

6D Object Pose Estimation using Keypoints and Part Affinity Fields

Moritz Zappel, Simon Bultmann, and Sven Behnke

Autonomous Intelligent Systems, Computer Science Institute VI
University of Bonn, Germany
s6mozapp@uni-bonn.de, {bultmann, behnke}@ais.uni-bonn.de
<https://www.ais.uni-bonn.de>

Abstract. The task of 6D object pose estimation from RGB images is an important requirement for autonomous service robots to be able to interact with the real world. In this work, we present a two-step pipeline for estimating the 6 DoF translation and orientation of known objects. Keypoints and Part Affinity Fields (PAFs) are predicted from the input image adopting the OpenPose CNN architecture from human pose estimation. Object poses are then calculated from 2D-3D correspondences between detected and model keypoints via the PnP-RANSAC algorithm. The proposed approach is evaluated on the YCB-Video dataset and achieves accuracy on par with recent methods from the literature. Using PAFs to assemble detected keypoints into object instances proves advantageous over only using heatmaps. Models trained to predict keypoints of a single object class perform significantly better than models trained for several classes.

Keywords: object pose estimation · robot perception · deep learning.

1 Introduction

Object pose estimation is essential for autonomous robots to be able to interact with their environment. It has numerous real-world applications, such as robotic manipulation and human-robot interaction for autonomous service robots which are the focus of RoboCup@Home [17].

The task addressed in this work consists of detecting known objects and estimating their 6DoF orientation and translation in 3D space from a single RGB image. In recent years, two-stage approaches, which first detect keypoints and then solve a Perspective-n-Point (PnP) problem to infer the object pose [14,15,16], have been shown to provide robust and accurate results. However, keypoint detection remains difficult for occluded or truncated objects.

In this paper, we propose a two-stage pipeline for 6D object pose estimation in real-world scenes. We adopt the OpenPose architecture [2], well known from person pose estimation, to detect keypoints and Part Affinity Fields (PAFs) of everyday objects. Keypoints are predicted as local maxima of a heatmap indicating the confidence of a part being present at the image location. PAFs are

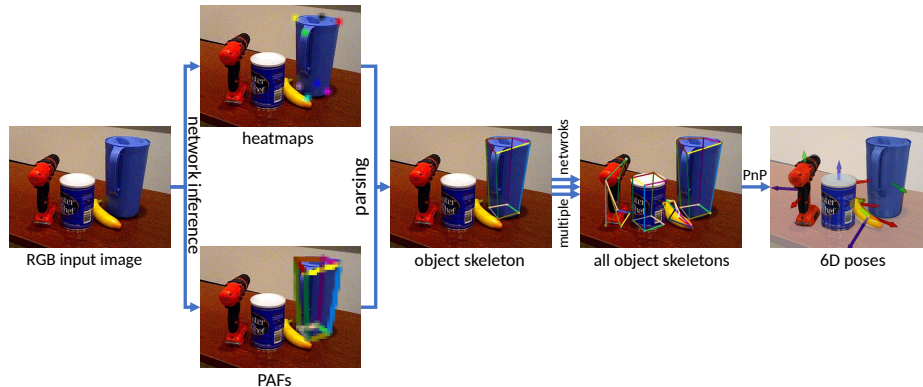


Fig. 1: Pose estimation approach: Heatmaps and PAFs are inferred for each object and 6D poses are calculated from detected keypoints via PnP.

vector fields connecting the keypoints of an object. They are used to assemble keypoints to object instances. As the OpenPose architecture is a bottom-up approach, keypoints are directly estimated from the input image. No prior object detection or segmentation is required, which is advantageous in terms of complexity and runtime [2]. Correspondences between the predicted 2D keypoints from the image and the keypoints defined on the 3D object models are then used to calculate the 6DoF object poses via the PnP-RANSAC algorithm. An overview of the proposed method is given in Fig. 1. Two different ways to define keypoints and PAFs on the object models are proposed in this work and the method is extensively evaluated on the challenging YCB-V dataset [19].

