

PFRL: Pose-Free Reinforcement Learning for 6D Pose Estimation

Jianzhun Shao Yuhang Jiang Gu Wang Zhigang Li Xiangyang Ji
Department of Automation and BNRist, Tsinghua University
Beijing, China

{sjz18,jiangyh19,wangg16,lzg15}@mails.tsinghua.edu.cn xyji@tsinghua.edu.cn

Abstract

6D pose estimation from a single RGB image is a challenging and vital task in computer vision. The current mainstream deep model methods resort to 2D images annotated with real-world ground-truth 6D object poses, whose collection is fairly cumbersome and expensive, even unavailable in many cases. In this work, to get rid of the burden of 6D annotations, we formulate the 6D pose refinement as a Markov Decision Process and impose on the reinforcement learning approach with only 2D image annotations as weakly-supervised 6D pose information, via a delicate reward definition and a composite reinforced optimization method for efficient and effective policy training. Experiments on LINEMOD and T-LESS datasets demonstrate that our Pose-Free approach is able to achieve state-of-the-art performance compared with the methods without using real-world ground-truth 6D pose labels.

1. Introduction

6D pose estimation aims to localize the 3D location and the 3D orientation of the object from a single image. It plays a crucial role in real-world applications including robot manipulation [5, 43], augmented reality [21] and self-driving cars [4, 19]. For instance, when a robot tries to grab an object, it is a prerequisite to accurately estimate the 6D object poses from the image captured by the equipped camera.

This problem was traditionally regarded as a geometric problem and can be solved with Perspective-n-Point (PnP) [14] algorithms by matching features between the 2D image and the 3D object model. Since rich textures are indispensable for feature matching, they cannot effectively handle the texture-less cases. With the rise of deep learning, considerable learning-based approaches are proposed for various perception cases (e.g., RGB-only, RGB-D) and have greatly promoted the evolvement of this field. Some people [24, 29, 25] followed the conventional way to build the 2D-3D correspondences and subsequently solve the pose via

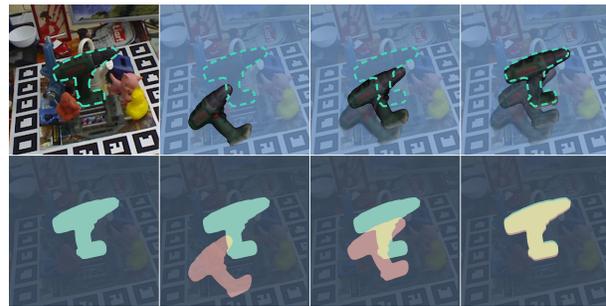


Figure 1. The sketch of Pose-Free Reinforcement Learning (PFRL). We formulate the pose refining procedure as a sequential decision-making process and exploit a composite reinforced optimization approach to tackle it. We calculate the reward using the similarity measurement of mask information to avoid the annotations of pose labels.

PnP. Others [12, 38] instead trained the deep model in an end-to-end manner to directly derive the pose from the image. Compared with traditional ones, these approaches can achieve remarkable performance even in challenging situations such as texture-less, clutter, occlusion, etc. A common way to further improve the pose accuracy is pose refinement. DeepIM [15] is able to predict the 6D pose difference between the current estimation and the ground-truth by comparing the rendered object with the observed one in 2D space. A similar idea is also explored by [20, 40].

Real-world pose annotations are essential for current approaches to achieve excellent performance. Unfortunately, acquisition of the 2D images annotated with real-world ground-truth 6D object poses is fairly cumbersome and expensive compared with that in 2D space (e.g. instance 2D mask), which greatly restricts the popularization of these approaches. Alternatively, some researchers [12, 33, 40] turned to getting rid of this problem by training on synthetic data only. However, the domain gap is prone to degrade pose estimation performance in real-world cases.

In this work, we focus on 6D pose estimation from RGB images only without requiring real-world pose annotations

and avoiding the domain gap caused by training the model on synthetic-only data.

To achieve this goal, we introduce reinforcement learning by formulating the 6D pose refinement problem as a Markov Decision Process (MDP). Given an initial pose estimate, our approach learns to rotate and translate the 3D object model sequentially to align it to the observation. To train the model, we design a 2D mask-based reward which is computed only from the rendered and the target 2D masks. For each episode, the accumulated rewards are maximized, resulting in a policy that can approach the target smoothly and quickly (Fig. 1). We also propose a composite reinforced optimization method to learn the operation policy efficiently and effectively. Compared with existing supervised approaches, our method utilizes a sequential decision-making process to optimize a delayed accumulated reward from the weakly-supervised mask similarity instead of the supervised loss from the ground-truth 6D pose. Since no ground-truth 6D pose information is involved, we call our proposed approach a Pose-Free Reinforcement Learning (PFRL) approach for 6D pose estimation.

Experiments on LINEMOD [9] and T-LESS [10] datasets demonstrate that our approach is able to achieve state-of-the-art performance compared with existing methods without using ground-truth 6D pose labels. Our work makes the following main contributions: i) We formulate the 6D pose estimation problem as a Markov Decision Process and propose a Pose-Free Reinforced Learning solution, which is able to exploit 2D image annotations as weakly supervised information and reward-based sequential decision making for 6D pose refinement. ii) We design a low-cost reward strategy based on the 2D mask and propose a composite reinforced optimization method for efficient and effective policy training. iii) On LINEMOD and T-LESS datasets we achieve state-of-the-art performance compared to the methods without using real-world ground-truth 6D pose labels.

