

“©2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

DSC-PoseNet: Learning 6DoF Object Pose Estimation via Dual-scale Consistency

Zongxin Yang^{1,2}, Xin Yu², Yi Yang^{2*}

¹ Baidu Research ² ReLER, University of Technology Sydney

zongxin.yang@student.uts.edu.au, {xin.yu,yi.yang}@uts.edu.au

Abstract

Compared to 2D object bounding-box labeling, it is very difficult for humans to annotate 3D object poses, especially when depth images of scenes are unavailable. This paper investigates whether we can estimate the object poses effectively when only RGB images and 2D object annotations are given. To this end, we present a two-step pose estimation framework to attain 6DoF object poses from 2D object bounding-boxes. In the first step, the framework learns to segment objects from real and synthetic data in a weakly-supervised fashion, and the segmentation masks will act as a prior for pose estimation. In the second step, we design a dual-scale pose estimation network, namely DSC-PoseNet, to predict object poses by employing a differential renderer. To be specific, our DSC-PoseNet firstly predicts object poses in the original image scale by comparing the segmentation masks and the rendered visible object masks. Then, we resize object regions to a fixed scale to estimate poses once again. In this fashion, we eliminate large scale variations and focus on rotation estimation, thus facilitating pose estimation. Moreover, we exploit the initial pose estimation to generate pseudo ground-truth to train our DSC-PoseNet in a self-supervised manner. The estimation results in these two scales are ensembled as our final pose estimation. Extensive experiments on widely-used benchmarks demonstrate that our method outperforms state-of-the-art models trained on synthetic data by a large margin and even is on par with several fully-supervised methods.

1. Introduction

The goal of object pose estimation is to estimate 6 degrees of freedom (DoF) of a given object, including 3D orientations and 3D translations, with respect to a camera. Considering an object may undergo various lighting changes and severe occlusions, estimating accurate poses from a single RGB image is quite challenging.

Thanks to the recent advance of deep neural networks, many deep learning based pose estimation algorithms (e.g., [47, 49]) have been proposed recently and achieved

promising performance. However, due to the notorious data-hunger nature of deep networks, state-of-the-art methods usually require a large number of real images with 3D pose annotations for training. Unlike 2D image labeling tasks, annotating 3D object poses is difficult, especially when only RGB images are provided to labelers. However, without sufficient accurately annotated 3D poses of RGB images, pose estimation methods may suffer severe performance degradation [41].

Different from real images of which poses are very difficult to obtain, pose labels in synthetic data are easily accessible. However, due to the domain gap between synthetic and real data, such as illumination variations and smoothed 3D reconstructed models, state-of-the-art fully-supervised pose estimation approaches would suffer drastic performance degradation when they are trained on synthesized images [41]. Therefore, we attempt to use minimal annotation efforts, i.e., 2D bounding boxes, to reduce the domain gap between real and synthetic images for pose estimation. We employ bounding-boxes to ensure objects of interest are correctly identified in images rather than highly relying on synthesized data's photorealism.

In this paper, we present a two-step object pose estimation framework to predict object poses by only leveraging easily obtained 2D annotations. In the first step, we employ a weakly supervised segmentation method to distinguish object pixels from background ones. Specifically, we firstly initialize a segmentation network by training it on synthetic images. Motivated by [4, 48], we use the segmentation network to predict pseudo masks for the unlabeled real data and re-train the network on both synthetic and pseudo-labeled data. Since some pixels in the background might be recognized as foreground ones, we fully exploit our 2D bounding-boxes to remove outliers and thus facilitate learning segmentation. The segmentation results, in return, provide a strong prior for pose estimation.

In the second step, we present a dual-scale pose estimation network, dubbed DSC-PoseNet. Our DSC-PoseNet is firstly initialized by training on synthetic images. Then, we employ DSC-PoseNet to estimate object poses in the original image scale. DSC-PoseNet regresses an offset of each pixel to an object keypoint and employs an attention map to obtain the keypoint locations. We adopt differen-

*This work was done when Zongxin Yang interned at Baidu Research. Yi Yang is the corresponding author.

tiable EPnP [20] and a renderer [32] to achieve rendered object masks, and the intersection over union (IoU) between the segmentation results and rendered visible masks is employed to train our DSC-PoseNet on real images. Once real image poses are attained, we resize object regions to a fixed scale and then project the keypoints computed by our estimated poses to the local regions as self-supervised signals to train DSC-PoseNet once again. By doing so, we improve the feature extraction ability of DSC-PoseNet and enforce the pose estimation consistency across different scales.

As DSC-PoseNet outputs pose estimation results on two scales, we ensemble these two-scale results to improve the estimation robustness. Extensive experiments on three popular widely-used datasets demonstrate that our method achieves superior performance compared to state-of-the-art RGB image based methods without real pose annotations.

Overall, our contributions are summarized as follows:

- We propose a weakly- and self-supervised learning based pose estimation framework to estimate object poses from single RGB images with easily obtained 2D bounding-box annotations.
- We present a self-supervised dual-scale pose estimation network, named DSC-PoseNet. The DSC-PoseNet significantly alleviates the domain gap between synthetic and real data by constructing cross-scale self-supervision with a differentiable renderer.
- To the best of our knowledge, our work is the first attempt to estimate 6DoF object poses from an RGB image without using 3D pose annotations and depth images during both training and testing phases. Contrast results show that DSC-PoseNet outperforms RGB based competitors trained on synthetic data.

2. Related Work

Since our method predicts object poses from a single RGB image, we briefly review fully-supervised and self-supervised RGB image based pose estimation approaches.

Conventional methods: Traditional object pose estimation methods use local feature/keypoint matching approaches [24, 45, 39] to associate 2D images with 3D textured models. Then, the object poses are obtained by solving a Perspective-n-Point (PnP) problem [20]. However, local feature based methods can only address textured objects. Image template [10, 7, 33, 50], edge [9, 23] based methods are developed for pose estimation to tackle textureless objects. When an object undergoes occlusions or drastic illumination changes, those methods might fail to estimate object poses accurately [29].

Fully-supervised deep model based methods: Deep learning based methods have demonstrated promising pose estimation performance [29, 43, 47, 46, 19, 22]. Some pioneering approaches, such as Render for CNN [35], View-

points & Keypoints [40], and SSD6D [17], discretize 3D pose space and classify object poses with discrete labels.