2 Related Work

6D object pose estimation methods from the literature can roughly be divided into two classes. Direct methods infer pose parameters directly from the image [9,19]. PoseCNN [19] defines a CNN architecture that segments objects in 2D images, predicts their depth, and regresses the 6D pose parameters. Capellen et al. [3] extend this approach to a fully convolutional network for dense prediction of pose parameters not depending on prior object segmentation. The direct pose regression is, however, difficult—especially for the rotation parameters as the 3D rotation space is highly nonlinear.

In contrast, Keypoint-based approaches adopt a two-step pipeline: First, 2D keypoints are predicted for each object instance in the image and the 6DoF pose is computed in a second step from 2D-3D correspondences with a variant of the PnP-Algorithm [11]. Approaches mainly differ in how the keypoints are defined on the object model and how they are inferred from the image. In BB8 [16], a segmentation mask is computed for each object and the keypoints are then inferred as the eight corners of the 3D object bounding box. The coordinates of the keypoints are directly regressed by the network. The corners of the 3D bounding box, however, often are not located on the object surface and are

thus difficult to infer from local object image features. Pavlakov et al. [14] define keypoints on the object surface and infer them as maxima of pixel-wise heatmaps. PVNet [15] also defines keypoints on the object surface but infers them in a dense manner: Each pixel in the object segmentation mask predicts vectors that point to every keypoint. The keypoint locations are then computed through RANSAC-based voting, choosing locations where the predicted directions intersect. This permits to also represent keypoints that are occluded or outside of the image.

Some approaches apply additional refinement after the initial estimation of the pose to further improve performance. Cosypose [10], e.g., implements an additional network that refines a given pose with the input image as additional input. When depth information is available, e.g., in the RGB-D variant of PoseCNN [19], ICP can also be used for pose refinement.

In this paper, we adopt a keypoint-based approach using keypoints on the object surface. The 6D object pose is computed via a combination of the PnP [11] and RANSAC [4] algorithms to increase the robustness of the estimation. Only RGB images are used as input and no further refinement steps are employed, keeping our pipeline simple and efficient. The OpenPose architecture [2] is adopted for 2D keypoint estimation. OpenPose is a keypoint-based bottom-up approach for human pose estimation in images. Together with heatmaps of the keypoints, the CNN computes vector fields, called Part Affinity Fields (PAFs), connecting the keypoints of an object instance. This permits to directly predict keypoints on the input image without prior segmentation or detection required. Local maxima in the heatmaps are then assembled into instances via the PAFs.

3 Method

In this section, we detail design choices and workflow of the proposed approach for 6DoF object pose estimation.

For evaluation, we employ the YCB-V dataset [19]. It comprises over 130k images at VGA resolution of 21 different object classes and is widely used for robot manipulation and object pose estimation tasks. For each object, a textured mesh is included as 3D model with the origin of the object coordinate frame defined at its center. All objects are household objects relevant for real-world robot experiments in domestic service scenarios (cf. Fig. 2). The images contain multiple objects in realistic settings with changing lighting conditions, significant image noise and cluttered backgrounds.

3.1 Selection of Keypoints and PAFs

The choice of keypoints and PAFs is an important design parameter of our method. They need to be well localized on the object geometry and texture, to facilitate their CNN-based detection, and should be spread out on the object surface such that a stable and well-defined solution of the PnP-problem can be found. Keypoints and PAFs are defined based on the 3D object models in two different ways: They are chosen manually or automatically. Eight keypoints and twelve PAFs are defined per object class. The manually defined keypoints are



Fig. 2: Objects of the YCB-V dataset with heatmaps of the manually defined keypoints and their interconnections.

located on easy-to-find spots of the object geometry and texture and represent the object contour. If applicable, the keypoints are placed to form a cuboid. This set of keypoints is shown in Fig. 2. The automatically defined set of keypoints is chosen with the farthest-point-algorithm, inspired by PVNet [15]: Starting with the object center, points on the object surface which are farthest from the already chosen points are added to the keypoint set. Eight points on the object surface are retained—the center point is not part of the final keypoint set. The set of automatically chosen keypoints is shown in Fig. 3. The PAFs for both sets of keypoints are defined by hand. The objective is to choose connections, which run along distinctive features and to form one upper and one lower polygon, which are connected with vertical PAFs. The PAFs run along the keypoint connections displayed in Figs. 2 and 3 and have a fixed width. The automatically picked keypoints are less intuitively placed and harder to find than the manually picked ones. The reason for this is, that the automatically picked keypoints are often located on edges instead of corners and on surfaces instead of edges. Furthermore, the texture is ignored in the automatic selection although it is important to localize keypoints. The inferior performance of the automatically chosen keypoints is confirmed by the evaluation results (cf. Sec. 4). Therefore, the manually chosen keypoints are used for the main results of this paper.