2. Related Work

RGB-based 6D Object Pose Estimation. Traditionally, the pose estimation was considered as a geometric problem and solved with Perspective-n-Point (PnP) [14] algorithms by matching the features between the 2D image and the 3D model. However, they rely heavily on the textures and cannot handle the textures-less cases. Recently, the deep learning technique has significantly advanced the development of pose estimation. Some works followed the traditional way by training the model to build the 2D-3D correspondences by 1) detecting the pre-defined keypoints from image [24, 29, 36, 41] or 2) predicting the 3D coordinates for object pixels [2, 37, 40, 16]. For the former, the sparse 2D-3D correspondences make the estimator sensitive to the occlusion. For the latter, RANSAC [23] is necessary

to solve the pose from the dense correspondences, which remarkably increase the computational complexity of the inference. Different from these PnP-based approaches, others trained the deep model in an end-to-end manner to predict the pose directly from an image. [12] trained a viewpoint classifier on the SSD [17] by discretizing the 3D rotation space into some bins. The translation was calculated from the 2D bounding box. [33] leveraged an Augmented Auto-Encoder (AAE) to encode features for each discrete rotation to build a dictionary. During test, the rotation is retrieved by matching the feature with the dictionary. [38] developed a CNN-based network to regress the quaternions from image directly. These direct approaches can enjoy fast inference speed, but they typically suffer from inferior performance when compared to PnP-based approaches.

Current RGB-based pose estimation approaches rely heavily on real-world annotated training data. For instance, with the aid of real data, the pose estimation precision of DPOD [40] can be improved from 50% to 82.98% on metric ADD. Without real training samples, the performance of current approaches is far from satisfactory. Unfortunately, the pose annotations of real data are quite expensive and time-consuming. We propose Pose-Free Reinforced Learning (PFRL) framework, which can significantly improve the pose estimation accuracy only based on the relatively low-cost 2D real-data annotations (i.e. segmentation masks).

6D Object Pose Refinement. To further promote the performance of 6D pose refinement, numerous approaches leveraged depth by ICP [1]. As a result, additional depth sensor is required and may be restricted by its the frame rate, resolution, depth range and light illumination, etc. For RGB-based pose refinement, a promising way is to train a pose refiner. For instance, DeepIM [15] was proposed to compare the rendered objects (accompanied with mask) with the observed ones, which can iteratively predict the relative transformation between the current pose estimation and the ground-truth. [20] and [40] developed the similar idea by introducing the parallel branches to extract features. These approaches can achieve remarkable pose estimation performance by utilizing the ground-truth pose annotations to provide fully-supervised training signals. More recently, [26] attempted to train a 6D object pose refiner without real 6D pose labels via an approximated differentiable renderer, however the utilization of the differentiable renderer was non-trivial, since they required an additional representation learning step where real 6D pose annotations were still needed.

Reinforcement Learning for Vision Tasks. Recently, deep reinforcement learning [35] has been successfully applied to a wide range of computer vision tasks, such as object detection [22], semantic segmentation [42], visual object tracking [34, 39, 30], vision-based robotic grasping [28], etc. Several works also try to solve 6D object pose

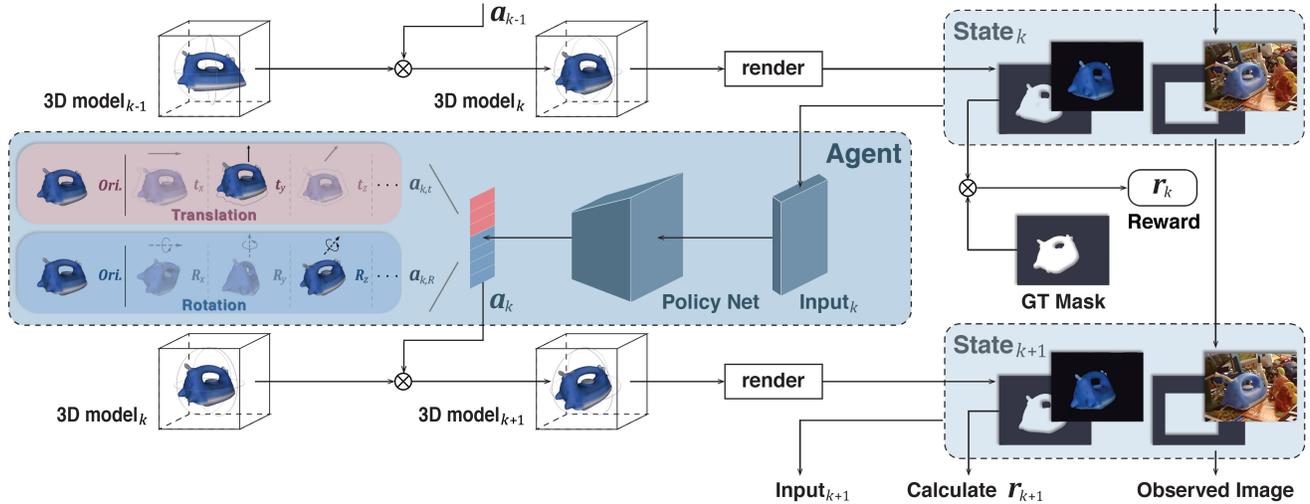


Figure 2. The PFRL framework. At each time step k , we use the cropped observed image, ground-truth bounding box, rendered image and rendered mask to form state s_k . The policy network (detailed in Sec. 3.4) takes s_k as input and generates disentangled action \mathbf{a}_k , which represents the relative SE(3) transformation for current pose (detailed in Sec. 3.1). The environment (i.e. renderer) then renders a new image according to the new pose, calculating the reward from current mask and the ground-truth mask (detailed in Sec. 3.2). Finally the composite policy optimization is executed using the reward (detailed in Sec. 3.3).

estimation related tasks via reinforcement learning. For instance, [13] proposes a policy gradient based method to save the inference budget for an existing 6D object pose estimation system. [32] selects the strategy of camera movements by policy gradient method, directing the camera’s attention to the most uncertain object.