Instead of treating pose estimation as a classification task, recent approaches directly regress 3D bounding-boxes [30, 38], local features [43, 29] or coordinate maps [47, 42, 18, 51] of objects, and then predict object poses via PnP. For example, YOLO6D [38] regresses 3D object bounding-boxes from input images. BB8 [30] first generates 2D object segmentation masks and then estimates 3D bounding-boxes from the 2D masks. CPDN [49], DPOD [47] and Pix2Pose [18] output the 2D UV coordinates or 3D coordinates of 3D object models from images, while PoseCNN [43] and PVNet [29] employ Hough voting to localize object keypoints from estimated vector fields. EPOS [12] uses symmetric segments of an objects to mitigate the symmetry ambiguity in pose estimation. Once 2D-3D correspondences are attained, object poses are achieved by solving a PnP problem.

Self-supervised deep model based methods: Prior arts learn to estimate object poses with ground-truth 3D pose annotations. However, when 3D object poses in real images are not available, those methods would suffer severe performance degradation [41]. Very recently, some self-supervised methods have been proposed to address the lack of ground-truth pose annotations in real images. Self6D [41] and DTPE [31] first localize objects and obtain object poses by minimizing the chamfer distances between rendered depth images and real depth images in a self-supervision manner. Although Self6D and DTPE do not require depth images during inference, they need depth images in training. AAE [37] trains a pose auto-encoder without depth images, but depth information is critical in refining pose results. Without depth information, these methods suffer dramatic performance degradation. NOL [28] generates novel views of a given object from a few cluttered real-world images in which object poses are needed. The generated views are self-supervised by the input images. Therefore, it is still an open question of how to estimate object poses only from RGB images for both training and testing.

Pose refinement: Depth images are usually used to refine initially estimated poses. For example, PoseCNN employs Iterative Closest Point (ICP) [1] to refine the estimated poses on depth images. DeepIM [21] and DPOD [47] exploit optical flow to modify initial pose estimation by minimizing the differences between the object appearance and the 2D projection of the 3D model. Moreover, Hu *et al.* [15] develop a standalone pose refinement network by removing outliers estimated by other methods and then refining initial pose estimation. Cosypose [19] leverages multi-view RGB images to refine each estimated pose in a scene via scene reconstruction. Considering the vast illumination changes, directly aligning rendered RGB images to input ones for pose

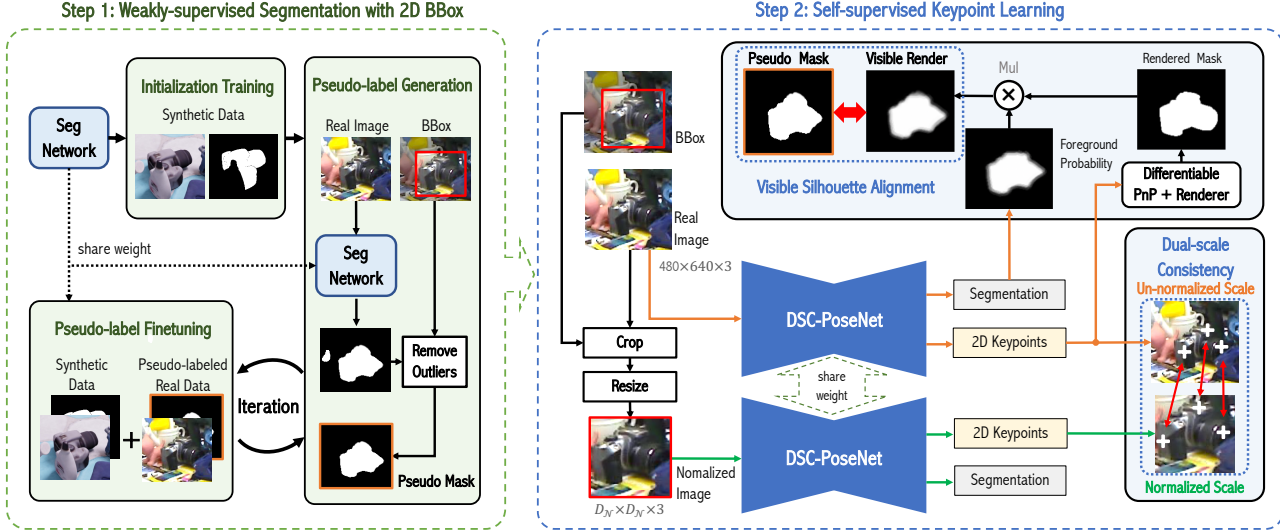


Figure 1. An **overview** of the training pipeline of DSC-PoseNet. In the first step, we utilize a weakly-supervised segmentation method to generate pseudo masks for real images with only bounding-box (BBox) annotations. In the second step, self-supervised keypoint learning is then developed for training DSC-PoseNet by constructing dual-scale self-supervision signals with a differentiable renderer.

refinement might be unreliable. Hence, single RGB image based pose refinement methods require ground-truth poses to bridge the domain gap between rendered images and real ones. In other words, when the ground-truth pose annotations of real images are not available, single RGB image based methods fail to refine initial pose estimation [19].

3. Proposed Method

Overview. In this section, we introduce our two-step pose estimation framework with 2D bounding-box (BBox) annotations. As illustrated in Fig. 1, the training paradigm includes (i) weakly-supervised segmentation learning and (ii) self-supervised object keypoint learning.

3.1. Weakly-supervised Segmentation with BBox

We aim to obtain object segmentation results from coarse 2D object BBox annotations for real images since segmentation can provide detailed object silhouettes from which object poses can be roughly determined. Thus, we design a weakly-supervised segmentation method to generate pixel-level object segmentation. Here, we do not assume the given 2D BBoxes tightly cover object regions, and this further relaxes the labeling requirements.

Our proposed weakly-supervised segmentation method takes an iterative learning strategy, similar to [4]. As shown in the left of Fig. 1, we first initialize the segmentation network by training it on synthetic data in a fully-supervised way. After initialization, we use the network to generate pseudo-labels for all the BBox-annotated real images and then fine-tune the network on pseudo-labeled real images. We iterate pseudo-label generation and fine-tuning until our segmentation network converges.

Pseudo-label generation and fine-tuning: After initializing the network on synthetic data or fine-tuning on pseudo-labeled real data, we generate (or update) the pseudo segmentation labels for all the real images along with a test-time augmentation, *i.e.*, multi-scale inputs and left-right flips. This augmentation is commonly used in segmentation tasks [5]. Since some pixels in the background might be recognized as foregrounds, we fully exploit 2D BBox annotations to remove those outliers and thus facilitate the network fine-tuning. We further set pixels with low confidence to be uncertain pixels and avoid calculating the segmentation loss for them.

