

Active 6D Multi-Object Pose Estimation in Cluttered Scenarios with Deep Reinforcement Learning

Juil Sock¹, Guillermo Garcia-Hernando¹ and Tae-Kyun Kim^{1,2}

Abstract—In this work, we explore how a strategic selection of camera movements can facilitate the task of 6D multi-object pose estimation in cluttered scenarios while respecting real-world constraints such as time and distance travelled, important in robotics and augmented reality applications. In the proposed framework, multiple object hypotheses inferred by an object pose estimator are accumulated both in space and time with a fusion function. At each time step, this fusion function makes use of a verification score to quantify the quality of the hypotheses in the absence of ground-truth annotations and passes this information to an agent. The agent reasons about these hypotheses, directing its attention to the object which it is most uncertain about, moving the camera towards such an object. Unlike previous works that propose short-sighted policies, our agent is trained in simulated scenarios using reinforcement learning, attempting to learn the camera moves that produce the most accurate object poses hypotheses for a given temporal and spatial budget, without the need of viewpoints rendering during inference. Our experiments show that the proposed approach successfully estimates the 6D object pose of a stack of objects in both challenging cluttered synthetic and real scenarios, showing superior performance compared to other baselines.

I. INTRODUCTION

Accurate 6D object pose estimation of multiple objects on a cluttered scenario may become essential in applications such as robotic manipulation and augmented reality which, currently lacking proper solutions, resort to either coarse approximations [1] or less interpretable solutions [2]. Such applications and scenarios are naturally framed in an active setting, where either a robot or a human has the capability of moving the camera to different viewpoints as presented in e.g. the Amazon Picking Challenge [3], [4]. The present work explores how moving the camera in the scene, reaching different viewpoints, can help to overcome the inherent challenges in 6D object pose estimation, such as clutter and occlusion, identifying two ends of the spectrum of such problem. On one end we can decide not moving the camera at all, reducing the problem to single-view object pose estimation and thus subject to the object pose estimator limitations [5], [6], [7], [8]. On the other end, the camera can be moved to cover all the possible viewpoints [4], however, this does not respect real-world constraints, such as time limitation in terms of a number of movements and energy expenditure in terms of distance traveled. We are interested in finding a compromised solution that moves the camera strategically

by reaching the most informative viewpoints given a limited number of camera movement and keeping the traveled distance as low as possible. Previous work in active 6D object pose estimation includes the work of Doumanoglou *et al.* [9], which explored the problem of estimating the next best camera viewpoint by aiming to reduce the uncertainty of the pose estimator in terms of entropy reduction. Sock *et al.* [10] take a similar entropy reduction approach, but instead of using pose inference, they propose a heuristic geometric approach to estimate view entropy. Both [9] and [10] present two severe limitations. First, both frameworks are short-sighted, meaning that they move the camera to the next best view independently of the past and future viewpoints and potentially reaching a redundant view. Second, they require rendering different views before making a camera movement, which is an important limitation on real-world applications where the number of possible viewpoints is high. This work proposes a framework to tackle the active 6D object pose estimation problem that overcomes the limitations of the previous work and it is depicted in Fig. 1. Given a set of objects hypotheses inferred by an object pose estimator, our framework first fuses the information taking into account the previous estimations and chooses the hypotheses that best explain the scene without needing ground-truth information with the use of a verification score. The result of this fusion function is given to an agent, which analyzes all the hypotheses. Using an attention mechanism, the agent directs its attention to the object which it is most uncertain about and moving the camera towards such an object. The agent is trained in simulated scenarios on a reinforcement learning framework, attempting to learn a policy that moves the camera producing the most accurate object poses hypotheses for a given temporal and spatial budget. The sequential nature of the reinforcement learning framework gives the agent the ability to reason temporally and make long term decisions based on previous movements. Furthermore, our careful state space design allows the agent to perform inference without the need for rendering the entire amount of viewpoints at each time step. We evaluate our framework on challenging synthetic and real scenarios, showing that our framework can produce camera movements that achieve robust 6D object pose estimation over other baselines. In summary, this paper contains the following main contributions:

