

6D Object Pose Estimation in Cluttered Scenes from RGB Images

Xiao-Long Yang^{1,2} (杨小龙), Xiao-Hong Jia^{1,2,*} (贾晓红), *Member, CCF*, Yuan Liang³ (梁 缘), and Lu-Bin Fan³ (樊鲁宾)

¹*Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China*

²*University of Chinese Academy of Sciences, Beijing 100049, China*

³*Alibaba DAMO Academy, Alibaba Group, Hangzhou 311121, China*

E-mail: yangxiaolong17@mails.ucas.ac.cn; xhjia@amss.ac.cn; {liangyuan.ly, lubin.flb}@alibaba-inc.com

Received January 22, 2021; accepted August 31, 2021.

Abstract We propose a feature-fusion network for pose estimation directly from RGB images without any depth information in this study. First, we introduce a two-stream architecture consisting of segmentation and regression streams. The segmentation stream processes the spatial embedding features and obtains the corresponding image crop. These features are further coupled with the image crop in the fusion network. Second, we use an efficient perspective-n-point (E-PnP) algorithm in the regression stream to extract robust spatial features between 3D and 2D keypoints. Finally, we perform iterative refinement with an end-to-end mechanism to improve the estimation performance. We conduct experiments on two public datasets of YCB-Video and the challenging Occluded-LineMOD. The results show that our method outperforms state-of-the-art approaches in both the speed and the accuracy.

Keywords two-stream network, 6D pose estimation, fusion feature

1 Introduction

The 6D object pose estimation has been widely used in computer vision tasks in daily life, such as robot grasping and manipulation, autonomous navigation and augmented/mixed reality, with the continuous development of sensor technology. High requirements are put forward for the following applications for various real-world needs: 1) proper handling of objects with irregular shapes, low-resolution textures, and different materials; 2) robustness to heavy occlusion, lighting changes in various environments, and potential technical noise; 3) real-time speed as possible.

Many RGB-D algorithms based on depth information can infer the pose relatively accurately when texture features of the object are blurred or the scene is blocked by other objects. The traditional method extracts features from RGB-D data, combines them with depth masks, and performs verification and

inference^[1–8]. However, reliance on manual functions and fixed matching procedures limits the empirical performance in the case of severe occlusion and lighting changes. RGB-D-based algorithms are generally restricted applications that require the hardware of the RGB-D sensor. Hence, the majority of the existing RGB-D-based methods present difficulty in meeting the requirements of accurate pose estimation and fast inference^[9].

We focus on the methods of image-based 2D recognition of object poses due to the increasing demands of source acquisition and the limitation of depth sensors. Traditional RGB-based methods rely on the correspondence between 3D and 2D keypoints. First, we find unified feature points through the known 3D model and 2D pixels of the object and then calculate regression parameters of rotation and translation for further processing using the perspective-*n*-point (PnP) algorithm^[10].

Regular Paper

A preliminary version of the paper was published in the Proceedings of ACM SIGGRAPH 2020.

This work was partially supported by the National Key Research and Development Program of China under Grant No. 2021YFB1715900, the National Natural Science Foundation of China under Grant Nos. 12022117 and 61802406, the Beijing Natural Science Foundation under Grant No. Z190004, the Beijing Advanced Discipline Fund under Grant No. 115200S001, and Alibaba Group through Alibaba Innovative Research Program.

*Corresponding Author

©Institute of Computing Technology, Chinese Academy of Sciences 2022

This approach is robust when the object has fruitful textures. However, it might fail if the texture becomes blurred or multiple objects block one another^[11–14]. Recent approaches have attempted to overcome these difficulties that usually apply deep neural networks to return 6D poses from images directly or detect keypoints related to objects. However, both situations endow the object with the property of the global entity, which makes the procedure susceptible to heavier occlusion.

In this paper, we propose a novel method for generating the high-quality real-time pose estimation of RGB images to address this occluded problem. The key idea is to build a two-stream architecture, which includes a segmentation stream to process various features of the input image and a regression stream to solve the 2D-3D correspondences via bounding corners. Different from existing methods that use image segmentation^[15] and bounding boxes^[16] to calculate global features separately, we further propose a fusion network for transforming and mapping RGB textures and location relationship at the per-pixel level. Abundant features ensure that our model considers both the texture and geometric information of the object. Moreover, we design an iterative refinement process with the end-to-end mechanism, which remarkably improves the accuracy of estimated poses. An example of pose estimation using our approach is shown in Fig.1.




Fig.1. Example of pose estimation using our approach. (a) Input image with occluded known objects. (b) Results of object detection and segmentation. (c) Locations of keypoints. (d) Final results of the pose estimation.

We conduct various comparisons to demonstrate the superiority of our approach over many representative competitors in terms of the pose estimation performance and efficiency. Our method presents superior performance on popular YCB-Video^[11] and Occluded-LineMOD^[12] datasets. Notably, our method still produces reliable estimation results when dealing with challenging samples in heavily cluttered scenes. Finally, we illustrate benefits of the iterative refinement process through an ablation study. This paper extends our recent ACM SIGGRAPH publication^[17], and shows the following additional contributions:

- additional qualitative and quantitative experiments to prove the superiority of our pose estimation method;
- improved loss function and other elaborate descriptions with further advanced performance;
- extended application: advertising replacement and wall decoration suggestion.

2 Related Work

We briefly discuss two relevant types of studies based on the input, that is, RGB images with or without depth.

2.1 Pose Estimation Based on RGB-D Images

Recent methods are typically data-driven. Song and Xiao^[18,19] enhanced 3D object detection through the discretization of 3D voxel spaces. The researchers proposed 3D BBox for the voxelized 3D model input and estimated poses. Although geometric information is encoded effectively, these time-consuming and space-expensive methods based on the voxel representation take nearly 20 seconds and 300 M+ memory for a single model frame in^[19].