Different from them, we propose a novel approach to learn the 6D pose refinement in a Pose-Free fashion without the need of real-world 6D pose annotations via reinforcement learning. By formulating the problem under the reinforcement learning framework, the information from the non-differentiable rasterization-based renderer can be better exploited during optimization. It enjoys both Pose-Free 2D annotations and pose refinement prediction capability.

3. Methodology

The overview of our framework is depicted in Fig. 2. In following sections, we first present the formulation of Reinforcement Learning (RL) for solving the 6D pose estimation problem in Sec. 3.1. We then introduce a 2D mask-based reward function in Sec. 3.2. Further, a composite reinforced optimization is proposed for the task-specific 6D pose estimation problem in Sec. 3.3. Finally, we discuss our policy network architecture in Sec. 3.4.

3.1. Problem Formulation

To achieve accurate 6D pose estimation, our core idea is to align the 2D projection of 3D object with its appearance observed in 2D image. We propose the PFRL framework to

tackle the problem by reinforcement learning. We hypothesize that the decision made by the agent at each time step will be only based on the current state and not be affected by previous states, which can be formulated as an MDP. The goal is to maximize the expected sum of future discounted rewards $V^\pi(\mathbf{s}) = \mathbb{E}[\sum_{k \geq 0} \gamma^k r_k]$ [27], where $\gamma \in [0, 1)$ is the discount factor, $r_k = r(\mathbf{s}_k, \mathbf{a}_k)$ is immediate reward at time k , \mathbf{s}_k is the state at time k and $\mathbf{a}_k \sim \pi(\cdot | \mathbf{s}_k)$ is the action generated by following some policy π conditioned on \mathbf{s}_k . The state and action space of our framework are defined as follows:

State: The state which encodes the knowledge of the environment should be instrumental for the agent to decide how to rotate and translate the 3D model. At current time step k , the state s_k is formed by a rendered RGB image and a projected mask of the 3D object model specified by the current pose, and the observed RGB image as well as the ground-truth 2D bounding box of the object. These four images are concatenated together as the input of the policy network. During test, since the ground-truth bounding box is not available, we then exploit it by projecting the object with the initial pose instead.

Action: Action is the relative SE(3) transformation the agent performs at each time step, which will affect the state of the environment. We exploit the same representation of $\Delta \mathbf{p}$ as in [15] for its characteristics of disentangled rotation and translation, i.e., the actions of rotation and translation are not mutually affected. The rotation and translation actions are chosen stochastically from the two distributions parameterized by policy network. Different from con-

ventional reinforcement learning approaches, we measure a single policy value for the two distributions by concatenating the features of rotation and translation together (Sec. 3.4). We evaluate continuous and discrete distributions for the policy model respectively:

$$\mathbf{a}_R \sim \mathcal{N}(\boldsymbol{\mu}_R(\theta), \boldsymbol{\sigma}_R^2), \mathbf{a}_t \sim \mathcal{N}(\boldsymbol{\mu}_t(\theta), \boldsymbol{\sigma}_t^2), \quad (1)$$

and

$$\mathbf{a}_R \sim \mathcal{C}(\boldsymbol{\zeta}_R(\theta)), \mathbf{a}_t \sim \mathcal{C}(\boldsymbol{\zeta}_t(\theta)). \quad (2)$$

Eq. 1 is the continuous distributions for rotation and translation actions. $\mathcal{N}(\cdot)$ represents multivariate Gaussian distribution. $\boldsymbol{\mu}_R(\theta)$ and $\boldsymbol{\mu}_t(\theta)$ represent the mean values of rotation (i.e. quaternion) and translation from the policy network respectively. Accordingly, $\boldsymbol{\sigma}_R$ and $\boldsymbol{\sigma}_t$ are the variance matrices. We assume the dimensions are independent of each other for simplicity and without loss of representation ability, hence both of $\boldsymbol{\sigma}_R$ and $\boldsymbol{\sigma}_t$ are diagonal matrices.

Eq. 2 indicates the discrete distribution for rotation and translation actions. $\mathcal{C}(\cdot)$ represents Categorical distribution, where $\boldsymbol{\zeta}_R(\theta)$ and $\boldsymbol{\zeta}_t(\theta)$ are the probabilities produced by the policy network.

We handcraft 13 operations for rotation and translation actions, where each of them includes rotating (or translating) 3D model along $\pm x, \pm y, \pm z$ with a fixed small and large degree (or step) as well as no rotation/translation action. See Sec. 1 of the supplementary material for more details.

Each operation for rotation and translation is encoded with a quaternion and a direction vector. The final form sampled from the continuous and discrete distribution is the same.

3.2. 2D Mask-based Reward

After rotating and translating the 3D model, the agent needs a reward signal r_A for policy update. Under the Pose-Free situation, we propose a 2D mask-based reward, which includes three components: IoU (Intersection over Union) Difference Reward r_I , Goal Reached Reward r_G , and Centralization Reward r_C , which are purely computed from the rendered and ground-truth masks in 2D space.

