Probability of choosing an inlier is set to the maximum inlier ratio of currently generated hypotheses.

D. Ablation Analysis

Extensive ablation studies are conducted to compare different design choices. Table I summaries the evaluation results. We incrementally modify the network of DenseFusion [13] and add in our innovations. First of all, we explore

TABLE II
COMPARISON ON LINEMOD DATASET. OBJECTS MARKED WITH * ARE CONSIDERED TO BE SYMMETRIC.

Method	RGB w/o refinement			RGB w/ refinement		RGB-D w/o refinement		RGB-D w/ refinement
	PoseCNN [5]	PVNet [7]	DPOD [8]	DeepIM [41]	DPOD+ [8]	Per-Pixel DF [13]	Ours	Iterative DF [13]
ape	21.62	43.62	53.28	77.0	87.73	79.5	85.03	92.3
benchvise	81.80	99.90	95.34	97.5	98.45	84.2	95.54	93.2
cam	36.57	86.86	90.36	93.5	96.07	76.5	91.27	94.4
can	68.80	95.47	94.10	96.5	99.71	86.6	95.18	93.1
cat	41.82	79.34	60.38	82.1	94.71	88.8	93.61	96.5
driller	63.51	96.43	97.72	95.0	98.80	77.7	82.56	87.0
duck	27.23	52.58	66.01	77.7	86.29	76.3	88.08	92.3
eggbox*	69.58	99.15	99.72	97.1	99.91	99.9	99.90	99.8
glue*	80.02	95.66	93.83	99.4	96.82	99.4	99.61	100.0
holepuncher	42.63	81.92	65.83	52.8	86.87	79.0	92.58	92.1
iron	74.97	98.88	99.80	98.3	100.0	92.1	95.91	97.0
lamp	71.11	99.33	88.11	97.5	96.84	92.3	94.43	95.3
phone	47.74	92.41	74.24	87.7	94.69	88.0	93.56	92.8
Average	55.95	86.27	82.98	88.6	95.15	86.2	92.87	94.3

the 3D RANSAC-based voting technique for translation estimation. DenseFusion predicts an offset to object center at each point. While our network outputs an unit vector pointing to the center, which is equivalent to the normalized offset. The overall detection accuracy increases by 3.2%, showing that 3D voting yields more accurate results due to the robustness of RANSAC. Then, we validate the benefit of incorporating local geometry by applying EdgeConv [38] layer. The accuracy increases by another 0.9%. To analyze the proposed discrete-continuous formulation for rotation, we replace the per-point prediction in "local geometry" variant with our anchor-based prediction. Correspondingly, ShapeMatch-Loss is also replaced by our probabilistic loss during training. When there are 60 anchors, our best performance on LINEMOD dataset is achieved. Compared with the "local geometry", our accuracy is 2.4% higher.

We also train the network with 24 and 12 anchors respectively. The accuracy decreases considerably along with the number of anchors. Additionally, we test the case in which rotation and translation branches share one global feature. This single modification makes the accuracy drops by 0.9% from our best model, proving that rotation and translation pay different attentions to color and geometric features. If we remove the regularization on the range of deviations, which means at each anchor our network can predict arbitrary rotation in $SO(3)$, the accuracy decreases significantly (e.g. 8.1%). Local prediction is empirically proven to be the key ingredient of our discrete-continuous formulation. Last but not the least, we re-implement the regression method proposed in [12] and replace the corresponding part in our best model. On YCB-Video dataset, the accuracy decreases from 83.8% to 71.1%. The reasons are twofold: their mutually exclusive classification and loss function do not differentiate between symmetric and asymmetric objects; they do not enforce local prediction for each anchor.

E. Comparison on LINEMOD

We compare our method with the state-of-the-art object pose detectors [5], [7], [8], [13] on the LINEMOD dataset.

TABLE III
AUC IN TERMS OF ADD-S METRIC ON YCB-VIDEO DATASET. * INDICATES SYMMETRIC OBJECT.

Method	RGB-D			RGB-D w/ Refinement		
	MCN [12]	Per-Pixel DF [13]	Ours	PoseCNN +ICP [5]	MCN+ ICP [12]	Iterative DF [13]
002_master_chef_can	89.4	95.2	93.9	95.8	96.0	96.4
003_cracker_box	85.4	92.5	92.9	91.8	88.7	95.5
004_sugar_box	92.7	95.1	95.4	98.2	97.3	97.5
005_tomato_soup_can	93.2	93.7	93.3	94.5	96.5	94.6
006_mustard_bottle	96.7	95.9	95.4	98.4	97.7	97.2
007_tuna_fish_can	95.1	94.9	94.9	97.1	97.6	96.6
008_pudding_box	91.6	94.7	94.0	97.9	86.2	96.5
009_gelatin_box	94.6	95.8	97.6	98.8	97.6	98.1
010_potted_meat_can	91.7	90.1	90.6	92.8	90.8	91.3
011_banana	93.8	91.5	91.7	96.9	97.5	96.6
019_ptcher_base	93.8	94.6	93.1	97.8	96.6	97.1
021_bleach_cleanser	92.9	94.3	93.4	96.8	96.4	95.8
024_bowl*	82.6	86.6	92.9	78.3	76.0	88.2
025_mug	95.3	95.5	96.1	95.1	97.3	97.1
035_power_drill	88.2	92.4	93.3	98.0	95.9	96.0
036_wood_block*	81.5	85.5	87.6	90.5	93.5	89.7
037_scissors	87.3	96.4	95.7	92.2	79.2	95.2
040_large_marker	90.2	94.7	95.6	97.2	98.0	97.5
051_large_clamp*	91.5	71.6	75.4	75.4	94.0	72.9
052_extra_large_clamp*	88.0	69.0	73.0	65.3	90.7	69.8
061_foam_brick*	93.2	92.4	94.2	97.1	96.5	92.5
Average	90.6	91.2	91.8	93.0	93.3	93.1