Other methods directly use 3D deep learning architectures from 3D point cloud data and perform 6-DoF detection pose estimation. On the basis of the pioneering work of PointNet^[13], both VoxelNet^[15] and Frustrum^[16] have demonstrated considerable progress in point cloud detection and surpassed a large number of competitors on the KITTI benchmark^[20].

Visual-recognition deep neural networks also perform pose estimation on the basis of RGB-D input. These methods require detailed steps of feature extraction and training to optimize 3D prior labels. For instance, DenseFusion^[9] uses the RGB-D data color and depth information of pixels, and PoseCNN^[11] applies the closest point (ICP) process iteratively.

The acquisition and the widespread use of these methods are obstacles to real-world applications because they require many sources with depth information. By comparison, our method relies only on RGB images to achieve high-quality pose recognition and real-time performance.

2.2 Pose Estimation Based on RGB Images

Existing methods based on RGB images are mainly divided into two categories. Traditional methods depend on keypoint detection and matching relationship with known object models [21–26]. Other methods, which use learn mechanisms to predict 2D keypoints, are presented to solve poses via the PnP algorithm [10] and address this challenge [1, 27–29]. Although being excellent and efficient in some tasks that require speed performance, the robustness and the accuracy of these methods reduce when faced with low-texture or low-resolution input images. Additional studies have focused on directly using CNN-based architectures for pose estimation due to deep learning [30, 31]. These studies mainly obtain different forms of orientation estimation based on RGB images. Xiang *et al.* [11] learned and clustered 3D features of the object model, and then extracted the orientation information via a viewpoint-aware strategy. Mousavian *et al.* [32] attempted to obtain geometric constraints and exploit 3D object parameters to restore poses while concentrating on a single-view. These methods present poor accuracy under low-texture or low-resolution input despite their important advantages in speed performance. Therefore, other technologies focus on the use of powerful deep learning based on CNN architectures to obtain target pose estimation directly from a single image [30, 31], and constantly concentrate on the prediction of orientation information. Xiang *et al.* [11] attempted to learn 3D features, cluster object models, and obtain perceptual predictions on the basis of viewpoints. Mousavian *et al.* [32] optimized the single view to extract potential geometric constraints, recover 3D parameters, and achieve pose estimation. Sundermeyer *et al.* [33] proposed a novel encoding method to convert orientation information into feature vectors implicitly, and then determine the optimal match for the test object frame.

Many methods focus on local information to solve the occlusion and improve accuracy performance. Seg-Driven [14] crops the input image in different sizes, and the final result is weighed by combining multiple local pose estimations. SilhoNet [34] continuously obtains the

direction information using the contour scanning mechanism. Pix2Pose [35] recovers the correspondence between 3D and 2D from the pixel level through the GAN generation technology. Although these studies have achieved excellent experimental design and acceleration network, they still exhibit poor accuracy. ACC [36] attempts to address this problem by reconstructing the mesh of the entire original model on the basis of the optimized DCNN model [37, 38]. However, this method can only be applied to a single target and demonstrates poor generalization in challenging situations.

Our method is inspired by DenseFusion [9], which proposes a heterogeneous architecture and jointly considers potential features of geometry and texture. However, the fusion method in our approach is used for spatial coordinates and object appearance textures without depth information following our conference publication [17]. We design a new loss function and perform extensive evaluation experiments. We show that our novel feature extraction and mapping scheme outperforms algorithms that rely solely on PnP or image cropping in this study. Furthermore, we introduce a general end-to-end refinement strategy to enhance our performance.

3 Methodology

In this paper, we address challenges with high occlusion or under barren light with a heterogeneous network (Subsection 3.2) that processes texture and location information differently based on DenseFusion [9] and our conference publication [17]. We also design a new and completely independent coupling method (Subsection 3.3). Moreover, we build a two-stream network based on image cropping and bounding box through an encoding decoding process, which shares some similarities with Seg-Driven [14] in segmentation processing (Subsection 3.4). Finally, the estimator of poses will be further optimized with a differentiable module (Subsection 3.5). Compared with the previous expensive optimization procedures [11, 39], our module accounts for only a fraction of the total reasoning time.

3.1 Architecture

Fig. 2 depicts the overall workflow of our method. The segmentation stream consists of the following steps: conducting semantic segmentation and cropping in each input frame. The existing technology can achieve excellent performance for known objects. We then feed segmentation labels for each segmented local RGB patch




Fig.2. Proposed network architecture. A two-stream architecture network is designed for feature processing, including segmentation and regression. A fusion network that can combine texture and spatial features also exists.

with the bounding box of each cropped image and the entire original image with different sizes to the fusion network, which couples spatial features after segmentation with texture features of the image itself.

The regression stream performs an efficient PnP ($EPnP$) solution on spatial channels (XYZ) of the obtained fusion feature, and returned 2D-3D correspondences (rotation R and translation t) and texture channels (RGB) jointly participate in confidence calculation. We can treat these correspondences as rough pose estimates according to the magnitude of the confidence and design a subsequent refinement iteration for performance improvement, as illustrated in Fig.3.

3.2 Segmentation Stream

We use classic backbone, which consists of two streams of segmentation and regression. We use the highly effective and efficient YOLOv3^[40] and

YOLOv4^[41] for initial object detection and obtain many areas in the segmentation. We then crop each detected area on the original input image in the segmentation stream. Overlapping may occur in divided parts due to the occlusion between objects. Furthermore, we label every unit of the $n \times n$ grid with the object or background information. Specifically, we can generate accurate semantic labels using the real 3D model and the corresponding depth information originally contained in datasets when training the model to greatly reduce the impact of occlusion on the image. The uneven distribution of positive and negative samples in practice is due to the significantly smaller area occupied by the investigated object than the background. Therefore, we utilize Focal Loss^[42] to address negative effects of sample distribution on the model training. Moreover, we apply a pixel-wise median frequency balancing technique^[38] because target objects demonstrate different sizes and changing the ratio of




Fig.3. Refinement procedure with the end-to-end mechanism. This additional iterative optimization directly follows the main network.

cropped images will affect the performance of the pose estimation task.