We denote the IoU Difference Reward r_I as:

$$r_I = f_\phi(\text{IoU}_{k+1}) - f_\phi(\text{IoU}_k), \quad (3)$$

where

$$f_\phi(x) = \begin{cases} x, & x < X_{\text{thr}} \\ \alpha x^2 - \beta x, & x \geq X_{\text{thr}} \end{cases}. \quad (4)$$

Let M_R and M_G represent the rendered mask at current state and the ground-truth mask respectively. In Eq. 3 $\text{IoU}_k = \frac{S(M_{R_k} \cap M_G)}{S(M_{R_k} \cup M_G)}$ measures their IoU. f_ϕ is a mapping function, whose shape is controlled by α , β , and a threshold X_{thr} . The motivation is that if the mask of an object without

high symmetry has a perfect overlap with the ground-truth mask, it is very likely for the object to be in the actual localization and orientation. Therefore the pose estimation problem can be converted into a 2D mask matching problem. Our objective is to maximize the IoU at the last image frame. Therefore, we design r_I in the form of IoU difference in adjacent frames corresponding to the RL objective: maximizing the cumulative reward $\sum_0^{k_{\text{end}}} r_k$. The mapping function f_ϕ is designed to give a larger reward when IoU is closer to 1. Sharper changes of reward at frames with large IoUs (e.g., 0.8 to 1.0), which is the most common situation during training, help the agent learn the objective much easier.

In practice, when the IoU reaches some threshold IoU_{thr} , we expect the environment to give the reward feedback immediately and stop refining. Thus we define an additional Goal Reached Reward r_G as:

$$r_G = \begin{cases} 1, & \text{IoU}_k \geq \text{IoU}_{\text{thr}} \\ 0, & \text{IoU}_k < \text{IoU}_{\text{thr}} \end{cases}. \quad (5)$$

Such reward gives the agent an evident arrival signal, which saves the budget in real-world applications.

In case of the object locating far away from the true position, which might happen if the network is not well initialized, we further add a Centralization Reward r_C to give an explicit constraint on the estimate of translation. r_C is denoted as:

$$r_C = \min(\|c_r - c_g\|_2^{-\frac{1}{2}}, 1), \quad (6)$$

where c_r and c_g are the center of the rendered mask and the ground-truth mask, respectively.

Above all, the final reward r_A can be summarized as:

$$r_A = r_I + \sigma_C r_C + \sigma_G r_G, \quad (7)$$

where σ_C and σ_G are weights of r_C and r_G .

Note that if an object exhibits some symmetry, naively training a pose estimator is problematic and unstable since the similar appearances in RGB can show different 6D poses and lead to distinct loss values [18]. Our design of mask-based reward can implicitly avoid this problem and work well on symmetric objects, similar to the idea in [33].

3.3. Composite Reinforced Optimization

The high dimensional state space brings high variance and instability for training the policy network. In usual, millions of samples are required to fully exploit the power of reinforcement learning model. However, in the 6D pose estimation, it is difficult to render such tremendous amount of the images. In this section, we propose a task-specific composite reinforced optimization method for 6D pose estimation. We combine the on-policy and off-policy strategies together to fully utilize the rendered images.

We use θ and ϕ to denote the current policy network and value function parameters. Given the tuple (s_k, \mathbf{a}_k, r_k) , we use θ_{old} to denote the parameters of the policy network where \mathbf{a}_k samples before θ . The corresponding policies are denoted as $\pi_\theta(\mathbf{a}_k|s_k)$ and $\pi_{\theta_{\text{old}}}(\mathbf{a}_k|s_k)$. The on-policy optimization refers to learning the value $V_\phi(s_k)$ of the policy $\pi_\theta(\mathbf{a}_k|s_k)$ being carried out by the agent with corresponding \mathbf{a}_k . In this paper, we employ the Proximal Policy Optimization (PPO) algorithm [31] for the on-policy optimization. The loss function of PPO is defined as:

$$L_{\text{on}} = L_{\text{clip}} + \lambda_v L_{\text{value}} + \lambda_e L_{\text{entropy}}. \quad (8)$$

In Eq. 8, L_{clip} is the clipped surrogate objective measuring minimum of clipped and original importance-weighted advantage value function. λ_v and λ_e are trade-off parameters. L_{value} measures an approximated squared-error loss between value function and cumulative reward. L_{entropy} is the entropy bonus of policy. For further details of the loss function, we refer readers to [31].

After updating policy network, the tuple (s_k, \mathbf{a}_k, r_k) are no longer associated with the π_θ , i.e., \mathbf{a}_k was sampled by a previous policy network $\pi_{\theta_{\text{old}}}$. Therefore the tuple (s_k, \mathbf{a}_k, r_k) can not be used for on-policy optimization. In order to further fully utilize “outdated” data samples and take advantage of data efficiency, we introduce an off-policy value update strategy similar to [7] to assist the on-policy optimization to accelerate the training process. We set up a priority queue replay buffer to store the “outdated” data samples. We then update the value function $V_\phi(s_k)$ with off-policy loss using samples from the replay buffer:

$$L_{\text{off}} = (V_\phi(s_k) - V_{\text{trace}})^2. \quad (9)$$

Definitions of V_{trace} in Eq. 9 can be found in [7]. The overall procedure of the composite reinforced optimization is summarized Algorithm 1.