- 1) An active 6D multi-object pose estimation within a reinforcement learning framework is proposed. State and action spaces are carefully designed and a reward function is tailored to the problem of interest.
- 2) Application specific modules are proposed: a hypothe-

¹ Department of Electrical and Electronic Engineering, Imperial College London, United Kingdom. ² School of Computing, KAIST, South Korea.
Corresponding author: ju-il.sock08@imperial.ac.uk
This work was part of Imperial College London-Samsung Research project, supported by Samsung Electronics.

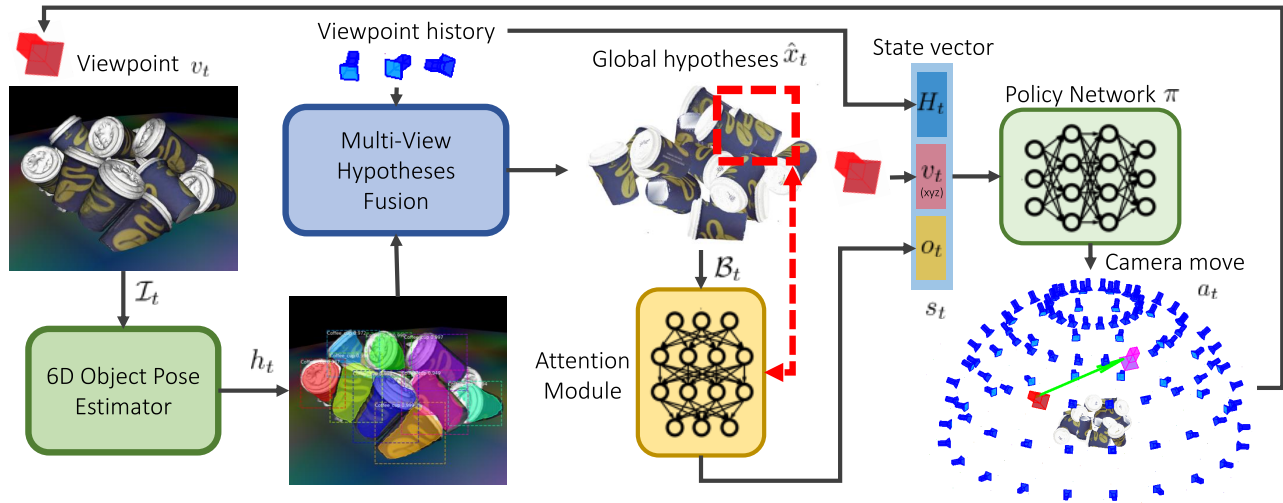


Fig. 1. An image from viewpoint v_t is passed through a 6D object pose estimator which generates pose hypotheses. Current and previous hypotheses are accumulated using a fusion function. The agent analyzes the scene, focusing its attention to the object which it is most uncertain about and moving the camera towards such object.

ses fusion function that works without ground-truth annotations by a proposed verification score and an object-specific attention module that directs the attention of the agent.

- 3) Extensive experimental evaluation of both the proposed method and baselines on both synthetic and real cluttered scenarios.

II. RELATED WORK

6D object pose estimation. 6D object pose estimation from a single image has been extensively researched in the past decades [11], [12], [13], [14]. Accurate pose estimation can be obtained under moderate occlusion and clutter with handcrafted features [15], [5]. To handle occlusion and truncation, keypoint detection with PnP [16], [17] or per-pixel regression/patch-based approaches [6], [9], [7] followed by Hough voting [18], [7] or RANSAC have been proposed. Most recent approaches attempt to jointly learn features and pose estimation using neural networks [12] in either RGB [19], [17], [20] or RGB-D images [1], [21]. However, many of the methods are designed and evaluated for Single instance of single object [8] and pose estimation of multiple instances of single object in cluttered scenarios remains a challenge. Estimating object poses on such scenarios using the above methods is challenging due to cascading errors such as severe occlusions and foreground clutter. Sock *et al.* [1] presented a work specifically for such scenarios by generating training dataset with occlusion patterns. Sundermeyer *et al.* [22] learned an implicit representation of object orientations defined by samples in a latent space to handle occlusions. However, the accuracy of a pose hypothesis is fundamentally limited by the visibility of the object in the input image. We end this section by reviewing methods that use multiple view information [23], [24], [25], [4], [26], [27]. However, these methods assume a set of images captured from pre-determined viewpoints are available. Our framework is agnostic to the choice of object pose estimator