3.3 Fusion Network

We will fuse individual multi-dimensional information for different detected objects with various sizes. The existing method DenseFusion^[9] also considers the fusion of features, and links the dense depth and color features on cropped images through a multilayer perceptron (MLP) to form a new global feature. However, this method retains unnecessary errors generated by the feature extraction in the previous step, such as the pixel information of other objects caused by occlusion and overlapping or the background information due to the segmentation accuracy. Therefore, undesigned fusion features will reduce the performance of pose estimation. A novel pixel-wise^[9] strategy is used in our implementation to achieve an efficient fusion network, especially under heavy occlusion and imperfect segmentation situations.

The proposed fusion network based on PointNet^[13] supplements texture features that correspond to spatial features on the processing object. Our fusion network performs local fusion and prediction on each $n \times n$ grid. Hence, we design different prediction weights according to visible and invisible parts of targets to minimize and decrease the negative influence of heavy occlusion and imperfect segmented noises.

We use the projection transformation to couple the spatial features and texture features on each grid, based on the real 3D model, ground-truth pose labels, and uv texture contained in original datasets combined with known internal camera parameters. We then expect to

generate dense feature vectors with a uniform size after passing the fusion network. Fig.4 shows the architecture of the fusion network in detail. Two features results with different dimensions are produced when the cropped image, which is resized to 128×128 , passes through a heterogeneous neural network. The upper network will obtain $(m+1)$ -channelled spatial information based on the background and occlusion whereas the lower network will obtain the texture information based on texture and appearance, and then couple them to obtain 128×6 fusion features, including XYZ and RGB intrinsic properties.

3.4 Regression Stream

Regression stream plays an important role in using the efficient PnP algorithm to obtain spatial features from the feature vector and link the position information of 2D and 3D key points. Following Seg-Driven^[14], we set key points to eight vertices of bounded boxes of each target, shown in Fig.5. We focus on the anchor point of the center in the bounding cube and calculate the offset vector and texture information deviation of the center and the corner points for every vertex.

C denotes the 2D coordinate of each anchor center point. An offset estimation $f_i(C)$ is proposed for the i -th vertex. Hence, we can express the vertex as $C + f_i(C)$ with the precise 2D coordinate C_i^{GT} from the original dataset. The texture of anchor point T , a corresponding offset $f_i(T)$ and precise texture T_i^{GT} are expressed as follows:

$$\begin{aligned}\Delta_i(C) &= C + f_i(C) - C_i^{\text{GT}}, \\ \Delta_i(T) &= T + f_i(T) - T_i^{\text{GT}}.\end{aligned}$$




Fig.4. Architecture of the fusion network. This detailed flowchart explains the process of extracting spatial and texture information on the basis of heterogeneous CNN, and then feeds them into another one to obtain fusion features. $m+1$ refers to the $(m+1)$ -channelled semantic segmentation map in Subsection 3.2 and 6 includes XYZ and RGB.



Fig.5. Simple example of our seg-reg two-stream network. (a) The segmentation stream labels every grid of the cropped image. (b) Offset vector and texture deviation of the center and corner points in the regression stream.

The novel loss function is expressed as follows:

$$E_{\text{pos}} = \sum_{\text{Grid } i=1}^n \|\Delta_i(C)\|_1 + \|\Delta_i(T)\|_1. \quad (1)$$

Here, we use the L_1 loss function instead of the L_2 one, because the former is less sensitive with ill-posed samples than the latter. We set equal weights because both spatial (x, y, z) and texture (r, g, b) features present the same dimensions.

We use the predicted multiple sets of EPnP results (subscript i), rotation matrix $\mathbf{R}$ and translation matrix $\mathbf{t}$, and their corresponding ground truth (superscript GT) to construct the reconstruction error:

$$\mathcal{L}_i = \frac{1}{8} \sum_n \|\mathbf{R}^{\text{GT}} p_n + \mathbf{t}^{\text{GT}} - (\mathbf{R}_i p_n + \mathbf{t}_i)\|, \quad (2)$$

where p_n denotes the selected n -th 3D point of corners from the real 3D model. However, we want to decide which pose estimation may be the optimal hypothesis because the PnP algorithm may lead to multiple sets of approximate solutions. Thus, we need to use RANSAC to balance the confidence among each prediction. A sigmoid-based function w is used as the activation function to produce the confidence Con_i for any group of EPnP algorithm results. Due to (2), the loss function to minimize becomes:

$$\mathcal{L} = \frac{1}{N} \sum_i (\mathcal{L}_i Con_i - w \log(Con_i)), \quad (3)$$

where Con_i represents the score of approximation between our pose estimation and the true three-dimensional transformation. After obtaining Con_i in (3), we then design a new loss penalty term as follows:

$$E_{\text{pro}} = \sum_{\text{Grid } i=1}^n \left\| -\exp(-\Theta \|\Delta_i(C) + \Delta_i(T)\|_2) + Con_i \right\|_1, \quad (4)$$

where Θ is a trade-off parameter. Finally, taking (1) and (4) together, the regression loss becomes:

$$E = \omega_{\text{pos}} E_{\text{pos}} + \omega_{\text{pro}} E_{\text{pro}}, \quad (5)$$

where ω_{pos} and ω_{pro} weight the effects of these two losses.