Algorithm 1: Composite Reinforced Optimization

Initialize policy network π_θ and value function V_ϕ ;
Initialize replay buffer R .
for episode = 1, M **do**
 Generate trajectories $T = (s_0, \mathbf{a}_0, r_0, s_1, \dots)$;
 Update π_θ and V_ϕ with L_{on} using T ;
 Store T into R ;
 Sample $B = (s, \mathbf{a}, r)$ pairs from R ;
 Compute V-trace target V_{trace} using B ;
 Update V_ϕ with L_{off} .

end

3.4. Disentangled Policy Architecture

Fig. 3 depicts the disentangled policy network structure. The backbone consists of the first 11 convolutional layers

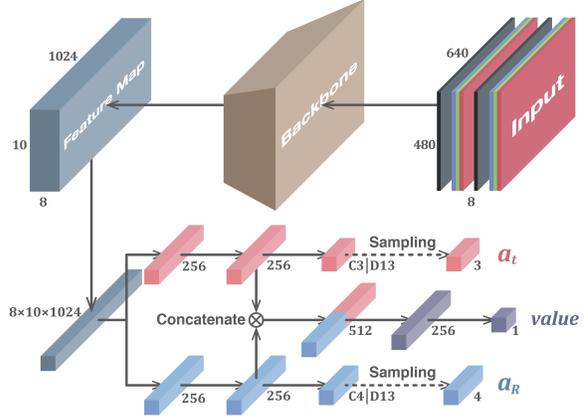


Figure 3. Disentangled Policy Architecture. The observed image, ground-truth bounding box, rendered image, rendered mask are concatenated and zoomed together as input. Branches for rotation and translation are separated and actions are sampled from continuous or discrete distributions. ‘C3’ and ‘C4’ denote the continuous distribution for translation and rotation while the double ‘D13’ denote the discrete distribution. The value function is computed together.

of FlowNet-S [6]. Similar to [15], the input images and masks are zoomed-in before feeding to the backbone. As discussed in Sec. 3.1, two separate branches are used to predict the distribution of the disentangled rotation and translation. Both branches contain 2 fully-connected (FC) layers with 256 hidden units, and are followed by one additional FC layer. Compared with [15], our network is lightweight both in training and inference, since we do not use the heavily computational flow and mask branches. The last FC layers of both branches are concatenated together followed by 2 FC layers to produce the value function. In this manner, the value function shares the same parameters as actions, and can be used for optimization of both actions.

4. Experiments

4.1. Data Preparation

We conduct our experiments on LINEMOD [9] and T-LESS [10] dataset. We split the LINEMOD dataset following [3], including about 15% for training and 85% for test, i.e., around 200 images per object for the training set. Note that no 6D pose annotations are directly used during training and testing. They are only used to generate ground-truth masks. Under the case that 6D pose is hard or impossible to obtain, we can get the ground-truth mask by semantic segmentation or manual annotation.

On LINEMOD dataset, we use results from AAE [33] and PoseCNN [38] as our initial pose for test. Random Gaussian noise with zero mean is added to each dimen-

Train data	Pose-Free Init		+Pose-Free Refine			Gt Pose Init		+Pose-Free Refine			
	AAE[33] ADD	DPOD-SYN[40] ADD	SSD6D[12] ADD	DPOD-SYN+Refine ADD	AAE+ours ADD	YOLO6D[36] ADD	PoseCNN[38] ADD	PoseCNN+DeepIM-SYN[15] Proj.2D	PoseCNN+DeepIM-SYN[15] ADD	PoseCNN+ours Proj.2D	PoseCNN+ours ADD
Ape	3.96	37.22	-	55.23	65.4	21.6	27.8	81.7	23.9	95.6	60.5
Benchvise	20.92	66.76	-	72.69	84.5	81.8	68.9	92.2	93.1	94.7	88.9
Camera	30.47	24.22	-	34.76	41.5	36.6	47.5	97.0	84.7	95.0	64.6
Can	35.87	52.57	-	83.59	80.9	68.8	71.4	89.6	91.5	93.1	91.3
Cat	17.90	32.36	-	65.10	80.4	41.8	56.7	96.1	79.5	99.3	82.9
Driller	23.99	66.60	-	73.32	77.6	63.5	65.4	85.9	82.3	94.8	92.0
Duck	4.86	26.12	-	50.04	52.5	27.2	42.8	92.6	24.0	98.2	55.2
Eggbox	81.01	73.35	-	89.05	96.1	69.6	98.3	90.8	88.3	97.8	99.4
Glue	45.49	74.96	-	84.37	76.7	80.0	95.6	81.2	96.9	97.1	93.3
Holepuncher	17.60	24.50	-	35.35	44.9	42.6	50.9	78.0	20.6	96.7	66.7
Iron	32.03	85.02	-	98.78	67.3	75.0	65.6	59.3	85.1	81.6	75.8
Lamp	60.47	57.26	-	74.27	91.1	71.1	70.3	75.6	85.5	96.0	96.6
Phone	33.79	29.08	-	46.98	52.7	47.7	54.6	88.3	66.1	91.0	69.1
Mean	31.41	50.0	34.1	66.43	70.1	56.0	62.7	85.3	70.9	94.7	79.7

Table 1. Comparison with state-of-the-art Pose-Free methods on LINEMOD with metrics ADD and Proj. 2D. Left part is results of initial pose trained on synthetic data + Pose-Free refiner, and right part is results of initial pose trained on ground-truth labels + Pose-Free refiner. PFRL (ours) outperforms the state-of-the-art method DPOD trained with synthetic data, despite that we use a much worse initialization method provided by AAE. When using the same initial pose from PoseCNN, PFRL outperforms DeepIM trained on synthetic data.

sion of the rotation and translation as initial pose for training. The variance is $(15^\circ)^2$ for each rotation axis and $[2^2, 2^2, 5^2](cm)^2$ for translation $[x, y, z]$. A noise is resampled if one of the rotation axis exceeds 45 degrees.