The YCB-V dataset contains several symmetric objects: 13, 16, 19, 20, and 21. The bowl (Obj. 13) is rotationally symmetric while the other objects possess discrete symmetry transformations. For each symmetric object, some poses are not distinguishable from each other. Hence, keypoints of the symmetric objects cannot be learned by the network if the symmetries are ignored. A simple elimination of symmetric poses during training is implemented in this work. All symmetry-equivalent poses are mapped to the same pose which ensures same relative position of keypoints on the image plane.

3.2 Network

We extend a public implementation of the OpenPose framework [18]. The network architecture is adjusted to the YCB-V dataset as well as the used keypoints



Fig. 3: Objects of the YCB-V dataset with heatmaps of the automatically defined keypoints and their interconnections.

and PAFs. For an input image of size $H \times W$ and C object classes, the output shape is $\frac{H}{8} \times \frac{W}{8} \times (C \cdot 8 + 1)$ for the heatmaps and $\frac{H}{8} \times \frac{W}{8} \times (C \cdot 12 \cdot 2)$ for the PAFs. Heatmaps consist of 8 keypoint channels per object class and background, while PAFs consists of x and y channels for the 12 keypoint connections per class. Training labels are generated from the object model keypoints projected into the images of the video sequences using the annotated poses. A Gaussian blob is rendered at the keypoint position in the respective heatmap channel and the respective PAF channels represent the unit vector in the direction of the connection, within a fixed width along the connecting line, as in [2].

To enable the network to estimate the poses for multiple object classes, the layer width of the intermediate stages needs to be scaled accordingly to the output layers. A network for all 21 object classes would require more GPU memory than most graphic cards possess. Because of this, we train a separate model for each object class. We verify in the evaluation (Sec. 4), that the performance of these 1-object models is superior to models trained for multiple object classes.

3.3 Workflow

Fig. 1 illustrates the general workflow of our pipeline for 6D object pose estimation. The input image is processed by the network and heatmaps and PAFs are estimated. The local maxima of each heatmap are candidates for the respective keypoint. These keypoint candidates are grouped into object instances using the PAFs, as in the OpenPose framework [2]. This step is repeated for every object class. The PnP and RANSAC algorithms are used to calculate the 6D poses of the found object instances with four or more valid keypoints¹ using correspondences between detected 2D image keypoints and 3D model keypoints. In datasets with only one instance of an object per image, the best guess will be kept. In multi-instance datasets, each estimated object pose will be assigned to the closest ground truth object pose of this class for evaluation.

¹ The PnP algorithm requires at least four correspondences for a unique solution.

4 Evaluation

We evaluate our approach on the YCB-Video dataset [19], compare the results to other approaches from the literature, and perform ablation studies to understand the influences of different components of our method. All experiments run on a workstation PC with RTX 2080 GPU, i7-8700K CPU, and 32 GB of RAM. Pose estimation takes 50 ms in average per object and image, thereof 8 ms for pre-processing, 8 ms for inference, and 34 ms for post-processing and PnP. The network requires 1.24 GB of GPU memory.