Let \mathbf{x} and \mathbf{y} denote an image and a segmentation map respectively, and $y_m \in \{0, 1\}$ is the pixel class (background or foreground) at position $m \in \{1, \dots, M\}$, the uncertain pixel set is defined as,

$$U_{\mathbf{x}} = \{m | 1 - \sigma < P(y_m = 1 | \mathbf{x}; \theta_{seg}) < \sigma\}, \quad (1)$$

where θ_{seg} indicates segmentation network parameters, P is a confidence score, and σ is a reliable segmentation confidence value. We empirically set σ to 0.7.

When fine-tuning on pseudo-labeled real images \mathbf{x}_{real} , the segmentation loss function is,

$$\mathcal{L}_{seg} = \sum_{m \notin U_{\mathbf{x}}} \log P(y_m = \hat{y}_m | \mathbf{x}_{real}, \theta_{seg}), \quad (2)$$

where \hat{y}_m denotes the pseudo segmentation label. To prevent the network from fitting to error-prone uncertain regions ($U_{\mathbf{x}}$), we avoid calculating the loss on them.

The pseudo-label generation and fine-tuning will iterate by T times until the segmentation network converges. Em-

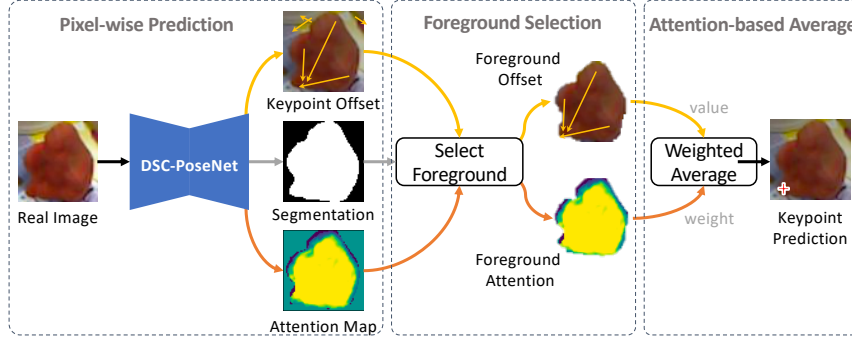


Figure 2. Differentiable 2D Keypoint Coordinate Prediction. We predict the mask of object, the keypoint offset for each pixel and the attention map. The final keypoint prediction is calculated by weighted-averaging the per-pixel predictions in the foreground.

pirically, we found that $T = 5$ makes a good balance between training efficiency and performance. The pseudo labels generated after the final iteration will serve as the segmentation labels in the following self-supervised pose estimation step. Note that some methods, such as Self6D [41], do not require object BBoxes in learning segmentation, but they request the photorealism of synthesized images. In contrast, we use coarsely labeled BBoxes to bridge the domain gap between real and synthetic data rather than requiring synthetic data to be similar to real ones.

3.2. Self-supervised DSC-PoseNet

Inspired by recent keypoint based estimation methods [29, 34], we use the intermediate object representation, *i.e.*, keypoints, for pose estimation. The final 6DoF poses are obtained by establishing 2D-3D correspondences and solving a PnP problem. However, the ground-truth positions of keypoints are not available in real images because we do not have 3D pose annotations. Therefore, in the second step of our framework, we design a novel self-supervised dual-scale consistency pose estimation network (DSC-PoseNet) to predict the locations of keypoints and then use the estimated keypoints to predict object poses in real images, as shown at the bottom of Fig 1.

In enforcing dual-scale keypoint consistency, we introduce a normalized scale by normalizing the object scales in real data. In other words, our keypoint network predicts keypoints on both un-normalized and normalized scales. Thus, we can improve the scale robustness by constraining the keypoint consistency between different scales.

To train our pose estimation network, we propose two self-supervised objective functions: visible silhouette alignment (VSA) and dual-scale keypoint consistency (DKC). VSA measures alignments between the predicted pose’s silhouette and the visible silhouette in the image, while DKC evaluates the keypoint estimation consistency in normalized image scale and un-normalized scale.

Differentiable keypoint regression: To regress keypoints in a differentiable way, DSNT [27] averages coordinates

based on the weights of a spatial heatmap. However, when a keypoint is occluded, DSNT will fail to predict a valid heatmap. PVNet [29] predicts a vector field for each keypoint and employs voting to determine keypoint locations. Since the voting procedure is not differentiable, using vector field representation is not applicable to designing a differentiable framework.

To solve all the above issues, we propose to regress 2D keypoint coordinates from all the object pixels (see Fig. 2). For each 2D keypoint $k_n (n = 1, \dots, N)$ and each object pixel coordinate p_m at position m , DSC-PoseNet will generate an attention weight a_{nm} and predict a keypoint offset Δk_{nm} . Then, the predicted keypoint \tilde{k}_n is defined as,

$$\tilde{k}_n = \frac{\sum_{m \in O} e^{a_{nm}} \tilde{k}_{nm}}{\sum_{m \in O} e^{a_{nm}}}, \quad \tilde{k}_{nm} = p_m + \Delta k_{nm}, \quad (3)$$

where O is the foreground segmentation predicted by DSC-PoseNet and \tilde{k}_{nm} denotes the keypoint prediction from only pixel p_m . Compared to the voting process and coordinate heatmap based average, our attention-based average is fully differentiable and can predict occluded keypoints.

For synthetic data \mathbf{x}_{syn} , we can obtain accurate keypoint k_n . Thus, we use the smooth ℓ_1 [6] to learn k_n directly. The keypoint loss function for synthetic data is,

$$\mathcal{L}_{key}^{syn} = \sum_{n=1}^N \ell_1 \left(\frac{\tilde{k}_n(\mathbf{x}_{syn}, \boldsymbol{\theta}_{pose}) - k_n(\mathbf{x}_{syn})}{\sigma S(\mathbf{x}_{syn})} \right), \quad (4)$$

where $\boldsymbol{\theta}_{pose}$ is the parameter of DSC-PoseNet, σ is a scale factor, and S is the object scale defined by the length of the longest side of the BBox. We observe that the larger an object is in terms of pixels, the greater the variance of the keypoint prediction will be (*i.e.*, pixel deviations). Thus, we employ S to normalize the keypoint estimation errors. We found the normalization with the object scale S is crucial in maintaining a stable training process.

Apart from the keypoint loss, we introduce an auxiliary offset loss,

$$\mathcal{L}_{off}^{syn} = \sum_{n=1}^N \sum_{m \in O} \ell_1 \left(\frac{\tilde{k}_{nm}(\mathbf{x}_{syn}, \boldsymbol{\theta}_{pose}) - k_n(\mathbf{x}_{syn})}{\sigma S(\mathbf{x}_{syn})} \right). \quad (5)$$

This loss enforces DSC-PoseNet to regress accurate keypoints from every foreground pixel. Before learning the keypoint regression, we use the offset loss to warm up DSC-PoseNet. This is helpful for initializing the keypoint regression and thus stabilizing the learning of keypoint attention.