On T-LESS dataset we initialize our method with results from AAE [33], the current state-of-the-art RGB-only method on T-LESS. No real images are available for training, so we render objects with random poses and light conditions on random background images from PASCAL VOC [8]. We perturb the ground-truth poses with the same noise used for LINEMOD as initial poses. Since the camera intrinsic matrix of each real image differs only in the principal points slightly, we translate and crop the images such that they have the same intrinsic matrix for ease of use.

4.2. Evaluation Metrics

In our experiments, we use three common metrics for evaluation: Proj. 2D, ADD and VSD. For Proj. 2D metric, a pose is regarded as correct if the average Proj. 2D distance (Eq.(10)) of the model points is less than 5 pixels:

$$\text{Proj. 2D} = \frac{1}{m} \sum_{x \in \mathcal{M}} \|\mathbf{K}(\mathbf{R}\mathbf{x} + \mathbf{t}) - \mathbf{K}(\hat{\mathbf{R}}\mathbf{x} + \hat{\mathbf{t}})\|. \quad (10)$$

For ADD metric, a pose is regarded as correct if the average distance of model points is less than 10% of the model diameter [9] (Eq.(11)). For symmetric objects, the distance is calculated by the closest model point (Eq.(12)).

$$\text{ADD} = \frac{1}{m} \sum_{x \in \mathcal{M}} \|(\mathbf{R}\mathbf{x} + \mathbf{t}) - (\hat{\mathbf{R}}\mathbf{x} + \hat{\mathbf{t}})\|, \quad (11)$$

$$\text{ADD-S} = \frac{1}{m} \sum_{x_1 \in \mathcal{M}} \min_{x_2 \in \mathcal{M}} \|(\mathbf{R}\mathbf{x}_1 + \mathbf{t}) - (\hat{\mathbf{R}}\mathbf{x}_2 + \hat{\mathbf{t}})\|. \quad (12)$$

Visible Surface Discrepancy (VSD) [11] is an ambiguity-invariant metric that depends on the visible surface, which can account for the symmetry better. The parameters used for VSD metric are the same as [33].

4.3. Results on LINEMOD

Table 1 shows the results of our method and some recent well-performed RGB-based methods. We achieve state-of-the-art on Pose-Free methods. Most Pose-Free methods like DPOD [40] and AAE [33] render objects on random backgrounds such as images from PASCAL VOC dataset [8] and train on synthetic data. When the initial pose estimation and the refiner are both Pose-Free, our method achieves average ADD 70.1% with initial pose estimated by AAE (ADD 31.41%), surpassing the previous state-of-the-art method DPOD’s ADD 66.43% with initial pose ADD 50.0%. And on 10 of 13 objects we get a higher score than DPOD despite the relatively lower initial pose score. We bring on average 38.7 percentages ADD improvement on AAE.

To evaluate our method with better initial poses, we also use the results of PoseCNN [38] trained with real 6D pose labels as initial poses. We bring 17.0 percentages improvement on PoseCNN. Our method also surpasses DeepIM-SYN [15], the DeepIM trained with pure synthetic data and initialized with PoseCNN results during test, on both metrics overall. We can also see that on Proj. 2D metric, our method performs better on 12 of 13 objects, which demonstrates the advantage of our design of the 2D mask based reward. We also provide results with DPOD initialization in Sec. 3 of the supplementary material.

4.4. Results on T-LESS

To evaluate our method on synthetic-only cases, we train on objects 19-23 of T-LESS dataset. We use AAE [33] with

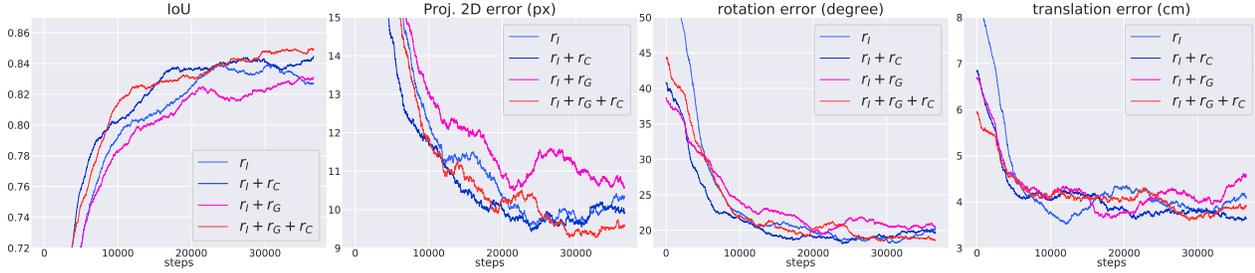


Figure 4. Curve of different reward styles.

ground-truth bounding box as our initial pose for testing. Objects 19-23 are the only objects trained with texture-less models for AAE. Since T-LESS dataset has no real images for training, we use synthetic data for training similar to AAE. The results are shown in Table 2.

Method Metric	AAE[33]		AAE+Ours	
	Proj. 2D	VSD	Proj. 2D	VSD
19	30.65	49.95	32.79	57.39
20	23.51	41.87	25.40	45.29
21	56.55	59.06	60.58	62.50
22	42.99	46.08	44.78	48.02
23	21.88	40.38	29.32	44.44
Mean	35.12	47.47	38.57	51.53

Table 2. Results on the T-LESS dataset. Accuracy in the left side is calculated from the models provided by the author of AAE [33] and is slightly better than the original reported results.