4.1 Metrics

We employ two standard metrics for evaluation: average 3D distance of model points (ADD) [5] and 2D projection error [1]. Both metrics employ the meshes of the object models to calculate the pose error. The ADD metric is defined as:

$$\epsilon_{\text{ADD}} = \frac{1}{|P|} \sum_{\mathbf{p} \in P} \|(\mathbf{R}\mathbf{p} + \mathbf{t}) - (\tilde{\mathbf{R}}\mathbf{p} + \tilde{\mathbf{t}})\|, \quad (1)$$

with $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{t}}$ being the estimated rotation and translation, \mathbf{R} and \mathbf{t} defining the ground-truth pose and P the set of vertices of the object model mesh. For symmetric objects, the point-to-point correspondences can be ambiguous and the metric is adapted to compute the average distance using the closest point from the mesh [19]:

$$\epsilon_{\text{ADD-S}} = \frac{1}{|P|} \sum_{\mathbf{p}_1 \in P} \min_{\mathbf{p}_2 \in P} \|(\mathbf{R}\mathbf{p}_1 + \mathbf{t}) - (\tilde{\mathbf{R}}\mathbf{p}_2 + \tilde{\mathbf{t}})\|. \quad (2)$$

The 2D projection metric computes the average 2D pixel distances between corresponding points projected onto the image plane of the evaluated view:

$$\epsilon_{\text{2DProj}} = \frac{1}{|P|} \sum_{\mathbf{p} \in P} \|\text{proj}(\mathbf{R}\mathbf{p} + \mathbf{t}) - \text{proj}(\tilde{\mathbf{R}}\mathbf{p} + \tilde{\mathbf{t}})\|. \quad (3)$$

The evaluation scores are given in terms of the area under the accuracy-threshold curve (AuC). For this, the threshold for the respective distance metric is varied and the pose accuracy is computed for each threshold value. The maximum thresholds are set to 10 cm for ADD(-S) and 40 px for the 2D projection metric.

4.2 Results on the YCB-Video Dataset

In Tab. 1, we give detailed evaluation results of the AuC scores for all 21 objects of the YCB-V dataset and compare them to the results of PoseCNN [19]. The proposed approach outperforms PoseCNN in terms of ADD for most of the non-symmetric objects and on average over all objects. The improvement is most significant for the box-shaped Objects 2, 3, 7, and 8 as well as for Objects 11,

Table 1: Area under accuracy Curve (AuC) for pose estimation of YCB-V objects. * denotes symmetric objects. Best results are marked bold.

Object	Ours			PoseCNN [19]	
	ADD	ADD-S	2D Proj.	ADD	ADD-S
1	49.9	80.7	54.8	50.9	84.0
2	80.5	88.4	84.3	51.7	76.9
3	85.5	92.4	88.8	68.6	84.3
4	68.5	81.4	84.8	66.0	80.9
5	87.0	93.3	89.8	79.9	90.2
6	79.3	89.7	81.7	70.4	87.9
7	81.8	89.5	88.7	62.9	79.0
8	89.4	94.0	92.9	75.2	87.1
9	59.6	70.0	69.0	59.6	78.5
10	36.5	58.3	55.0	72.3	85.9
11	78.1	86.9	78.0	52.5	76.8
12	56.7	67.1	66.2	50.5	71.9
*13	12.2	23.5	4.1	6.5	69.7
14	54.0	76.9	75.2	57.7	78.0
15	82.8	91.0	88.2	55.1	72.8
*16	16.7	29.6	29.5	31.8	65.8
17	46.0	64.1	76.7	35.8	56.2
18	9.8	11.9	20.8	58.0	71.4
*19	20.0	47.4	8.9	25.0	49.9
*20	14.1	45.5	3.5	15.8	47.0
*21	12.1	29.7	2.3	40.4	87.8
average	59.0	72.7	65.0	53.7	75.9

15, and 17 which have a more complex shape (cf. Fig. 2). PoseCNN, on the other hand, achieves better results for the symmetric objects in terms of ADD-S. The proposed approach provides less accurate results for these objects, where several keypoint configurations can result in visually equivalent poses. This makes the keypoint estimation harder to learn and cannot be fully compensated by the symmetry handling during training (cf. Sec. 3). Also, keypoints are difficult to infer for objects with little prominent geometric features (e.g., edges or corners) such as Objects 10 and 18. The 2D projection metric scores per object are not reported by the authors of PoseCNN.