Dual-scale keypoint consistency. For real data \mathbf{x}_{real} , we only have pseudo segmentation labels $\tilde{\mathbf{y}}$. Therefore, we need a self-supervised way to learn keypoints from the real data without pose annotations. Intuitively, the keypoint locations of the same object predicted by our DSC-PoseNet should be consistent across different scales. Thus, we introduce a dual-scale keypoint prediction consistency constraint for pose estimation, as shown in the bottom right of Fig. 1.

Given an image \mathbf{x} , we can transform it into a normalized scale $\mathcal{N}(\mathbf{x})$ by cropping out the object and resizing it to a fixed size, $D_N \times D_N$. The 2D keypoints should follow the same transformation,

$$k_n(\mathbf{x}) = \mathcal{N}^{-1}(k_n(\mathcal{N}(\mathbf{x}))). \quad (6)$$

Such a relation can serve as a dual-scale consistency regularization for our keypoint prediction.

Apart from the dual-scale consistency, we also utilize an augmentation consistency. Data augmentations (such as random cropping and resizing) are commonly used in pose estimation task for improving the network robustness and relieving over-fitting. The 2D keypoints should be consistent after the augmentation \mathcal{A} , *i.e.*,

$$k_n(\mathbf{x}) = \mathcal{A}^{-1}(k_n(\mathcal{A}(\mathbf{x}))). \quad (7)$$

To further utilize the 3D structure information while reducing prediction errors, we employ BPnP [3] to compute the object pose from the predicted 2D keypoints, and then re-project the 3D keypoints on a CAD model back to 2D image space using the computed pose. To be specific, given a set of keypoint predictions, $\tilde{\mathbf{k}} = \{\tilde{k}_1, \dots, \tilde{k}_N\}$, corresponding 3D keypoint set, \mathbf{k}^{3D} , on CAD model, and camera intrinsic matrix, \mathcal{K} , the re-projected 2D keypoints $\tilde{\mathbf{k}}^P$ are,

$$\tilde{\mathbf{k}}^P = \mathcal{P}(\tilde{\mathbf{k}}) = \tilde{R}\mathbf{k}^{3D} + \tilde{t}, \quad (8)$$

$$(\tilde{R}, \tilde{t}) = BPnP(\tilde{\mathbf{k}}, \mathbf{k}^{3D}, \mathcal{K}), \quad (9)$$

where \tilde{R} and \tilde{t} are the predicted 3D rotation and translation. Combining Eqn. (6), (7), and (8), the loss function of our

dual-scale consistency for real data is,

$$\begin{aligned} \mathcal{L}_{dual}^{real} = & \sum_{n=1}^N \left(\ell_1 \left(\frac{\mathcal{A}^{-1}(\tilde{k}_n(\mathcal{A}(\mathbf{x}_{real}), \boldsymbol{\theta}_{pose})) - \bar{k}_n}{\sigma S(\mathbf{x}_{real})} \right) \right. \\ & \left. + \ell_1 \left(\frac{\mathcal{A}^{-1} \cdot \mathcal{N}^{-1}(\tilde{k}_n(\mathcal{N} \cdot \mathcal{A}(\mathbf{x}_{real}), \boldsymbol{\theta}_{pose})) - \bar{k}_n}{\sigma S(\mathbf{x}_{real})} \right) \right), \end{aligned} \quad (10)$$

where

$$\bar{k}_n = \mathcal{P} \left(\frac{\tilde{k}_n(\mathbf{x}_{real}, \boldsymbol{\theta}_{pose}) + \mathcal{N}^{-1}(\tilde{k}_n(\mathcal{N}(\mathbf{x}_{real}), \boldsymbol{\theta}_{pose}))}{2} \right) \quad (11)$$

is the 2D re-projection of the average keypoint prediction over normalized and un-normalized (original) scales of \mathbf{x}_{real} . Similar to Eqn. (4), we use the object scale S to normalize keypoint errors.

Visible silhouette alignment supervision: To estimate object poses from real images, we further optimize DSC-PoseNet by aligning the predicted object pose silhouette with its counterpart in real images. Given a predicted 3D rotation \tilde{R} , a translation \tilde{t} , the 3D CAD model \mathcal{M} , and the camera intrinsic matrix \mathcal{K} , we use PyTorch3D [32] to render the predicted object mask $\tilde{\mathbf{y}}^R = \mathcal{R}(\tilde{R}, \tilde{t}, \mathcal{M}, \mathcal{K}) = \{\tilde{y}_1^R, \dots, \tilde{y}_M^R\}$.

We directly use the pseudo labels $\hat{\mathbf{y}}$ from our weakly-supervised segmentation step as the object silhouette for real images. However, the segmentation result only provides visible regions of an object, while the rendered mask covers the entire object structure. In other words, the alignment will fail when the object is occluded.

To solve the above problem, we use the segmentation branch of DSC-PoseNet to select visible regions of the rendered mask. Let $\tilde{y}_m = P(y_m = 1 | \mathbf{x}; \boldsymbol{\theta}_{pose})$ denote the foreground probability predicted by the segmentation branch, the visible rendered mask is,

$$V(\tilde{\mathbf{y}}^R) = \{\tilde{y}_m^R \tilde{y}_m | m \in 1, \dots, M\}, \quad (12)$$

which is a combination of foreground probability and rendered mask.

After achieving the visible rendered mask, we use the IoU-based Dice loss [36] to align it with the pseudo segmentation mask, $\hat{\mathbf{y}} = \{\hat{y}_1, \dots, \hat{y}_M\}$, as follows,

$$\mathcal{L}_{align}^{real} = Dice(\hat{\mathbf{y}}(\mathbf{x}_{real}, \boldsymbol{\theta}_{seg}), V(\tilde{\mathbf{y}}^R(\mathbf{x}_{real}, \boldsymbol{\theta}_{pose}))), \quad (13)$$

$$Dice(\hat{\mathbf{y}}, V(\tilde{\mathbf{y}}^R)) = 1 - \frac{\sum_{m=1}^M 2\hat{y}_m \tilde{y}_m^R \tilde{y}_m + \epsilon}{\sum_{m=1}^M (\hat{y}_m + \tilde{y}_m^R \tilde{y}_m) + \epsilon}, \quad (14)$$

where ϵ is a small number for avoiding division by 0.