On texture-less objects without real training images, the improvement of metric recall is not as obvious as in LINEMOD due to the large domain gap. However we still get 3.5 percentages improvement on Proj. 2D and 4.1 percentages on the more widely used metric VSD. We believe further improvement can be achieved with the availability of real images with mask annotations, which is much cheaper to obtain than 6D pose annotations.

4.5. Ablation study

We do ablation study on the reward style, action style, testing speed and optimization strategy. For each ablation we use several objects from LINEMOD [9] for training and testing. We set batch size to 256, $X_{thr} = 0.5$, $IoU_{thr} = 0.98$, $\alpha = 3.8$, $\beta = 1.8$, $\sigma_C = 1$, $\sigma_G = 2$, $\lambda_v = 0.5$, $\lambda_e = 0.001$. All hyper-parameters are fixed unless otherwise stated.

Reward style. Our basic reward is the IoU Difference Reward r_I . We train on object Iron with the same parameters and training steps using r_I and the additional Goal Reached Reward r_G and the Centralization Reward r_C . Fig. 4 shows the comparison of four reward styles in four eval-

uation metrics. Although we can not see obvious difference from the average rotation and translation error since rewards are designed from the mask information, the IoU of $r_I + r_G + r_C$ increases fastest and maintains the highest level, while it has the lowest Proj. 2D error. Comparing $r_I + r_C$ with r_I , r_C brings obvious acceleration on the IoU improvement. Although adding r_G to r_I independently seems bad, it performs well when combined with r_C . Therefore in the rest part of our experiment, we all use reward $r_I + r_G + r_C$.

Action Style Metric	Continuous		Discrete	
	Proj. 2D	ADD	Proj. 2D	ADD
Initial	14.2	35.6	14.2	35.6
Epoch: 800	37.6	43.8	58.9	52.8
Epoch:1600	45.7	46.7	63.3	55.4
Epoch:2400	57.5	56.2	75.3	70.7
Epoch:3200	60.4	60.9	80.6	79.2

Table 3. Ablation on action style.

Action style. We train 2 models on the same conditions except that the action styles are continuous and discrete separately and test on the Benchvise object. Table 3 shows the results for different training epochs on the test set. We can tell that discrete action space outperforms continuous one for the whole training procedure, and performs around 20% better in convergence. Discretization makes great contribution to our method.

Effect of refining steps and testing speed. Our experiments were conducted on a desktop with an Intel Core i7 3.60GHz CPU and a NVIDIA GeForce RTX 2080 Ti GPU. We set episode length to 50 for one pose estimation process during training. We find that the object usually gets very close to the target pose in a few steps and keeps motionless in the next steps. Therefore we plot the curve of Proj. 2D accuracy with respect to refining steps in Fig. 6. Curves include 3 objects (Ape, Benchvise, Can) with 2 kinds of initial poses (AAE, PoseCNN). We can see that all objects reach their best accuracy at step 15-20. We summarize the



Figure 5. Qualitative results of AAE initial poses [33] and our refined poses on the LINEMOD dataset. The red and green lines represent the contours of the initial poses and our refined poses respectively.

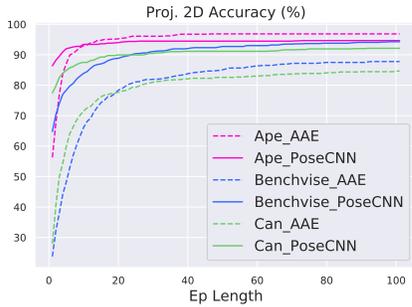


Figure 6. The Proj. 2D accuracy with respect to test episode length curves of 3 objects: Ape, Benchvise, Can.

accuracy and test speed(ms per frame) with respect to the episode length using AAE [33] as initial pose in Table 4. Our method is flexible in terms of accuracy and refining

Ep Length	Ape		Benchvise		Can		
	A	P	A	P	A	P	
Acc (%)	0	39.0	83.1	14.2	50.0	16.8	69.6
	5	87.0	92.0	48.2	78.4	59.8	84.8
	10	93.0	93.4	66.1	83.9	71.6	87.5
	20	95.2	94.0	78.4	88.9	77.7	90.0
	50	96.9	94.2	85.1	92.7	82.6	91.1
Time (ms)	5	62	64	67	65	68	68
	10	110	114	126	121	128	127
	20	218	220	243	234	245	243
	50	529	534	588	578	588	584

Table 4. Accuracy and testing time with respect to episode length. ‘A’ and ‘P’ denote initial poses from AAE [33] and PoseCNN [38], respectively.

steps. We find that a relatively good pose can be estimated in 20 steps, that is, in about 240ms per frame.

Optimization strategy. We compare the composite optimization strategy with the pure on-policy optimization strategy on objects Ape and Lamp under the same number of samples in Fig. 7. The 4 curves in each figure represents 2 objects with 2 optimization strategies separately. We use the same learning rate 10^{-4} for on-policy and off-policy update. The on-policy and off-policy update times are 4:1.