Results are reported in Table II. Of all those methods without pose refinement, PVNet [7] and Per-Pixel DF [13] are the state of the art on LINEMOD dataset. PVNet utilizes a similar voting technique to detect keypoints on color image, and pose is obtained through solving a Perspective-n-Point problem. Although Per-Pixel DF uses the depth information, their performance is only comparable to PVNet. Comparing with them, our detection accuracy is 6.6% higher, and close to the results after pose refinement. Especially, our method can handle small objects very well, such as "ape" (5.5% higher), "duck" (11.8% higher) and "holepuncher" (13.6% higher). These small textureless objects are inherently diffi-




Fig. 3. **Accuracy-threshold curves.** Top-left: pose error of all objects measured by ADD(-S) metric. Top-right: translation error of all objects. Bottom-left: rotation error of all symmetric objects in terms of ADD-S metric. Bottom-right: rotation error of all asymmetric objects in terms of relative rotation angle. Dashed line indicates that methods have pose refinement.

cult for those methods based on establishing dense or sparse 2D-3D correspondences, even after pose refinement.

F. Comparison on YCB-Video

Table III and IV compare our method with [5], [9], [10], [12], [13] on the more challenging YCB-Video dataset. For fair comparisons, we use the segmentation masks from [5] and [13]. In Table III, our average AUC in terms of ADD-S metric is 0.6% higher than the state-of-the-art method (e.g. Per-Pixel DF [13]). Since the ADD-S metric computes the average distances between closest point pairs, regardless of whether the object is symmetric or not, the ability of different detectors is not well reflected through this evaluation. Therefore, we also compare the detection accuracy in terms of ADD(-S) metric. our method achieves a notably improvement over Per-Pixel DF (4.6% higher). In general, accuracies of RGB-D based methods are much higher than RGB-only methods, demonstrating the robustness induced by geometric features.

Fig.3 decouples the pose errors and provides detailed analysis on rotation and translation respectively. Our method is marginally better on estimating translation (top-right) and rotation of asymmetric objects (bottom-right). However, we achieve superior performance on estimating rotation of symmetric objects (bottom-left). Both PoseCNN and DenseFusion naively implement the ShapeMatch-Loss. Our remarkably higher AUC on symmetric objects shows that the proposed discrete-continuous formulation is able to handle the local-optimum problem of ShapeMatch-Loss. Furthermore, pose refinement can effectively refine a rotation only when the initial prediction is near the ground truth. For large errors which usually caused by local optimal predictions, the refinement performs no better than its baseline. To this

TABLE IV
AVERAGE ACCURACIES OF OUR METHOD AND BASELINE METHODS ON YCB-VIDEO DATASET IN TERMS OF ADD(-S) METRIC.

Method	RGB		Seg-Driven	RGB-D		RGB-D w/ Refinement	
	PoseCNN [5]	Heatmaps [9]	[10]	Per-Pixel DF [13]	Ours	PoseCNN +ICP [5]	Iterative DF [13]
Accuracy	22.9	53.1	39.0	79.2	83.8	73.7	85.6




Fig. 4. Qualitative results on YCB-Video dataset. Some cases in which our method (bottom row) can accurately predict the poses, while PoseCNN [5] (top row) and Per-Pixel DF [13] (middle row) gives inaccurate or local optimal predictions.

end, it is important to purposely deal with the local-optimum problem. Fig.4 shows some qualitative comparisons between our method, PoseCNN and Per-Pixel DF.

G. Runtime Analysis

On single Nvidia GPU (GTX 1080Ti), our pose estimation network infers at 0.04s per instance. Instance segmentation at the first stage takes about 0.03s per frame. By concurrently running multiple pose estimation networks (one for each object instance), our method could run at about 14 FPS (0.07s per frame), which is very promising for real-time applications.

V. CONCLUSION

We present a deep learning method for 6D object pose estimation. Specially, a novel discrete-continuous formulation is proposed for rotation to handle the local-optimum problem during training. In the future, we would consider to use the uncertainty values for pose refinement, as well as introduce them to grasping algorithms. Another extension is to test whether our approach can achieve similarly good results if trained purely on synthetic data.