The overall self-supervision loss used to train DSC-PoseNet is a combination of all the above losses,

$$\mathcal{L}_{self} = \mathcal{L}_{key}^{syn} + \mathcal{L}_{off}^{syn} + \lambda_1 \mathcal{L}_{dual}^{real} + \lambda_2 \mathcal{L}_{align}^{real}, \quad (15)$$

| Train data | w/o Real Pose Labels | | | | | | with Real Pose Labels | | | | |
|------------|----------------------|----------|-----------|-------------|-------------|-------------|-----------------------|-------------|-----------------------|-------------|-------------|
| Method | RGB-based | | RGB-based | | | | RGB-based | | with Real Pose Labels | | |
| | AAE [37] | MHP [25] | Self6D-LB | DPOD [47] | Ours | Ours+ | DTPE [31] | Self6D [41] | Tekin [38] | DPOD [47] | CDPN [49] |
| Ape | 4.2 | 11.9 | 14.8 | 35.1 | 35.9 | 31.2 | 19.8 | 38.9 | 21.6 | 53.3 | 64.4 |
| Bvise | 22.9 | 66.2 | 68.9 | 59.4 | 83.1 | 83.0 | 69.0 | 75.2 | 81.8 | 95.2 | 97.8 |
| Cam | 32.9 | 22.4 | 17.9 | 15.5 | 51.5 | 49.6 | 37.6 | 36.9 | 36.6 | 90.0 | 91.7 |
| Can | 37.0 | 59.8 | 50.4 | 48.8 | 61.0 | 56.5 | 42.3 | 65.6 | 68.8 | 94.1 | 95.9 |
| Cat | 18.7 | 26.9 | 33.7 | 28.1 | 45.0 | 57.9 | 35.4 | 57.9 | 41.8 | 60.4 | 83.8 |
| Drill | 24.8 | 44.6 | 47.4 | 59.3 | 68.0 | 73.7 | 54.7 | 67.0 | 63.5 | 97.4 | 96.2 |
| Duck | 5.9 | 8.3 | 18.3 | 25.6 | 27.6 | 31.3 | 29.4 | 19.6 | 27.2 | 66.0 | 66.8 |
| Eggbox | 81.0 | 55.7 | 64.8 | 51.2 | 89.2 | 96.0 | 85.2 | 99.0 | 69.6 | 99.6 | 99.7 |
| Glue | 46.2 | 54.6 | 59.9 | 34.6 | 52.5 | 63.4 | 77.8 | 94.1 | 80.0 | 93.8 | 99.6 |
| HoleP | 18.2 | 15.5 | 5.2 | 17.7 | 26.4 | 38.8 | 36.0 | 16.2 | 42.6 | 64.9 | 85.8 |
| Iron | 35.1 | 60.8 | 68.0 | 84.7 | 56.3 | 61.9 | 63.1 | 77.9 | 75.0 | 99.8 | 97.9 |
| Lamp | 61.2 | - | 35.3 | 45.0 | 68.7 | 64.7 | 75.1 | 68.2 | 71.1 | 88.1 | 97.9 |
| Phone | 36.3 | 34.4 | 36.5 | 20.9 | 46.3 | 54.4 | 44.8 | 58.9 | 47.7 | 71.4 | 90.8 |
| Mean | 32.6 | 38.8 | 40.1 | 40.5 | 54.7 | 58.6 | 51.6 | 59.0 | 56.0 | 82.6 | 89.9 |

Table 1. Comparisons with the state-of-the-art on LINEMOD dataset. We present the results for the Average Recall (%) of ADD(-S). The results except ours and DTPE are copied from [41]. Ours⁺: our results trained with additional 10K synthetic data rendered by OpenGL.

where λ_1 and λ_2 are the balance factors for $\mathcal{L}_{dual}^{real}$ and $\mathcal{L}_{align}^{real}$, respectively. \mathcal{L}_{key}^{syn} , \mathcal{L}_{off}^{syn} and $\mathcal{L}_{align}^{real}$ are employed at both normalized and un-normalized scales.

4. Experiments

We conduct experiments on three popular widely-used datasets designed for evaluating 6DOF pose estimation algorithms. We also present a detailed analysis of each key component in our two-step framework. For better understanding, we provide our network only trained on synthetic data as well as real data with ground-truth poses as our lower-bound and upper-bound baselines.

4.1. Datasets

Synthetic PBR training dataset [14] Our segmentation and pose estimation networks are both trained on synthetic and real images. We employ 50K PBR images rendered with LINEMOD models [11] as our synthetic data. Note that PBR images exhibit an obvious domain gap with LINEMOD dataset [11].

LINEMOD [11] consists of 13 objects and each object has around 1,200 images. The CAD models of those objects are also provided. It is regarded as a standard 6DOF pose estimation benchmark. For each object category, 15% images are used for training and the rest for testing [30]. LINEMOD exhibits many challenges, such as textureless objects, lighting variations, and cluttered scenes.

Occluded LINEMOD [2] provides additional annotations on the bench vise subset in LINEMOD dataset. All the LINEMOD objects are annotated in each image. As scenes are cluttered, many objects undergo severe occlusion.

HomebrewedDB [16] Following Self-6D [41], we also employ HomebrewedDB dataset to test the generalization of

our method. Specifically, we adopt the second sequence of HomebrewedDB for evaluation as it contains three objects from LINEMOD (*i.e.*, bench vise, driller, and phone), and those objects are captured in a new environment.

4.2. Evaluation Metrics

We employ a commonly used metric, *i.e.*, ADD score [11], to evaluate 6DoF pose estimation accuracy. ADD score measures the average 3D distance between 3D model point clouds transformed by a predicted pose and its corresponding ground-truth. When the distance is less than 10% of the diameter of a 3D model, the estimated pose is considered correct. For symmetric objects, ADD-S [43] is used, where the closest 3D point distance is employed.

4.3. Implementation Details

In our DSC-PoseNet, we choose 8 keypoints on each object model by the farthest point sampling algorithm. We also apply data augmentation, including random cropping, resizing, rotation, and color jittering, to avoid overfitting. We use DeepLabV3+ [5] as the segmentation network in the weakly-supervised step. The backbone of DeepLabV3+ is GCT-ResNet-50 [44]. To relax the requirements of accurate labeling, we apply perturbations to ground-truth BBoxes by randomly expanding them up to 15% of their sizes.