In Tab. 2, we further compare the overall results of our method with the recent Benchmark for 6D Object Pose Estimation (BOP) challenge 2020 [6,7]. The BOP challenge defines slightly different evaluation metrics: The MSSD and MSPD metrics are similar to the ADD and 2D projection metrics, but give the maximum value instead of the average error and deal with symmetries. The VSD metric is the percentage of pixels, which are visible in the estimated and ground truth pose and are close in the image space. The formal definitions are given in [7]. We achieve the third-best result. In comparison to CosyPose [10], which achieves the best result by a significant margin, we do not refine the initially estimated pose, which could further improve our result.

Table 2: Results on YCB-V of the five best candidates using only RGB images of the BOP challenge 2020 [7] in comparison with our work.

Name	AR	AR_{VSD}	AR_{MSSD}	AR_{MSPD}	training data
CosyPose [10]	0.821	0.772	0.842	0.850	pbr+real
EPOS [8]	0.696	0.626	0.677	0.783	pbr
Ours	0.575	0.506	0.567	0.654	pbr+real
CosyPose [10]	0.574	0.516	0.554	0.653	pbr
Leaping from 2D to 6D [13]	0.543	0.443	0.499	0.687	pbr+real
CDPNv2 [12]	0.532	0.396	0.570	0.631	pbr+real



Fig. 4: Qualitative results on the YCB-V dataset: (a) keypoints and connections, (b) transformed object models overlay on input image.

In Fig. 4, qualitative results on the YCB-V dataset are shown. 6D poses are estimated accurately despite occlusions and outlier keypoint detections.

4.3 Ablation Studies

Several systematic ablation studies have been conducted in this work to evaluate the influences of different components and parameters of the proposed method.

PAFs: We investigate the benefit of PAFs for the YCB-V dataset, where a maximum of one instance per object class is present in an image. For this, we infer the keypoints of an object instance as the global maximum of the respective heatmaps and do not use the PAFs computed by the CNN to assemble keypoints into object instances. This heatmaps-only approach is compared to the full approach using the PAF output. The results are presented in Tab. 3. Using the PAFs improves the pose estimation result for almost all objects as well as the average accuracy. Without PAFs, recovering from wrong or ambiguous heatmap maxima is not possible, leading to inaccurate results especially in the case of occlusions and truncation. An exception is Object 18, where the pose estimation is more accurate without using PAFs. This is probably due to the small size of the object leading to very short PAF vectors at the ends of the marker (cf. Fig. 2). These cannot be well detected by the model, leading to problems parsing the object skeleton. Fig. 5 shows the accuracy-threshold curves for each metric with and without using PAFs. The improvement using PAFs is most significant for small accuracy thresholds, demanding a precise estimation of the object pose. The two curves approach each other for higher thresholds.

Table 3: AuC when using PAFs to connect keypoints into instances vs. only using global heatmap maxima. * denotes symmetric objects. Best results marked bold.

Object	using PAFs			heatmaps only		
	ADD	ADD-S	2D Proj.	ADD	ADD-S	2D Proj.
1	49.9	80.7	54.8	48.0	80.1	50.2
2	80.5	88.4	84.3	73.1	83.3	72.7
3	85.5	92.4	88.8	79.2	89.4	76.6
4	68.5	81.4	84.8	59.5	75.6	75.0
5	87.0	93.3	89.8	80.4	90.8	79.3
6	79.3	89.7	81.7	68.3	84.5	72.4
7	81.8	89.5	88.7	74.0	84.3	81.2
8	89.4	94.0	92.9	81.8	90.3	81.0
9	59.6	70.0	69.0	54.3	66.3	60.4
10	36.5	58.3	55.0	35.2	55.3	53.0
11	78.1	86.9	78.0	72.9	84.3	67.4
12	56.7	67.1	66.2	51.7	64.3	57.8
*13	12.2	23.5	4.1	11.5	23.2	4.1
14	54.0	76.9	75.2	49.1	72.0	63.3
15	82.8	91.0	88.2	76.8	88.3	77.1
*16	16.7	29.6	29.5	15.7	27.4	23.2
17	46.0	64.1	76.7	37.3	56.2	65.1
18	9.8	11.9	20.8	36.1	43.1	65.1
*19	20.0	47.4	8.9	17.8	45.1	10.1
*20	14.1	45.5	3.5	11.8	43.0	4.7
*21	12.1	29.7	2.3	6.2	18.3	0.2
average	59.0	72.7	65.0	54.5	70.4	58.9

Selection of Keypoints on Object Models: In Tab. 4, we compare evaluation results using manually and automatically chosen object keypoints for an exemplary subset of the object classes. The estimated pose generally is more accurate using the manually defined keypoints. As discussed in Sec. 3, these keypoints are easier to find and can be more precisely localized by the CNN architecture, as they are placed on distinct spots of the object geometry and texture (cf. Figs. 2 and 3). The locations of the automatically chosen keypoints, on the other hand, are often weakly constrained along edges or on the object surface, thus being predicted less precisely.