The large number of result candidates generated by the iterative algorithm for every target will consume unnecessary computing resources. Hence, we use the EPnP algorithm with the RANSAC mechanism ($n = 10$) [43] to obtain the 6D pose with consideration for efficiency. Fig.6 illustrates the 2D-3D corresponding procedure between the image and the 3D model.




Fig.6. Process of regression stream. (a) Detected Coke can use the segmentation stream. (b) The EPnP algorithm will generate many series of pose estimations, shown as pink dots. (c) We select the $n = 10$ groups with the maximum confidence for RANSAC repetition. (d) Contour result after iterative optimization.

3.5 Iterative Refinement

We propose an iterative optimal procedure with the network architecture that evidently advances the pose estimation performance in a robust and effective manner.

Optimization adjustment aims to reduce the noise caused by calculation and trade-off as extensively as possible through the iterative form and improve the

performance of pose estimation. New predictions from the beginning are unnecessary because we have obtained a satisfactory result. However, we add a module behind the main network to polish the previous results. Therefore, we use the result of the main network as the initial value, perform refinement iterations, and then utilize the iteration result as the new input and repeat this step. The important operation of iterative optimization aims to combine the 3D model from the original dataset with the obtained pose estimation, which presents the maximum confidence, and conducts a projection transformation to output the corresponding image. The image continues to iterate in the main network due to the segmentation stream to improve the accuracy of performance.

This module is illustrated in Fig.3. We obtain the pose estimation after K ($K = 2$ in our experiment) iterations as follows:

$$Pos = M_{R_K, t_K} \cdot M_{R_{K-1}, t_{K-1}} \cdots M_{R_1, t_1} \cdot M_{R_0, t_0}, \quad (6)$$

where M_{R_K, t_K} represents the K -th Euclidean transformation, including rotation and translation.

4 Experimental Results

We introduce the evaluation on the two popular datasets of Occluded-LINEMOD^[12] and YCB-Video^[11] that outperforms those of other methods. The YCB dataset consists of various models with distinguishable textures under different surroundings while presenting many rich labels. Hence, this information is a commonly used 6oF datasets in many existing methods. The Occluded-LINEMOD dataset is widely used in challenging situations with serious occlusion and messy background information. The performance of existing SOTA on this dataset is unsatisfactory.

We compare our proposed approach with the following recent advanced methods: PoseCNN^[11], Seg-Driven^[14], Tekin^[28], SilhoNet^[34], Pix2Pose^[35], BB8^[44], Heatmaps^[45], PVnet^[46], CDPN^[47], and our

previous work (OCP)^[17]. The commonly-used metric ADD-S^[11] proposed by the author of the YCB-Video dataset^[11] is used in this study. We set ADD-0.1d in our experiments to adjust whether a pose estimation will be positive under the condition that the metric should be below 10% of the model diameter.

4.1 Ablation Study on Different Architectures

We investigate the effectiveness of different architectures in this subsection. The architecture we propose can differ in the following aspects: 1) fusion feature, and 2) iterative refinement. We compare different features in the first part and omit the second item $\Delta_i(c)$ related to texture information in (1) (Subsection 3.4). The module that only uses spatial features is similar to the Seg-Driven^[14] method. Both approaches rely only on segmenting images and use a regression algorithm (e.g., EPnP) to calculate the 2D-3D correspondence. The experimental results of different modules are reported in Table 1. The fusion feature (the third row) strongly promotes the improvement of performance based on only spatial features (the first row). The symmetric Eggbox model with unclear texture may cause a slight drop in performance, but the average result of the fusion feature presents an increasing of more than 5% compared with the use of spatial feature alone. The iterative refinement part clearly shows that the results with iterative refinement (even rows) are constantly better than those without iterative refinement (odd rows). The ablation study proves the superiority of the fusion feature and iterative refinement module.

4.2 Quantitative Evaluation on Occluded-LINEMOD

We produce deep-fake images on random PASCAL VOC^[48] images to find suitable training and testing objects. Fig.7 illustrates some sampled results of our algorithm. Table 2 presents the quantitative evaluations on Occluded-LINEMOD^[12].

Table 1. Ablation Study on Occluded-LINEMOD^[12]

Method	SF	TF	IR	Ape	Can	Cat	Driller	Duck	Eggbox	Glue	Holepun	Average
XYZ only	✓			14.9	60.1	16.5	48.2	25.1	35.6	44.4	35.9	35.1
XYZ+refine	✓		✓	17.5	63.1	16.6	48.3	25.4	35.8	44.6	38.6	36.2
Fusion	✓	✓		37.4	68.6	26.0	48.5	25.0	30.3	45.2	45.1	40.8
Fusion+refine	✓	✓	✓	39.1	69.8	26.9	49.0	25.6	31.9	46.5	53.1	42.7

Note: The metric is ADD-0.1d. Symmetric objects have bold-name. SF: spatial feature; TF: texture feature; IR: iterative refinement.




Fig.7. Results on the Occluded-LINEMOD dataset. The entire scene is largely obscured and presents challenging camera viewpoints.

Table 2. Quantitative Evaluation on the Occluded-LINEMOD Dataset^[12]

Method	Ape	Can	Cat	Driller	Duck	Eggbox	Glue	Holepun	Average
PoseCNN ^[11]	9.60	45.20	0.90	41.40	19.60	22.00	38.50	22.10	24.90
Seg-Driven ^[14]	12.10	39.90	8.20	45.20	17.20	22.10	35.80	36.00	27.00
Tekin ^[28]	7.00	1.20	3.60	1.40	5.10	9.60	6.50	8.30	5.30
Pix2Pose ^[35]	8.30	12.10	9.30	10.90	6.30	13.80	11.30	10.70	10.30
BB8 ^[44]	28.50	11.20	9.60	0.20	6.80	4.00	4.70	8.10	9.10
Heatmap ^[45]	16.50	42.50	2.80	47.10	11.00	24.70	39.50	21.90	25.80
PVnet ^[46]	15.81	63.31	16.68	65.65	25.24	50.17	49.62	39.67	40.77
CDPN ^[47]	28.92	55.98	13.24	51.37	22.97	35.98	39.68	51.06	37.40
OCP ^[17]	28.77	49.81	6.94	48.90	22.61	11.92	16.51	53.10	29.82
Ours	39.10	69.80	26.90	49.00	25.60	31.90	46.50	53.10	42.70

Note: Symmetric objects are presented in bold font. Red and blue labels denote the best and the second best results, respectively.