For our DSC-PoseNet, we use ResNet-50 [8] as the backbone following Self-6D. We set the batch size to 16 and employ Adam optimizer with an initial learning rate of $1e^{-3}$, which is decayed to $1e^{-5}$ by a factor of 0.85 every 5 epochs. The maximum training epoch is set to 120. When training on synthetic images, we only use ground-truth keypoints as supervision. When real images are given, we employ the Eqn. (10) and (13) to optimize our network. In particular,

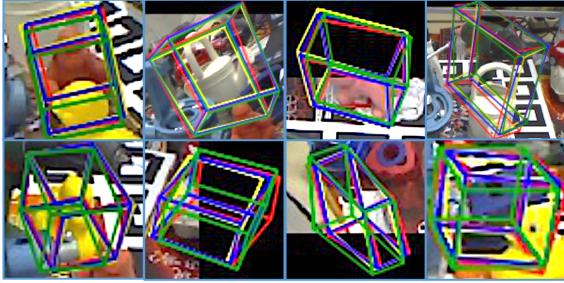


Figure 3. Qualitative results on **Occluded LINEMOD** dataset. *Green*: the ground truth pose. *Red*: un-normalized scale prediction. *Yellow*: normalized scale prediction. *Blue*: ensemble prediction by averaging the keypoints predicted at both the scales.

the resolution of the un-normalized scale is 480×640 pixels, and the normalized one is 64×64 pixels. During inference, only DSC-PoseNet is used, and the cropping BBoxes for the normalized scale are generated by bounding the object segmentation prediction from the un-normalized scale. Besides, we ensemble the keypoint predictions from both the scales by averaging their coordinates. Our method runs at 7fps on 480×640 images on 1 NVIDIA V100 GPU.

4.4. Comparison with SOTA

In this section, we mainly compare our method with state-of-the-art RGB image based methods trained on synthetic datasets, *i.e.*, AAE [37], MHP [25], DPOD [47] and PVNet [29], as well as RGBD based methods, *i.e.*, DTPE [31] and Self6D [41]. Note that, DTPE and Self6D both use RGBD images during training and only use RGB images in testing. Thus, we categorize Self6D and DTPE as RGBD based methods. Moreover, Self6D also provides an RGB only version, namely Self6D-LB, and we also include it for comparisons.

Comparisons on LINEMOD: As expected, the performance of state-of-the-art methods decreases dramatically on LINEMOD dataset when they are trained on synthetic data, as indicated in Table 1. This also implies that the domain gap is the main issue rather than different pose representations. When depth images are available, RGBD based methods can localize objects and estimate their poses via ICP [1]. In this way, they significantly ease the difficulty of pose estimation. As seen in Table 1, RGBD based methods indeed outperform RGB based methods. In contrast, our method even outperforms RGBD based method DTPE and is slightly inferior to Self6D. In particular, our method outperforms Self6D on some objects, such as bench vise, cam, and duck, although we did not use depth information during training. We also provide our results with additional synthetic data rendered by OpenGL following Self6D. Under this fairer comparisons, our performance is on par with the results of RGBD based Self6D, as seen in Table 1.

| Method | RGB-based | | | | | RGBD |
|--------|-----------|-------------|-------------|-------------|-------------|-------------|
| | DPOD [47] | Self6D-LB | CDPN [49] | Ours | Ours+ | Self6D [41] |
| Ape | 2.3 | 7.4 | 20.0 | 13.9 | 9.1 | 13.7 |
| Can | 4.0 | 14.1 | 15.1 | 15.1 | 21.1 | 43.2 |
| Cat | 1.2 | 7.6 | 16.4 | 19.4 | 26.0 | 18.7 |
| Drill | 7.2 | 18.0 | 22.2 | 40.5 | 33.5 | 14.4 |
| Duck | 10.5 | 12.2 | 5.0 | 6.9 | 12.2 | 32.5 |
| Eggbox | 4.4 | 18.3 | 36.1 | 38.9 | 39.4 | 57.8 |
| Glue | 12.9 | 31.4 | 27.9 | 24.0 | 37.0 | 54.3 |
| HoleP | 7.5 | 11.5 | 24.0 | 16.3 | 20.4 | 22.0 |
| Mean | 6.3 | 11.5 | 20.8 | 21.9 | 24.8 | 32.1 |

Table 2. Comparisons with the state-of-the-art on **Occluded LINEMOD** dataset.

| Method | RGB-based | | | | RGBD |
|--------|------------|-----------|---------------|-------------|-------------|
| | Self6D-Syn | DPOD [47] | SSD+Ref. [17] | Ours | Self6D [41] |
| Bvise | 37.7 | 52.9 | 82.0 | 72.9 | 72.1 |
| Drill | 19.2 | 37.8 | 22.9 | 40.6 | 65.1 |
| Phone | 20.9 | 7.3 | 24.9 | 18.5 | 41.8 |
| Mean | 25.9 | 32.7 | 43.3 | 44.0 | 59.7 |

Table 3. Comparisons with state-of-the-art on **HomebrewedDB** dataset. Note that we did not finetune our model.

Comparisons on Occluded LINEMOD: Following previous methods (*i.e.*, DPOD, CDPN and Self6D), we directly test the models trained on LINEMOD dataset on the BOP [13] split of Occluded LINEMOD dataset. The comparisons of the state-of-the-art approaches are shown in Table 2. Our method achieves the best overall performance compared to RGB based methods. Note that when Self6D does not use depth images in training, named Self6D-LB, our method even outperforms Self6D-LB. Fig. 3 demonstrates our qualitative results on Occlusion LINEMOD.

Comparisons on HomebrewedDB: As aforementioned, we employ HomebrewedDB dataset to test the generalization ability of our method. We also compare state-of-the-art method, *i.e.*, DPOD, Self6D, and SSD [17] after refinement [26] (SSD+Ref.). Note that Self6D re-trains their models on HomebrewedDB and also provides a version only trained on synthetic RGBD data, named Self6D-Syn. As indicated in Table 3, our method still outperforms the methods trained on synthetic data. This also demonstrates that by employing a few 2D annotations on real images, our method significantly minimizes domain gaps and improves the performance of the pose estimation.