and we propose a fusion function to incorporate multi-view information.

Active vision, object detection and poses. We review recent work on the active vision that aims to either improve the detection of objects, their pose estimation or both. Several methods have been proposed to select glimpse on static image to accelerate object detection [28], [29], [30] and more recently view selection for a moving observers such as visual navigation [31], [32], [33] and classification [34]. In [35] an active vision dataset and a reinforcement learning based baseline to explore the environment to detect objects are proposed. [36] presents a method to jointly learn a policy for both grasping and viewing. The method uses a simulated environment which detects objects depending on the occlusion rate. Unlike our framework where multiple objects are present, both methods include one object per scene. [37] shows a geometry-aware neural networks which integrate different views to a latent feature tensor which is also used to select views for the purpose of object reconstruction and segmentation. Similarly [38] uses reinforcement learning framework to select views to reconstruct 3D volume from RGB images of single object. More related to our work, [9] proposes a Hough forest approach and an entropy-based viewpoint selection. An important drawback of this method is that it does not have a mechanism to resolve conflicting next best view selection in the presence of multiple hypotheses. [10] integrates different components to build a complete active system which detects and pose estimates multiple objects. However, both frameworks decide the next best view independently of the past and future decision and also requires rendering of different viewpoints.

Attention to objects. Recent work has shown that employing attention mechanism is useful in the presence of multiple objects, distractors, and clutter. [39] used attention to select discriminative features for the purpose of view selection for object classification given a single detected object, whereas we use attention module to select a hypothesis from a set of

detected instances. [40] uses a pixel-wise attention module for the self-driving system to build more interpretable agent. [41] presents a system for object grasping with object-level attention where an attention mechanism is used to select the object to be manipulated on specific tasks. Similarly, in the proposed framework we propose an attention mechanism that makes the agent focus on the most uncertain object.

III. PROPOSED FRAMEWORK

Problem definition. Given a target stack of n objects, our aim is to infer the 6D pose $x^i \in \mathbb{R}^6$ of every i visible object in the stack, consisting of its 3D location and 3D orientation in the scene. In other words, our objective is to find the set of object hypotheses $\hat{x} = \{\hat{x}^1, \dots, \hat{x}^N\}$ that best describes the scene. We formulate the problem in an active setting, where an agent has the ability to navigate through a finite set of viewpoints $\mathcal{V} = \{v^1, \dots, v^M\}$, where $v^i \in \mathbb{R}^6$ is the 6D camera pose, and $M = |\mathcal{V}|$ is the number of viewpoints accessible to the agent. At each time step t , the agent observes the state of the scene s_t from a viewpoint $v_t \in \mathcal{V}$ and decides an action a_t following a policy function π that moves the agent to next viewpoint v_{t+1} . The agents proceeds until it reaches the maximum episode length T generating a trajectory $\tau = (s_0, a_0, \dots, s_{T-1}, a_{T-1})$. Moving an agent in the real world is costly in terms of time T and energy consumption, which we indirectly measure as a function of total traveled distance d . Both d and T are budget-constrained and low values are desired.