Our method demonstrates superior performance with global inference technologies, such as PoseCNN, Tekin, BB8, and Pix2Pose. Compared with other methods, the proposed method also slightly outperforms Seg-Driven, Heatmaps, and PVnet, which use deep information in the data-build stage. Due to using a more reasonable loss function design and making full use of 3D information of the iterative module, our results also bring a significant performance improvement compared with the conference version (OCP)^[17], with the averaged performance from 29.82 to 42.70. Furthermore, we compare the time efficiency in Table 3. Our method is approximately five times faster than other methods due to its succinct network architecture and EPnP-RANSAC mechanism. Moreover, our approach only needs less than 30 ms for each image. Notably, our proposed approach shows high accuracy and efficiency even with large occlusions.

4.3 Quantitative Evaluation on YCB-Video

Fig. 8 presents additional test results in a large number of scenarios. Our method exhibits satisfactory estimation results in all tests. We compare our

method with the baseline PoseCNN and four latest approaches Seg-Driven, Heatmap, CDPN and SilhoNet. Our method outperforms both PoseCNN and SilhoNet. Furthermore, our method is only slightly better than Heatmap and Seg-Driven because symmetrical objects have similar appearance and are texture-less.

Table 3. Efficiency Comparison on the Occluded-LINEMOD Dataset

Method	FPS
PoseCNN	4
Tekin	40
BB8	3
Pix2Pose	—
Heatmap	4
Seg-Driven	22
PVnet	8
Ours	30

4.4 Ablation Study on Iterative Refinement

Fig.9 illustrates some sampled predictions and comparisons, which intuitively shows the effect with and without iterative optimization ($K = 1$). The results show that our refinement module can improve the performance of the original main network.




Fig.8. YCB-Video results. (a)–(f) Each pair of images represents different perspectives of the same scene. Our method generates accurate pose estimations, even in the presence of large occlusions. Furthermore, it can process multiple objects in real time.

Table 4. Quantitative Evaluation on the YCB-Video Dataset

Object	PoseCNN	SilhoNet	Heatmap	Seg-Driven	CDPN	Ours
master_chef_can	3.6	23.8	32.9	33.0	35.5	37.1
cracker_box	25.1	20.1	62.6	44.6	45.6	45.8
sugar_box	40.3	48.5	44.5	75.6	71.5	69.4
tomato_soup_can	25.5	25.1	31.1	40.8	49.6	52.3
mustard_bottle	61.9	60.8	42.0	70.6	74.8	78.2
tuna_fish_can	11.4	25.3	6.8	18.1	25.2	24.1
pudding_box	14.5	17.0	58.4	12.2	48.4	32.6
gelatin_box	12.1	26.2	42.5	59.4	57.8	46.9
potted_meat_can	18.9	22.2	37.7	33.3	37.6	40.1
banana	30.3	32.8	16.8	16.6	25.1	27.5
pitcher_base	15.6	25.9	57.2	90.0	80.8	82.0
bleach_cleanser	21.2	20.8	65.3	70.9	81.9	82.1
bowl	12.1	22.5	25.6	30.5	22.6	23.0
mug	5.2	12.3	11.6	40.7	51.2	55.6
power_drill	29.9	26.0	46.1	63.5	57.4	59.8
wood_block	10.7	18.7	34.3	27.7	25.1	27.5
scissors	2.2	3.4	0.1	17.1	11.6	12.4
large_marker	3.4	3.0	3.2	4.8	4.5	6.1
large_clamp	28.5	29.7	10.8	25.6	25.9	27.7
extra_large_clamp	19.6	20.4	29.6	8.8	13.4	16.4
foam_brick	54.5	42.0	51.7	34.7	40.7	44.8
Average	21.3	25.1	33.6	39.0	42.2	42.4

Note: The metric is ADD-0.1d. Symmetric objects are presented in bold font. Red and blue labels denote the best and the second best results.

4.5 Applications

The 6D object pose estimation is a key technology for real-world applications, as mentioned in Section 1. We apply our method in object placement. We skin known textures at a specified location and pretend that the object is in the dataset because our method targets known objects. We translate and rotate the target object according to the pose, and then place it in the specified position after estimating the pose by our algorithm. As shown in Fig.10, Fig.10(a) is the advertising space of the most popular e-sports game in the world, and all viewers will see it during the game. Fig.10(b) shows that the official game logo is placed on the adver-

tising space. The image in Fig.10(c) uses our algorithm with the texture of the cracker box in the YCB dataset.

Notably, many recommended decoration algorithms can also be viewed as an implementation of object placement. Liang *et al.* [49] determined the optimal decoration recommendation on the wall. Meanwhile, we also applied our proposed method to achieve the same performance in the real environment.

5 Conclusions

We proposed a method for object pose estimation from RGB images in occluded scenes. Our approach can simultaneously process spatial features of the ob-

ject position and geometric features of the appearance texture. We achieved high-precision object pose estimation without relying on depth information even in challenging situations such as heavy occlusion or poor lighting conditions.




Fig.9. Comparison of the iterative refinement module. (a) Original estimation of the main network. (b) Results presenting the benefits of our end-to-end procedure for large occlusions.