Since state-of-the-art RGB based methods do not have mechanisms to use 2D bounding-boxes to train their networks, they may fail to use 2D annotations even though real images are provided. On the contrary, our method is able to leverage 3D and 2D annotations. Therefore, our framework is more generic and more applicable in practice.

| | Ape | Bvise | Cam | Can | Cat | Drill | Duck | Eggbox | Glue | Holep | Iron | Lamp | Phone | Mean |
|---|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| <i>Self-supervised Keypoint Learning (ADD metric)</i> | | | | | | | | | | | | | | |
| Ours [†] (with Real Pose Labels) | 59.2 | 98.1 | 88.0 | 92.1 | 79.4 | 94.5 | 51.7 | 98.5 | 93.9 | 78.4 | 96.2 | 96.3 | 90.0 | 85.9 |
| Ours | 35.9 | 83.1 | 51.5 | 61.0 | 45.0 | 68.0 | 27.6 | 89.2 | 52.5 | 26.4 | 56.3 | 68.7 | 46.3 | 54.7 |
| w/o Real RGB images | 23.4 | 75.6 | 11.7 | 40.1 | 26.7 | 53.8 | 14.0 | 73.6 | 26.7 | 19.5 | 56.2 | 39.4 | 20.0 | 37.0 |
| w/o Normalized Scale | 24.7 | 47.0 | 11.9 | 30.4 | 35.7 | 40.1 | 19.5 | 53.7 | 28.7 | 6.4 | 25.0 | 38.8 | 22.7 | 29.6 |
| w/o Keypoint Attention | 10.2 | 60.0 | 26.4 | 33.6 | 29.3 | 39.6 | 8.1 | 43.4 | 40.7 | 8.1 | 16.5 | 45.7 | 32.5 | 30.3 |
| w/o $\mathcal{L}_{real}^{align}$ | 14.8 | 84.9 | 25.9 | 47.2 | 47.6 | 56.6 | 13.0 | 88.5 | 40.7 | 8.6 | 66.0 | 66.3 | 50.7 | 47.0 |
| w/o $\mathcal{L}_{real}^{dual}$ | 32.8 | 76.5 | 42.0 | 49.5 | 47.5 | 58.2 | 21.0 | 69.5 | 51.9 | 11.2 | 51.5 | 63.7 | 41.5 | 47.4 |
| Larger Normalized Scale (128 ²) | 15.0 | 85.1 | 40.4 | 50.6 | 58.6 | 72.5 | 35.2 | 94.0 | 58.7 | 17.4 | 43.7 | 79.8 | 50.4 | 54.0 |
| Un-normalized Scale Prediction | 27.9 | 70.6 | 34.9 | 22.2 | 36.2 | 54.0 | 20.4 | 86.8 | 43.5 | 4.9 | 51.4 | 54.4 | 33.9 | 41.6 |
| Normalized Scale Prediction | 22.2 | 63.2 | 46.8 | 11.1 | 37.6 | 51.4 | 20.3 | 80.2 | 31.0 | 6.9 | 37.0 | 40.7 | 35.8 | 37.2 |
| <i>Weakly-supervised Segmentation (IoU metric)</i> | | | | | | | | | | | | | | |
| Ours | 95.7 | 92.5 | 89.4 | 93.9 | 92.6 | 93.9 | 94.2 | 95.9 | 87.5 | 90.5 | 90.4 | 76.2 | 85.2 | 90.6 |
| w/o 2D BBox | 95.2 | 91.9 | 88.7 | 93.4 | 91.8 | 87.5 | 88.4 | 94.9 | 87.1 | 85.8 | 85.9 | 68.8 | 82.0 | 87.8 |
| w/o Real RGB images | 70.9 | 91.5 | 85.9 | 92.8 | 90.6 | 85.9 | 85.4 | 93.3 | 84.9 | 84.9 | 84.8 | 66.4 | 80.8 | 84.5 |

Table 4. Ablation Study on LINEMOD dataset. For self-supervised keypoint learning, we report the Average Recall of ADD(-S). For weakly-supervised segmentation, we report the Intersection over Union (IoU) score.

4.5. Ablation Study

In this section, we analyze the main components of our methods and evaluate the impact of each component on our final pose estimation performance in Table 4. All the experiments are implemented on LINEMOD dataset.

Training data: Since we use both synthetic and real data for training, we dissect the contributions of each modality of data. As seen in Table 4, without real data, our performance degrades significantly and is similar to other RGB based methods. This demonstrates the large domain gap between our synthetic and real data. Moreover, we also incorporate 3D labels of real images to show our method’s upper bound, marked as Ours[†].

Attention map: We introduce an attention map to compute the weighted average of a keypoint location. For the ablation study, we remove the attention map and directly average all the estimated offsets. Table 4 manifests that our attention map improves the keypoint robustness against the estimation noise.

Impacts of different losses: Our DSC-PoseNet employs a dual-scale strategy to estimate poses. In the training stage, we regularize the keypoint predictions to be consistent for both un-normalized and normalized scales. The network is also optimized by aligning pseudo-labeled and visible rendered masks. Without either of the above self-supervised losses, $\mathcal{L}_{real}^{dual}$ or $\mathcal{L}_{real}^{align}$, the performance of our method decreases significantly, as seen in Table 4.

Effects of the normalized scale: As shown in Table 4, the performance of our method degrades significantly without the normalized scale. Moreover, our method is robust for a larger size of the normalized scale (*i.e.*, 128×128), which leads to a similar result. In the inference stage, we

ensemble the normalized scale with the un-normalized scale by averaging their keypoint predictions. Table 4 shows that our method achieves much better performance by using ensemble keypoints.

Importance of 2D annotations: 2D bounding-box annotations are the only labels for DSC-PoseNet learning. We also investigate how our framework performs when 2D bounding-box annotations are not available. In this case, we employ PBR images to train our segmentation and then generate pseudo labels to finetune the segmentation network iteratively. As expected, some background pixels would be misclassified as the foreground regions, and the IoU score of the segmentation network degrades. Inaccurate segmentation also eliminates the advantage of using real images. Hence, our bounding-box annotations guarantee the performance of pose estimation while reducing domain gaps.

5. Conclusion

This paper proposes a novel two-step object pose estimation approach and significantly improve the performance of state-of-the-art RGB based models that are trained without real pose annotations. Our proposed pose estimation network, dubbed DSC-PoseNet, only uses easily attained 2D bounding-boxes annotations in training. Thanks to our visible silhouette alignment and dual-scale consistency self-supervision losses, DSC-PoseNet can be trained to estimate object poses without ground-truth supervision and also provides a solution to involve real images when their 3D pose labels are not available. Furthermore, although our network is designed for unlabeled real images, it can also be trained with real pose annotations. Hence, our DSC-PoseNet can be adapted to different scenarios, making it more favorable.