ACKNOWLEDGMENT

This work was supported in part by the Singapore MOE Tier 1 grant R-252-000-A65-114 and the National Research Foundation, Prime Ministers Office, Singapore, under its CREATE programme, Singapore-MIT Alliance for Research and Technology (SMART) Future Urban Mobility (FM) IRG.

REFERENCES

- [1] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Analysis and observations from the first amazon picking challenge," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 1, pp. 172–188, 2016.
- [2] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *arXiv preprint arXiv:1809.10790*, 2018.
- [3] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1521–1529.
- [4] B. Tekin, S. N. Sinha, and P. Fua, "Real-time seamless single shot 6d object pose prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 292–301.
- [5] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *arXiv preprint arXiv:1711.00199*, 2017.
- [6] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 699–715.
- [7] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao, "Pvnet: Pixel-wise voting network for 6dof pose estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4561–4570.
- [8] S. Zakharov, I. Shugurov, and S. Ilic, "Dpod: Dense 6d pose object detector in rgb images," *arXiv preprint arXiv:1902.11020*, 2019.
- [9] M. Oberweger, M. Rad, and V. Lepetit, "Making deep heatmaps robust to partial occlusions for 3d object pose estimation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 119–134.
- [10] Y. Hu, J. Hugonot, P. Fua, and M. Salzmann, "Segmentation-driven 6d object pose estimation," *arXiv preprint arXiv:1812.02541*, 2018.
- [11] T. Hodan, F. Michel, E. Brachmann, W. Kehl, A. GlentBuch, D. Kraft, B. Drost, J. Vidal, S. Ihrke, X. Zabulis, et al., "Bop: benchmark for 6d object pose estimation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 19–34.
- [12] C. Li, J. Bai, and G. D. Hager, "A unified framework for multi-view multi-class object pose estimation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 254–269.
- [13] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese, "Densefusion: 6d object pose estimation by iterative dense fusion," *arXiv preprint arXiv:1901.04780*, 2019.
- [14] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 652–660.
- [15] D. Novotny, S. Albanie, D. Larlus, and A. Vedaldi, "Self-supervised learning of geometrically stable features through probabilistic introspection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3637–3645.
- [16] J. Li and G. H. Lee, "Usip: Unsupervised stable interest point detection from 3d point clouds," *arXiv preprint arXiv:1904.00229*, 2019.
- [17] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes," in *Asian conference on computer vision*. Springer, 2012, pp. 548–562.
- [18] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [19] D. G. Lowe et al., "Object recognition from local scale-invariant features," in *iccv*, vol. 99, 1999, pp. 1150–1157.
- [20] E. Brachmann, A. Krull, F. Michel, S. Gumhold, J. Shotton, and C. Rother, "Learning 6d object pose estimation using 3d object coordinates," in *European conference on computer vision*. Springer, 2014, pp. 536–551.
- [21] E. Brachmann, F. Michel, A. Krull, M. Ying Yang, S. Gumhold, et al., "Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3364–3372.
- [22] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [23] M. Rad and V. Lepetit, "Bb8: a scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3828–3836.
- [24] A. Krull, E. Brachmann, F. Michel, M. Ying Yang, S. Gumhold, and C. Rother, "Learning analysis-by-synthesis for 6d pose estimation in rgb-d images," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 954–962.
- [25] F. Michel, A. Kirillov, E. Brachmann, A. Krull, S. Gumhold, B. Savchynskyy, and C. Rother, "Global hypothesis generation for 6d object pose estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 462–471.
- [26] A. Tejani, D. Tang, R. Kouskouridas, and T.-K. Kim, "Latent-class hough forests for 3d object detection and pose estimation," in *European Conference on Computer Vision*. Springer, 2014, pp. 462–477.
- [27] W. Kehl, F. Milletari, F. Tombari, S. Ilic, and N. Navab, "Deep learning of local rgb-d patches for 3d object detection and 6d pose estimation," in *European Conference on Computer Vision*. Springer, 2016, pp. 205–220.
- [28] A. Dai and M. Nießner, "3dmv: Joint 3d-multi-view prediction for 3d semantic scene segmentation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 452–468.
- [29] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [30] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask r-cnn," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [31] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, "Path aggregation network for instance segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8759–8768.
- [32] Y. Yan and G. S. Chirikjian, "Almost-uniform sampling of rotations for conformational searches in robotics and structural biology," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 4254–4259.
- [33] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [34] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, "3d bounding box estimation using deep learning and geometry," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7074–7082.
- [35] F. Manhardt, D. M. Arroyo, C. Ruppert, B. Busam, N. Navab, and F. Tombari, "Explaining the ambiguity of object detection and 6d pose from visual data," *arXiv preprint arXiv:1812.00287*, 2018.
- [36] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [38] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *arXiv preprint arXiv:1801.07829*, 2018.
- [39] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.
- [40] T. Hodaň, J. Matas, and Š. Obdržálek, "On evaluation of 6d object pose estimation," in *European Conference on Computer Vision*. Springer, 2016, pp. 606–619.
- [41] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox, "Deepim: Deep iterative matching for 6d pose estimation," in *The European Conference on Computer Vision (ECCV)*, September 2018.