Fig.10. Examples of object placement. The top row is found in virtual media: (a) advertising space, (b) official logo implant, and (c) our result. The bottom row is located in the real environment: (d) advertising space, (e) recommended decoration algorithm, and (f) our result.

However, our method may produce unsatisfactory results on objects with symmetry or texture-less models. Fig.11 shows two examples of this scenario. If an excessive number of occlusions exist in the scene, our method fails to obtain sufficient information to estimate accurate poses. This issue will be our main concern in a future investigation.




Fig.11. Unsatisfactory examples (highlighted in the green box). (a) Heavily occluded symmetrical Eggbox model in the Occluded-LINEMOD dataset. (b) Textureless symmetrical Bowl model in the YCB-Video dataset.

References

- [1] Brachmann E, Krull A, Michel F, Gumhold S, Shotton J, Rother C. Learning 6D object pose estimation using 3D object coordinates. In *Proc. the 13th European Conference on Computer Vision*, Sept. 2014, pp.536-551. DOI: [10.1007/978-3-319-10605-2.35](https://doi.org/10.1007/978-3-319-10605-2.35).
- [2] Hinterstoisser S, Holzer S, Cagniart C, Ilic S, Konolige K, Navab N, Lepetit V. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *Proc. the 2011 IEEE International Conference on Computer Vision*, Nov. 2011, pp.858-865. DOI: [10.1109/ICCV.2011.6126326](https://doi.org/10.1109/ICCV.2011.6126326).
- [3] Hinterstoisser S, Lepetit V, Ilic S, Holzer S, Bradski G, Konolige K, Navab N. Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes. In *Proc. the 11th Asian Conference on Computer Vision*, Nov. 2012, pp.548-562. DOI: [10.1007/978-3-642-37331-2.42](https://doi.org/10.1007/978-3-642-37331-2.42).
- [4] Kehl W, Milletari F, Tombari F, Ilic S, Navab N. Deep learning of local RGB-D patches for 3D object detection and 6D pose estimation. In *Proc. the 14th European Conference on Computer Vision*, Oct. 2016, pp.205-220. DOI: [10.1007/978-3-319-46487-9.13](https://doi.org/10.1007/978-3-319-46487-9.13).
- [5] Rios-Cabrera R, Tuytelaars T. Discriminatively trained templates for 3D object detection: A real time scalable approach. In *Proc. the 2013 IEEE International Conference on Computer Vision*, Dec. 2013, pp.2048-2055. DOI: [10.1109/ICCV.2013.256](https://doi.org/10.1109/ICCV.2013.256).
- [6] Tejani A, Tang D, Kouskouridas R, Kim T K. Latent-class hough forests for 3D object detection and pose estimation. In *Proc. the 13th European Conference on Computer Vision*, Sept. 2014, pp.462-477. DOI: [10.1007/978-3-319-10599-4.30](https://doi.org/10.1007/978-3-319-10599-4.30).
- [7] Wohlhart P, Lepetit V. Learning descriptors for object recognition and 3D pose estimation. In *Proc. the 2015 IEEE Conference on Computer Vision and Pattern Recognition*, June 2015, pp.3109-3118. DOI: [10.1109/CVPR.2015.7298930](https://doi.org/10.1109/CVPR.2015.7298930).
- [8] Cao Y, Ju T, Xu J, Hu S. Extracting sharp features from RGB-D images. *Computer Graphics Forum*, 2017, 36(8): 138-152. DOI: [10.1111/cgf.13069](https://doi.org/10.1111/cgf.13069).