Framework overview. In our framework, presented in Fig. 1, at each time step t an input image \mathcal{I}_t is acquired from a viewpoint v_t and it is passed through a 6D object pose estimator that provides pose hypotheses h_t on the currently observed objects. These hypotheses are accumulated using a fusion function that considers both the current hypotheses and the previously observed ones and a selection of such is given to the agent. In a cluttered scenario, there are multiple objects with different challenges such as occlusion or measuring pose estimation confidence. This indicates there may not be a single next view ideal for all the objects and thus the agent reasons about the scene and focus its attention to the object which it is most uncertain about and moves the camera towards such object. This process is repeated until the maximum fixed number of time steps T and then the 6D pose hypotheses for the scene x is obtained. In the following sections, we describe the different components of our framework.

A. 6D object pose estimation

Given an image \mathcal{I}_t acquired from viewpoint v_t , the 6D object pose estimator outputs a set of object hypotheses for the current view $h_t = \{h_t^1, \dots, h_t^k\}$, where $h_t^i \in \mathbb{R}^6$ and k might differ from the actual number of objects in the scene N . Our framework is agnostic to the 6D object pose estimator of use and in this work, we opted to regress pixel-wise object-coordinate maps in a similar way to [21]. We use a 6D object pose estimator working on the RGB-D domain which, in addition to object hypotheses, also provides 2D bounding boxes and segmentation masks for

each detected object, which will be used on different stages of our framework detailed below.

B. Multi-view hypotheses fusion

At each viewpoint, v_t , the agent obtains a different set of object hypotheses h_t from the 6D object pose estimator. To obtain a global scene hypothesis \hat{x} that accumulates the information obtained from different viewpoints, we define a hypotheses fusion function f . Given a history of object hypotheses h_1, \dots, h_t and camera viewpoints up to time step t , the function f outputs the accumulated hypotheses \hat{x}_t in the global coordinate system. Intuitively, f gathers all the object hypotheses that have been observed so far, selecting the ones that better explain the scene among all the pool of hypotheses, hence $\hat{x}_t \subseteq \bigcup_{i=1}^t \bigcup_j h_i^j$. In our framework, f makes use of hierarchical clustering to group hypotheses and the hypothesis with the highest object pose confidence score is selected from each cluster in a similar way to [10]. In our framework, the hypothesis with the highest verification score from each cluster is chosen, as we use verification score to represent quality of the hypothesis.

Verification score. The hypotheses fusion function requires quantifying the quality of a hypothesis in the absence of ground-truth information. For this, we introduce a verification score that estimates hypothesis confidence inspired by the work of [42] and [9]. These works assume that a scene point q that belongs to an object -is an *inlier*- if the distance to the nearest neighboring point p of a rendered 3D model of a hypothesis is lower than a certain threshold ϵ . This assumption is not generally true in the presence of multiple objects in clutter. To overcome this limitation, we use the segmentation mask from the object pose estimator and consider points on the depth map that lie inside the mask as inliers. The local fitting between p and q is measured by $\delta(p, q)$ and is defined as follows:

$$\delta(p, q) = \begin{cases} \frac{1}{2}(1 - \frac{\|p-q\|_2}{\epsilon}) + \frac{1}{2}(n_p \cdot n_q), & \text{if } \|p - q\|_2 < \epsilon \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where n_p and n_q denote the normal vectors at p and q respectively. The verification score c is calculated as the mean value of $\delta(p, q)$ for all the inlier points. In Fig. 2, a graph of verification score against the e_{ADD} for coffee cup is presented showing how the proposed verification score correlates with e_{ADD} , making it suitable to be used in the absence of ground-truth annotations. For coffee cup object, 10% of the diameter of the object is 14mm which is used to decide whether the object pose hypothesis is of an acceptable quality.