Number of Object Classes per Model: We also investigate the influence of training the CNN model to detect keypoints and PAFs for objects of one or of several classes. The results of this comparison are shown in Tab. 5 exemplary for a subset of the object classes. The performance of models trained for one object class only is significantly better. The plot of the accuracy-threshold curves shown in Fig. 6 confirms these results. Also, rescaling the width of the CNN model, i.e., doubling the number of channels at each stage of the network, does not significantly improve the results of the 2-object model, as is shown in Tab. 6. Therefore, a separate model per object class has been used in this work.




Fig. 5: Accuracy curves for all YCB-V objects using PAFs vs. heatmaps only.

Table 4: AuC for manually and automatically picked keypoints.

Object	manually picked keypoints			automatically picked keypoints		
	ADD	ADD-S	2D Proj.	ADD	ADD-S	2D Proj.
1	49.9	80.7	54.8	47.7	78.5	56.7
2	80.5	88.4	84.3	81.0	89.8	87.1
3	85.5	92.4	88.8	82.4	90.8	89.6
4	68.5	81.4	84.8	68.4	80.8	84.3
5	87.0	93.3	89.8	83.7	92.0	92.0
6	79.3	89.7	81.7	77.3	88.9	85.2
7	81.8	89.5	88.7	75.9	85.0	85.2
8	89.4	94.0	92.9	81.0	89.6	92.4

Table 5: AuC for ADD(-S) and 2D-projection metrics comparing models trained to detect one object class versus models detecting two object classes.

Object	1-Object models			2-Object models		
	ADD	ADD-S	2D Proj.	ADD	ADD-S	2D Proj.
1	49.9	80.7	54.8	18.0	31.5	23.0
4	68.5	81.4	84.8	38.5	43.6	46.8
5	87.0	93.3	89.8	20.3	23.8	23.9
6	79.3	89.7	81.7	53.6	63.1	60.3
7	81.8	89.5	88.7	38.5	42.3	42.5
8	89.4	94.0	92.9	39.8	44.4	45.3
9	59.6	70.0	69.0	25.2	33.3	35.5
10	36.5	58.3	55.0	8.3	12.4	16.9
average	66.2	81.3	75.2	32.2	39.8	39.2

5 Conclusion

In this work, we introduce a pipeline for 6D object pose estimation adopting the OpenPose CNN architecture [2] to predict keypoints and Part Affinity Fields (PAFs) on objects. Keypoints are defined on prominent geometric features of the object contour (i.e., corners) and interconnected by PAFs to form a cuboid-like structure. They are predicted as local maxima of a pixel-wise heatmap and assembled into instances using the PAFs. This bottom-up approach permits to infer keypoints directly from the input image, without prior object detection or segmentation. Object poses are then calculated via the PnP-RANSAC algorithm using 2D-3D correspondences between detected and model keypoints.

Table 6: AuC using a 2-object model with and without rescaling layer width.

Object	doubled layer width			normal layer width		
	ADD	ADD-S	2D Proj.	ADD	ADD-S	2D Proj.
1	19.1	31.5	25.8	18.0	31.5	23.0
4	41.5	47.1	47.4	38.5	43.6	46.8
average	32.2	40.7	38.5	32.2	39.8	39.2




Fig 6: Accuracy curves for all YCB-V objects using models that detect one object class vs. models detecting two object classes.