- [9] Wang C, Xu D, Zhu Y, Martin R, Lu C, Li F, Savarese S. DenseFusion: 6D object pose estimation by iterative dense fusion. In *Proc. the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019, pp.3343-3352. DOI: [10.1109/CVPR.2019.00346](https://doi.org/10.1109/CVPR.2019.00346).
- [10] Fischler M A, Bolles R C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981, 24(6): 381-395. DOI: [10.1145/358669.358692](https://doi.org/10.1145/358669.358692).
- [11] Xiang Y, Schmidt T, Narayanan V, Fox D. PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. In *Proc. the 14th Robotics: Science and Systems*, June 2018. DOI: [10.15607/RSS.2018.XIV.019](https://doi.org/10.15607/RSS.2018.XIV.019).
- [12] Krull A, Brachmann E, Michel F, Yang M Y, Gumhold S, Rother C. Learning analysis-by-synthesis for 6D pose estimation in RGB-D images. In *Proc. the 2015 IEEE International Conference on Computer Vision*, Dec. 2015, pp.954-962. DOI: [10.1109/ICCV.2015.115](https://doi.org/10.1109/ICCV.2015.115).
- [13] Qi C R, Su H, Mo K, Guibas L J. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, July 2017, pp.652-660. DOI: [10.1109/CVPR.2017.16](https://doi.org/10.1109/CVPR.2017.16).
- [14] Hu Y, Hugonot J, Fua P, Salzmann M. Segmentation-driven 6D object pose estimation. In *Proc. the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019, pp.3385-3394. DOI: [10.1109/CVPR.2019.00350](https://doi.org/10.1109/CVPR.2019.00350).
- [15] Xu D, Anguelov D, Jain A. PointFusion: Deep sensor fusion for 3D bounding box estimation. In *Proc. the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp.244-253. DOI: [10.1109/CVPR.2018.00033](https://doi.org/10.1109/CVPR.2018.00033).
- [16] Qi C R, Liu W, Wu C, Su H, Guibas L J. Frustum PointNets for 3D object detection from RGB-D data. In *Proc. the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp.7918-7927. DOI: [10.1109/CVPR.2018.00102](https://doi.org/10.1109/CVPR.2018.00102).
- [17] Yang X L, Jia X H. 6D pose estimation with two-stream net. In *Proc. the 2020 ACM SIGGRAPH*, Aug. 2020, Article No. 40. DOI: [10.1145/3388770.3407423](https://doi.org/10.1145/3388770.3407423).
- [18] Song S, Xiao J. Sliding shapes for 3D object detection in depth images. In *Proc. the 13th European Conference on Computer Vision*, Sept. 2014, pp.634-651. DOI: [10.1007/978-3-319-10599-4_41](https://doi.org/10.1007/978-3-319-10599-4_41).
- [19] Song S, Xiao J. Deep sliding shapes for Amodal 3D object detection in RGB-D images. In *Proc. the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, June 2016, pp.808-816. DOI: [10.1109/CVPR.2016.94](https://doi.org/10.1109/CVPR.2016.94).
- [20] Geiger A, Lenz P, Urtasun R. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proc. the 2012 IEEE Conference on Computer Vision and Pattern Recognition*, June 2012, pp.3354-3361. DOI: [10.1109/CVPR.2012.6248074](https://doi.org/10.1109/CVPR.2012.6248074).
- [21] Aubry M, Maturana D, Efros A A, Russell B C, Sivic J. Seeing 3D chairs: Exemplar part-based 2D-3D alignment using a large dataset of CAD models. In *Proc. the 2014 IEEE Conference on Computer Vision and Pattern Recognition*, June 2014, pp.3762-3769. DOI: [10.1109/CVPR.2014.487](https://doi.org/10.1109/CVPR.2014.487).
- [22] Collet A, Martinez M, Srinivasa S S. The MOPED framework: Object recognition and pose estimation for manipulation. *International Journal of Robotics Research*, 2011, 30(10): 1284-1306. DOI: [10.1177/0278364911401765](https://doi.org/10.1177/0278364911401765).
- [23] Ferrari V, Tuytelaars T, Gool L V. Simultaneous object recognition and segmentation from single or multiple model views. *International Journal of Computer Vision*, 2006, 67(2): 159-188. DOI: [10.1007/s11263-005-3964-7](https://doi.org/10.1007/s11263-005-3964-7).
- [24] Rothganger F, Lazebnik S, Schmid C, Ponce J. 3D object modeling and recognition using local affine-invariant image descriptors and multi-view spatial constraints. *International Journal of Computer Vision*, 2006, 66(3): 231-259. DOI: [10.1007/s11263-005-3674-1](https://doi.org/10.1007/s11263-005-3674-1).
- [25] Zhu M, Derpanis K G, Yang Y, Brahmabhatt S, Zhang M, Phillips C, Lecce M, Daniilidis K. Single image 3D object detection and pose estimation for grasping. In *Proc. the 2014 IEEE International Conference on Robotics and Automation*, May 31-June 7, 2014, pp.3936-3943. DOI: [10.1109/ICRA.2014.6907430](https://doi.org/10.1109/ICRA.2014.6907430).
- [26] Nakajima Y, Saito H. Robust camera pose estimation by viewpoint classification using deep learning. *Computational Visual Media*, 2017, 3(2): 189-198. DOI: [10.1007/s41095-016-0067-z](https://doi.org/10.1007/s41095-016-0067-z).
- [27] Suwajanakorn S, Snaveley N, Tompson J J, Norouzi M. Discovery of latent 3D keypoints via end-to-end geometric reasoning. In *Proc. the 2018 Annual Conference on Neural Information Processing Systems*, Dec. 2018, pp.2059-2070.
- [28] Tekin B, Sinha S N, Fua P. Real-time seamless single shot 6D object pose prediction. In *Proc. the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp.292-301. DOI: [10.1109/CVPR.2018.00038](https://doi.org/10.1109/CVPR.2018.00038).
- [29] Tremblay J, To T, Sundaralingam B, Xiang Y, Fox D, Birchfield S. Deep object pose estimation for semantic robotic grasping of household objects. In *Proc. the 2nd Conference on Robot Learning*, Oct. 2018, pp.306-316.
- [30] Schwarz M, Schulz H, Behnke S. RGB-D object recognition and pose estimation based on pre-trained convolutional neural network features. In *Proc. the 2015 IEEE International Conference on Robotics and Automation*, May 2015, pp.1329-1335. DOI: [10.1109/ICRA.2015.7139363](https://doi.org/10.1109/ICRA.2015.7139363).
- [31] Tulsiani S, Malik J. Viewpoints and keypoints. In *Proc. the 2015 IEEE Conference on Computer Vision and Pattern Recognition*, June 2015, pp.1510-1519. DOI: [10.1109/CVPR.2015.7298758](https://doi.org/10.1109/CVPR.2015.7298758).
- [32] Mousavian A, Anguelov D, Flynn J, Kosecka J. 3D bounding box estimation using deep learning and geometry. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, July 2017, pp.7074-7082. DOI: [10.1109/CVPR.2017.597](https://doi.org/10.1109/CVPR.2017.597).
- [33] Sundermeyer M, Marton Z C, Durner M, Brucker M, Triebel R. Implicit 3D orientation learning for 6D object detection from RGB images. In *Proc. the 15th European Conference on Computer Vision*, Sept. 2018, pp.699-715. DOI: [10.1007/978-3-030-01231-1_43](https://doi.org/10.1007/978-3-030-01231-1_43).
- [34] Billings G, Johnson-Roberson M. SilhoNet: An RGB method for 6D object pose estimation. *IEEE Robotics and Automation Letters*, 2019, 4(4): 3727-3734. DOI: [10.1109/LRA.2019.2928776](https://doi.org/10.1109/LRA.2019.2928776).