Objects feature representation. When encoding object features, 6D object pose hypothesis information is not directly used given that some estimated object poses are of poor quality to provide useful enough information. Instead, we propose to use a more reliable and simpler feature representation to encode object hypotheses. The output of the hypotheses selector is a set of object features \mathcal{B}_t for k object hypotheses. Each element in \mathcal{B}_t consists of a tuple (b^i, d^i, c^i) , where b^i is the normalized 4 dimensional object 2D bounding

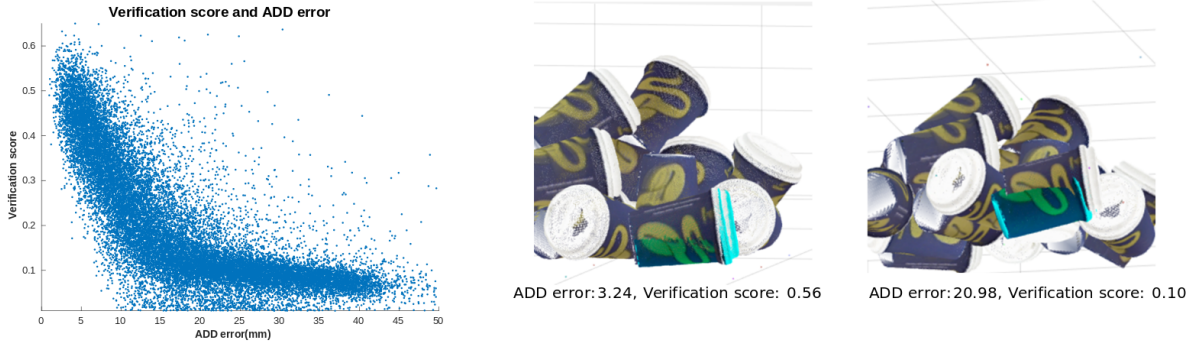


Fig. 2. (Left) A graph showing a relation between verification score and e_{ADD} . (Middle and right) The figures show the scene and hypotheses rendered in blue.

box coordinates, d^i is the distance of the hypothesis from the viewpoint in depth values on the camera coordinates and c^i is the verification score of the hypothesis. Bounding box coordinates are directly used for b^i and d^i is approximated by averaging the mean value of depth map values corresponding to the segmentation mask of the 6D object pose estimator. The value from the object hypothesis in depth axis is not directly used as it is often incorrect notably when the object is highly occluded.

C. Object attention module

At a given time step t , the agent observes a number of object hypotheses and makes a choice on which next viewpoint to visit at $t + 1$. To guide the agent on that decision, we propose an attention mechanism that makes the agent to focus on the most uncertain object and move accordingly. In this section, we describe how we design this attention mechanism and in Sec. III-D we provide details on its learning function.

Our attention mechanism receives as input the set of object features \mathcal{B}_t from the previous module and returns an individual object o_t feature of the object to reason about and its index m_t in a similar way to [41] and [43]. First, a representation of each object in \mathcal{B}_t is extracted using a fully connected network. These individual features are aggregated using a mean-pool layer leading to a global object representation that is concatenated to the individual feature. A Selector network takes these concatenated features and outputs a scalar score indicating the importance of the object which is normalized using a softmax layer. The object with the highest score is then selected and its features are given to the policy network. Note that gradients flow from the policy network, guiding the attention module to the object of interest in our active setting. A description of the algorithm can be found in Alg. 1. Selector(\cdot) is a linear mapping from object features to attention score which is optimized jointly with the policy.

D. Policy learning

Formulation of reinforcement learning problem. We formulate the decision process of choosing an action a_t leading to a viewpoint v_{t+1} within a reinforcement learning (RL) framework. The next viewpoint is strategically chosen to lead to an optimal set of hypotheses \hat{x}^* that best describes

Algorithm 1: Attention module

Data: \mathcal{B}_t features of k objects at time step t
Result: o_t object feature and object index m_t for the policy network

```

for  $b^i \in \mathcal{B}_t$  do
     $g^i := \text{FC}(b^i, d^i, c^i)$  // fully connected
    network
end
 $g^{\text{global}} := \text{meanpool}(g^1, \dots, g^k)$ 
for  $b^i \in \mathcal{B}_t$  do
     $w^i := \text{Selector}(\text{concat}(g^i, g^{\text{global}}))$  // object
    score
end
 $\bar{w}^1, \dots, \bar{w}^k = \text{softmax}(w^1, \dots, w^k)$ 
 $m_t = \text{argmax}(\bar{w}^1, \dots, \bar{w}^k)$  // attended object
    index
 $o_t = \text{concat}(g^m * \bar{w}^m, b^m, d^m, c^m)$ 
    // individual object feature