The proposed approach is evaluated on the YCB-Video dataset, containing 21 typical household objects important for domestic service robot applications. Our method achieves accuracy comparable to recent state-of-the-art methods using RGB images only, without employing further pose refinement. The usage of PAFs is shown to be advantageous compared to directly using global maxima of the heatmaps as keypoint detections. Models trained for a single object class perform significantly better than models trained for multiple object classes.

Directions for future work include improved handling of symmetric objects and evaluating the method in scenarios where multiple instances of the same object class occur. Furthermore, a method for automatic keypoint and PAF selection without sacrificing accuracy compared to manual selection should be investigated, enabling to efficiently extend the method with novel object models.

Acknowledgments

This work was funded by grant BE 2556/18-2 of the German Research Foundation (DFG).

References

1. Brachmann, E., Michel, F., Krull, A., Yang, M.Y., Gumhold, S., Rother, C.: Uncertainty-driven 6D pose estimation of objects and scenes from a single RGB image. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3364–3372 (2016)
2. Cao, Z., Hidalgo Martinez, G., Simon, T., Wei, S., Sheikh, Y.A.: OpenPose: Real-time multi-person 2D pose estimation using part affinity fields. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **43**(1), 172–186 (2019)

3. Capellen, C., Schwarz, M., Behnke, S.: ConvPoseCNN: Dense convolutional 6D object pose estimation. In: International Conference on Computer Vision Theory and Applications (VISAPP). pp. 162–172 (2020)
4. 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* **24**(6), 381–395 (1981)
5. 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: Asian Conference on Computer Vision (ACCV). pp. 548–562 (2012)
6. Hodan, T., Michel, F., Brachmann, E., Kehl, W., GlentBuch, A., Kraft, D., Drost, B., Vidal, J., Ihrke, S., Zabulis, X., Sahin, C., Manhardt, F., Tombari, F., Kim, T.K., Matas, J., Rother, C.: BOP: Benchmark for 6D object pose estimation. In: European Conference on Computer Vision (ECCV). pp. 19–34 (2018)
7. Hodan, T., Sundermeyer, M., Drost, B., Labbe, Y., Brachmann, E., Michel, F., Rother, C., Matas, J.: BOP challenge 2020 on 6D object localization. arXiv:2009.07378 [cs] (2020)
8. Hodaň, T., Baráth, D., Matas, J.: EPOS: Estimating 6D pose of objects with symmetries. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11700–11709 (2020)
9. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In: IEEE International Conference on Computer Vision (ICCV). pp. 2938–2946 (2015)
10. Labbé, Y., Carpentier, J., Aubry, M., Sivic, J.: CosyPose: Consistent multi-view multi-object 6D pose estimation. In: European Conference on Computer Vision (ECCV). pp. 574–591 (2020)
11. Lepetit, V., Moreno-Noguer, F., Fua, P.: EPnP: An accurate $O(n)$ solution to the PnP problem. *International Journal of Computer Vision* **81**(2), 155 (2008)
12. Li, Z., Wang, G., Ji, X.: CDPN: Coordinates-based disentangled pose network for real-time RGB-based 6-DoF object pose estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7677–7686 (2019)
13. Liu, J., Zou, Z., Ye, X., Tan, X., Ding, E., Xu, F., Yu, X.: Leaping from 2D detection to efficient 6DoF object pose estimation. In: European Conference on Computer Vision (ECCV) Workshops. pp. 707–714 (2020)
14. Pavlakos, G., Zhou, X., Chan, A., Derpanis, K.G., Daniilidis, K.: 6-DoF object pose from semantic keypoints. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 2011–2018 (2017)
15. Peng, S., Liu, Y., Huang, Q., Zhou, X., Bao, H.: PVNet: Pixel-wise voting network for 6DoF pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4561–4570 (2019)
16. 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: IEEE International Conference on Computer Vision (ICCV) (2017)
17. Stückler, J., Schwarz, M., Behnke, S.: Mobile manipulation, tool use, and intuitive interaction for cognitive service robot cosero. *Frontiers in Robotics and AI* **3**, 58 (2016)
18. Wangpeng, A.: Multi person pose estimation by pytorch. https://github.com/tensorboy/pytorch_Realtime_Multi-Person_Pose_Estimation (Jul 2020)
19. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. In: Robotics: Science and Systems (RSS) (2018)