- [35] Park K, Patten T, Vincze M. Pix2Pose: Pixel-wise coordinate regression of objects for 6D pose estimation. In *Proc. the 2019 IEEE/CVF International Conference on Computer Vision*, Oct. 27-Nov. 2, 2019, pp.7668-7677. DOI: [10.1109/ICCV.2019.00776](https://doi.org/10.1109/ICCV.2019.00776).
- [36] Castro P, Armagan A, Kim T K. Accurate 6D object pose estimation by pose conditioned mesh reconstruction. arXiv:1910.10653, 2019. <https://arxiv.org/pdf/1910.10653.pdf>, Jan. 2022.
- [37] Kalchbrenner N, Grefenstette E, Blunsom P. A convolutional neural network for modelling sentences. arXiv:1404.2188, 2014. <https://arxiv.org/pdf/1404.2188.pdf>, Jan. 2022.
- [38] Badrinarayanan V, Kendall A, Cipolla R. SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017, 39(12): 2481-2495. DOI: [10.1109/TPAMI.2016.2644615](https://doi.org/10.1109/TPAMI.2016.2644615).
- [39] Li C, Bai J, Hager G D. A unified framework for multi-view multi-class object pose estimation. In *Proc. the 15th European Conference on Computer Vision*, Sept. 2018, pp.254-269. DOI: [10.1007/978-3-030-01270-0_16](https://doi.org/10.1007/978-3-030-01270-0_16).
- [40] Redmon J, Farhadi A. YOLOv3: An incremental improvement. arXiv:1804.02767, 2018. <https://arxiv.org/pdf/1804.02767.pdf>, Jan. 2022.
- [41] Bochkovskiy A, Wang C, Liao H M. YOLOv4: Optimal speed and accuracy of object detection. arXiv:2004.10934, 2020. <https://arxiv.org/pdf/2004.10934.pdf>, Jan. 2022.
- [42] Lin T, Goyal P, Girshick R, He K, Dollar P. Focal loss for dense object detection. In *Proc. the 2017 IEEE International Conference on Computer Vision*, Oct. 2017, pp.2980-2988. DOI: [10.1109/ICCV.2017.324](https://doi.org/10.1109/ICCV.2017.324).
- [43] Lepetit V, Moreno-Noguer F, Fua P. EPnP: An accurate $O(n)$ solution to the PnP problem. *International Journal of Computer Vision*, 2009, 81(2): Article No. 155. DOI: [10.1007/s11263-008-0152-6](https://doi.org/10.1007/s11263-008-0152-6).
- [44] Rad M, Lepetit V. BB8: A scalable, accurate, robust to partial occlusion method for predicting the 3D poses of challenging objects without using depth. In *Proc. the 2017 IEEE International Conference on Computer Vision*, Oct. 2017, pp.3828-3836. DOI: [10.1109/ICCV.2017.413](https://doi.org/10.1109/ICCV.2017.413).
- [45] Oberweger M, Rad M, Lepetit V. Making deep heatmaps robust to partial occlusions for 3D object pose estimation. In *Proc. the 15th European Conference on Computer Vision*, Sept. 2018, pp.119-134. DOI: [10.1007/978-3-030-01267-0_8](https://doi.org/10.1007/978-3-030-01267-0_8).
- [46] Peng S, Liu Y, Huang Q, Zhou X, Bao H. PVNet: Pixel-wise voting network for 6DoF pose estimation. In *Proc. the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019, pp.4561-4570. DOI: [10.1109/CVPR.2019.00469](https://doi.org/10.1109/CVPR.2019.00469).
- [47] Li Z, Wang G, Ji X. CDPN: Coordinates-based disentangled pose network for real-time RGB-based 6-DoF object pose estimation. In *Proc. the 2019 IEEE/CVF International Conference on Computer Vision*, Oct. 27-Nov. 2, 2019, pp.7678-7687. DOI: [10.1109/ICCV.2019.00777](https://doi.org/10.1109/ICCV.2019.00777).
- [48] Everingham M, Van Gool L, Williams C K I, Winn J, Zisserman A. The Pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 2010, 88(2): 303-338. DOI: [10.1007/s11263-009-0275-4](https://doi.org/10.1007/s11263-009-0275-4).


- [49] Liang Y, Fan L, Ren P, Xie X, Hua X. DecorIn: An automatic method for plane-based decorating. *IEEE Transactions on Visualization and Computer Graphics*, 2021, 27(8): 3438-3450. DOI: [10.1109/TVCG.2020.2972897](https://doi.org/10.1109/TVCG.2020.2972897).



Xiao-Long Yang received his B.S. degree in information and computing science from Northwestern Polytechnical University, Xi'an, in 2017, and is currently pursuing his Ph.D. degree at Key Laboratory of Mathematics Mechanization, Academy of Mathematics and Systems Science, Chinese Academy of Sciences (CAS), Beijing. His research interest is in computer graphics and computer vision.



Xiao-Hong Jia is a professor at Key Laboratory of Mathematics Mechanization, Academy of Mathematics and Systems Science, Chinese Academy of Sciences (CAS), Beijing. She received her Ph.D. and Bachelor's degrees in mathematics from the University of Science and Technology of China, Hefei, in 2009 and 2004, respectively. Her research interests include computer graphics, computer aided geometric design, and computational algebraic geometry.



Yuan Liang is an algorithm engineer at Alibaba DAMO Academy, Hangzhou. He received his B.Sc. and Ph.D. degrees in computer science from Tsinghua University, Beijing, in 2014 and 2019, respectively. His research interests include interactive multimedia analysis and geometric processing from visual media.



Lu-Bin Fan is a senior algorithm engineer at Alibaba DAMO Academy, Hangzhou. He received his B.Sc. and Ph.D. degrees in mathematics from Wuhan University, Wuhan, in 2009, and Zhejiang University, Hangzhou, in 2014, respectively. From 2014 to 2017, he was a postdoctoral fellow at the Visual Computing Center, King Abdullah University of Science and Technology (KAUST), Jeddah. He joined the Alibaba Group in 2017. His research interests include image and video processing, geometric processing and procedural modeling.