```

the scene. The optimization problem is defined as the maximization of a parametric function $J(\theta)$ defined as $J(\theta) = \mathbb{E}_{\tau \sim p_\theta(\theta)} \left[\sum_{t=0}^T \gamma^t r_t \right]$, where $p_\theta(\theta)$ is the distribution over all possible trajectories following the parametric policy π_θ . The term $\sum_{t=0}^T \gamma^t r_t$ represents the total return of a trajectory for a horizon of T time steps and a discount factor $\gamma \in [0, 1]$. **Reward function.** Our reward function involves three different terms: (i) individual pose estimation accuracy; (ii) change in distance from camera to the attended object and (iii) penalization for long distance camera movements. The first term is defined as follows:

$$r_t^{e_{ADD}} = e_{ADD}(\hat{x}_{t+1}^m, x_{t+1}^m) - e_{ADD}(\hat{x}_t^m, x_t^m), \quad (2)$$

where index m_t denotes the object index selected by attention network described in sec. III-C at time t . e_{ADD} is the most widely used 6D object pose error function in the literature and is the average Euclidean distance of model points proposed by Hinterstoisser *et al.* [44]. It is defined as the average Euclidean distance of the estimated pose \hat{x} with respect to the ground-truth pose x of an object model.

We hypothesize that a viewpoint closer to the object of interest is more likely to give better pose estimation. It can be used to guide and accelerate the policy learning and this

leads to the second term defined as:

$$r_t^{dist} = \|\hat{x}_{t+1}^m[xyz] - v_{t+1}[xyz]\|_2 - \|\hat{x}_t^m[xyz] - v_t[xyz]\|_2, \quad (3)$$

where $[xyz]$ is the translation vector from a 6D pose. The full reward function is the combination of the two above and another term penalizing long distance camera motions aiming to minimize the total distance traveled d :

$$r_t = (1 - \alpha)r_t^{eADD} + \alpha\beta r_t^{dist} - (1 - \alpha)(1 - \beta)(\|v_{t+1}[xyz] - v_t[xyz]\|_1). \quad (4)$$

where α and β are hyperparameters in $[0, 1]$ weighting the influence of the different reward terms.

State space. The state vector s_t in our framework consists of three concatenated components. First, the object feature o_t that comes from the attention module presented in Algorithm 1. Second, the current camera pose position $v_t[xyz]$. Last, a history vector H_t that encodes previous camera positions for T steps, which is initialized to zeros and it is filled at each time step. This history vector is needed in the absence of a memory module to respect the Markov property on a reinforcement learning framework.

Action space. We define the action space as a two-dimensional continuous space spanning azimuth and elevation angles the camera viewpoint of a hemisphere centered in the stack of objects. More specifically, an action at time step t is defined as $a_t = \{\phi_t^a, \phi_t^e\}$ where ϕ^a and ϕ^e represent azimuth and elevation angles of the camera respectively. Given the finite nature of the viewpoint space, the angles from a_t are mapped to the closest viewpoint in \mathcal{V} .

Policy network We represent the policy function π_θ as a neural network with parameters θ that outputs the mean and standard deviation of a Gaussian distribution. To optimize θ different methods can be used, however, in our framework, we use policy gradients method [45]. These methods optimize $J(\theta)$, where the gradient of the expected return $\nabla_\theta J(\theta)$ is estimated with trajectories sampled by following the policy.