References

- [1] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. International Society for Optics and Photonics, 1992. 2, 7
- [2] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *ECCV*, pages 536–551. Springer, 2014. 6
- [3] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *CVPR*, pages 8100–8109, 2020. 5
- [4] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *ECCV*, pages 695–714. Springer, 2020. 1, 3
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018. 3, 6
- [6] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 4
- [7] Chunhui Gu and Xiaofeng Ren. Discriminative mixture-of-templates for viewpoint classification. In *ECCV*, pages 408–421. Springer, 2010. 2
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [9] Stefan Hinterstoisser, Cedric Cagniart, Slobodan Ilic, Peter Sturm, Nassir Navab, Pascal Fua, and Vincent Lepetit. Gradient response maps for real-time detection of textureless objects. *TPAMI*, 34(5):876–888, 2011. 2
- [10] Stefan Hinterstoisser, Stefan Holzer, Cedric Cagniart, Slobodan Ilic, Kurt Konolige, Nassir Navab, and Vincent Lepetit. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *ICCV*, pages 858–865. IEEE, 2011. 2
- [11] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *ACCV*, pages 548–562. Springer, 2012. 6
- [12] Tomas Hodan, Daniel Barath, and Jiri Matas. Epos: Estimating 6d pose of objects with symmetries. In *CVPR*, pages 11703–11712, 2020. 2
- [13] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: Benchmark for 6d object pose estimation. In *ECCV*, pages 19–34, 2018. 7
- [14] Tomáš Hodaň, Vibhav Vineet, Ran Gal, Emanuel Shalev, Jon Hanzelka, Treb Connell, Pedro Urbina, Sudipta N Sinha, and Brian Guenter. Photorealistic image synthesis for object instance detection. In *ICIP*, pages 66–70. IEEE, 2019. 6
- [15] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. *arXiv preprint arXiv:1911.08324*, 2019. 2
- [16] Roman Kaskman, Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Homebreweddb: Rgb-d dataset for 6d pose estimation of 3d objects. In *ICCV Workshops*, 2019. 6
- [17] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *ICCV*, pages 1521–1529, 2017. 2, 7
- [18] Park Kiru, Patten Timothy, and Vincze Markus. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *ICCV*, 2019. 2
- [19] Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. Cosypose: Consistent multi-view multi-object 6d pose estimation. In *ECCV*, 2020. 2, 3
- [20] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *IJCV*, 81(2):155, 2009. 2
- [21] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *ECCV*, pages 683–698, 2018. 2
- [22] Jinhui Liu, Zhikang Zou, Xiaoqing Ye, Xiao Tan, Errui Ding, Feng Xu, and Xin Yu. Leaping from 2d detection to efficient 6dof object pose estimation. In *ECCV*, pages 707–714. Springer, 2020. 2
- [23] Ming-Yu Liu, Oncel Tuzel, Ashok Veeraraghavan, and Rama Chellappa. Fast directional chamfer matching. In *CVPR*, pages 1696–1703. IEEE, 2010. 2
- [24] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 2
- [25] Fabian Manhardt, Diego Martin Arroyo, Christian Rupprecht, Benjamin Busam, Nassir Navab, and Federico Tombari. Explaining the ambiguity of object detection and 6d pose from visual data. *arXiv preprint arXiv:1812.00287*, 2018. 6, 7
- [26] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6d pose refinement in rgb. In *ECCV*, pages 800–815, 2018. 7
- [27] Aiden Nibali, Zhen He, Stuart Morgan, and Luke Prendergast. Numerical coordinate regression with convolutional neural networks. *arXiv preprint arXiv:1801.07372*, 2018. 4
- [28] Kiru Park, Timothy Patten, and Markus Vincze. Neural object learning for 6d pose estimation using a few cluttered images. In *ECCV*, 2020. 2
- [29] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *CVPR*, pages 4561–4570, 2019. 2, 4, 7
- [30] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *ICCV*, pages 3828–3836, 2017. 2, 6
- [31] Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Domain transfer for 3d pose estimation from color images without manual annotations. In *ACCV*, pages 69–84. Springer, 2018. 2, 6, 7

- [32] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 2, 5
- [33] Reyes Rios-Cabrera and Tinne Tuytelaars. Discriminatively trained templates for 3d object detection: A real time scalable approach. In *ICCV*, pages 2048–2055, 2013. 2
- [34] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6d object pose estimation under hybrid representations. *arXiv preprint arXiv:2001.01869*, 2020. 4
- [35] Hao Su, Charles R Qi, Yangyan Li, and Leonidas J Guibas. Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In *ICCV*, pages 2686–2694, 2015. 2
- [36] Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 240–248. Springer, 2017. 5
- [37] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *ECCV*, pages 699–715, 2018. 2, 6, 7
- [38] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *CVPR*, pages 292–301, 2018. 2, 6
- [39] Yurun Tian, Xin Yu, Bin Fan, Fuchao Wu, Huub Heijnen, and Vassileios Balntas. Sosnet: Second order similarity regularization for local descriptor learning. In *CVPR*, pages 11016–11025, 2019. 2
- [40] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *CVPR*, pages 1510–1519, 2015. 2
- [41] Gu Wang, Fabian Manhardt, Jianzhun Shao, Xiangyang Ji, Nassir Navab, and Federico Tombari. Self6d: Self-supervised monocular 6d object pose estimation. In *ECCV*, 2020. 1, 2, 4, 6, 7
- [42] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *CVPR*, pages 2642–2651, 2019. 2
- [43] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *Robotics: Science and Systems*, 2017. 2, 6
- [44] Zongxin Yang, Linchao Zhu, Yu Wu, and Yi Yang. Gated channel transformation for visual recognition. In *CVPR*, pages 11794–11803, 2020. 6
- [45] Xin Yu, Yurun Tian, Fatih Porikli, Richard Hartley, Hongdong Li, Huub Heijnen, and Vassileios Balntas. Unsupervised extraction of local image descriptors via relative distance ranking loss. In *ICCV Workshops*, pages 0–0, 2019. 2
- [46] Xin Yu, Zheyu Zhuang, Piotr Koniusz, and Hongdong Li. 6dof object pose estimation via differentiable proxy voting regularizer. In *BMVC*, pages 1–13, 2020. 2
- [47] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *ICCV*, 2019. 1, 2, 6, 7
- [48] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. Weakly-supervised salient object detection via scribble annotations. In *CVPR*, pages 12546–12555, 2020. 1
- [49] Li Zhigang, Wang Gu, and Ji Xiangyang. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *ICCV*, 2019. 1, 2, 6, 7
- [50] Menglong Zhu, Konstantinos G Derpanis, Yinfei Yang, Samarth Brahmabhatt, Mabel Zhang, Cody Phillips, Matthieu Lecce, and Kostas Daniilidis. Single image 3d object detection and pose estimation for grasping. In *ICRA*, pages 3936–3943. IEEE, 2014. 2
- [51] Zheyu Zhuang, Xin Yu, and Robert Mahony. Lyrn (lyapunov reaching network): A real-time closed loop approach from monocular vision. In *ICRA*, pages 8331–8337. IEEE, 2020. 2