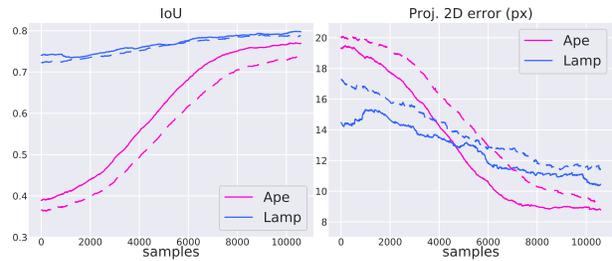


Figure 7. Composite optimization strategy (solid line) versus on-policy optimization strategy (dashed line).

The left and right shows the change of IoU and Proj. 2D error with the number of updated samples. We can tell that for the same object and the same update samples, mixed update rule has higher IoU and lower Proj. 2D error than on-policy update rule, which demonstrates that extra updates on value functions help the policy get an accurate value earlier without more samples acquired from the environment.

For more experiment results such as generalization ability and class-agnostic training, we kindly refer readers to supplementary material.

5. Conclusion

In this work we formulate the RGB-based 6D pose estimation problem as an MDP and introduce PFRL framework in a Pose-Free fashion without the need of real-world 6D pose annotations. We design a task-specified 2D mask-based reward which is purely computed from the object mask information and employ a composite reinforced optimization rule to learn the operation policy efficiently and effectively. The experiments demonstrate that our approach is able to achieve the state-of-the-art performance compared with the methods without using real-world ground-truth 6D pose labels on LINEMOD and T-LESS dataset.

6. Acknowledgement

This work was supported by the National Key R&D Program of China under Grant 2018AAA0102800 and 2018AAA0102801.

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] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *European conference on computer vision*, pages 536–551. Springer, 2014.
- [3] Eric Brachmann, Frank Michel, Alexander Krull, Michael Ying Yang, Stefan Gumhold, et al. Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3364–3372, 2016.
- [4] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017.
- [5] Alvaro Collet, Manuel Martinez, and Siddhartha S Srinivasa. The moped framework: Object recognition and pose estimation for manipulation. *The International Journal of Robotics Research*, 30(10):1284–1306, 2011.
- [6] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2758–2766, 2015.
- [7] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, pages 1406–1415, 2018.
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [9] 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 *Asian conference on computer vision*, pages 548–562. Springer, 2012.
- [10] Tomáš Hodan, Pavel Haluza, Štěpán Obdržálek, Jiri Matas, Manolis Lourakis, and Xenophon Zabulis. T-less: An rgb-d dataset for 6d pose estimation of texture-less objects. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 880–888. IEEE, 2017.
- [11] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6d object pose estimation. In *European Conference on Computer Vision*, pages 606–619. Springer, 2016.
- [12] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1521–1529, 2017.
- [13] Alexander Krull, Eric Brachmann, Sebastian Nowozin, Frank Michel, Jamie Shotton, and Carsten Rother. Poseagent: Budget-constrained 6d object pose estimation via reinforcement learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6702–6710, 2017.
- [14] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Eppn: An accurate o (n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009.
- [15] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–698, 2018.
- [16] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [17] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [18] Fabian Manhardt, Diego Martin Arroyo, Christian Rupprecht, Benjamin Busam, Tolga Birdal, Nassir Navab, and Federico Tombari. Explaining the ambiguity of object detection and 6d pose from visual data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6841–6850, 2019.
- [19] Fabian Manhardt, Wadim Kehl, and Adrien Gaidon. Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2069–2078, 2019.
- [20] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6d pose refinement in rgb. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 800–815, 2018.
- [21] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: a hands-on survey. *IEEE transactions on visualization and computer graphics*, 22(12):2633–2651, 2015.
- [22] Stefan Mathe, Aleksis Pirinen, and Cristian Sminchisescu. Reinforcement learning for visual object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2894–2902, 2016.
- [23] David Nistér. Preemptive ransac for live structure and motion estimation. *Machine Vision and Applications*, 16(5):321–329, 2005.
- [24] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018. IEEE, 2017.
- [25] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019.

- [26] Arul Selvam Periyasamy, Max Schwarz, and Sven Behnke. Refining 6d object pose predictions using abstract render-and-compare. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2019.
- [27] Martin L. Puterman. *Markov Decision Processes.: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.
- [28] Deirdre Quillen, Eric Jang, Ofir Nachum, Chelsea Finn, Julian Ibarz, and Sergey Levine. Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6284–6291. IEEE, 2018.
- [29] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3836, 2017.
- [30] Liangliang Ren, Xin Yuan, Jiwen Lu, Ming Yang, and Jie Zhou. Deep reinforcement learning with iterative shift for visual tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 684–700, 2018.
- [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [32] Juil Sock, Guillermo Garcia-Hernando, and Tae-Kyun Kim. Active 6d multi-object pose estimation in cluttered scenarios with deep reinforcement learning, 2019.
- [33] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018.
- [34] James Supancic III and Deva Ramanan. Tracking as online decision-making: Learning a policy from streaming videos with reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 322–331, 2017.
- [35] Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. 2011.
- [36] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 292–301, 2018.
- [37] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019.
- [38] 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 (RSS)*, 2018.
- [39] Sangdoon Yun, Jongwon Choi, Youngjoon Yoo, Kimin Yun, and Jin Young Choi. Action-decision networks for visual tracking with deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2711–2720, 2017.
- [40] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1941–1950, 2019.
- [41] Xingyi Zhou, Arjun Karapur, Linjie Luo, and Qixing Huang. Starmap for category-agnostic keypoint and viewpoint estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018.
- [42] Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, and Wenjun Zeng. Context-reinforced semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4046–4055, 2019.
- [43] 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 *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3936–3943. IEEE, 2014.