IV. EXPERIMENTAL RESULTS

Dataset and experimental setup. For synthetic experiments, Bin-picking dataset with two different object models are used for training and testing: Coffee Cup from [9] dataset and Bunny model from [46] dataset. For each model, we generate 40 scenes with different number of objects and random pose configurations. Objects with low visibility not detected in any viewpoints are removed from the ground-truth and are not included in the evaluation. For each object, 30 scenarios are used for training and 10 scenarios are used for testing. Each scenario has 100 views on the upper hemisphere with radius of 80 cm. The view grid consists of 5 elevation and 20 azimuth levels and both RGB and depth images are rendered with OpenGL. The number of instances varies randomly between 15-20 for Coffee Cup and 7-12 for Bunny due to the size and shape difference. For e_{ADD} , all detectable objects are taken into account, and the undetected object are assigned the error of 50mm which is more than 3 times the acceptable threshold.

TABLE I

DETECTION RATE, $e_{ADD/I}$ AND DISTANCE TRAVELLED EVALUATED WITH THE DIFFERENT BASELINES ON SYNTHETIC ENVIRONMENT. MAXIMUM EPISODE LENGTH SET TO 5.

Policy	Coffee cup			Bunny		
	Distance $d \downarrow$	e_{ADI} (mm) \downarrow	Detection rate \uparrow	Distance $d \downarrow$	e_{ADD} (mm) \downarrow	Detection rate \uparrow
Random	4.97	13.48	0.76	3.97	26.69	0.19
Maximum Distance	7.66	14.02	0.75	7.68	26.83	0.46
Unidirectional	2.88	14.32	0.74	2.88	27.12	0.44
<i>Proposed</i>	3.71	11.12	0.80	2.43	25.35	0.46

We use the coffee cup scenario from the dataset of Doumanoglou *et al.* [9] for the real experiment. Since the viewpoints of the dataset are not evenly distributed, a view grid is constructed in the same ways as the synthetic environment and the image from viewpoint closest to the grid viewpoint is used as the observation. During training and testing, 2D bounding box coordinate is rotated around the depth axis to correct the in-plane rotation. The inference time for the pose estimator depends on the number of objects in the scene. In the case of 14 objects the pose estimator outputs hypotheses at 2 Hz. The rest of the pipeline including object accumulation, clustering and policy network operate at 25 Hz. It is possible to design a pipeline to leverage intermediate image data acquired while travelling from v_t to v_{t+1} , however it is less practical due to the processing time. The results are obtained with Intel i5-6400 CPU and GTX 1080Ti.

Our framework implementation is built in PyTorch and the learning algorithm, PPO [45].

Baseline policies. We test our active pose estimation system against a variety of baselines. The first baseline **“Random”** samples random actions from a fixed Gaussian distribution. The second baseline **“Unidirectional”** moves the camera in one horizontal direction around the scene. The trajectory is generated such that it completes one full revolution around the scene. Elevation of the viewpoint is fixed to be 45 degrees which provide a balanced view between occlusion and variability in each view. The third baseline, **“Maximum distance”**, selects viewpoints with the longest distance from previous viewpoints. This baseline test whether the most informative views are simply far-apart views. Lastly, an **“Entropy-based”** baseline which selects the next best view based on entropy [10], [9] is presented to be able to compare to a recent evaluation. For this baseline, the pipeline is identical to the proposed system except the action is generated based on view entropy, which is pre-computed for every view. Segmentation masks inferred from the pose estimator are used for the view entropy calculation. Sample trajectories of baselines are visualized in Fig. 3.

Evaluation metrics: Pose accuracy is evaluated with e_{ADD} for bunny and e_{ADI} for coffee cup due to symmetry. Following the standard practice [44], objects with $e_{ADD/I}$ error less than 10% of the object diameter is considered to be a correct hypothesis. All pose errors are measured in mm. The detection rate is defined as the ratio of number of objects with correctly estimated pose to the total number of objects




Fig. 3. (Left) Sample image of Maximum Distance policy baseline (Middle) Sample image of Random policy baseline (Right) Sample image of Unidirectional policy baseline




Fig. 4. Performances with different hyperparameters. (Left) Distance travelled and mean pose error with respect to α values. Change in mean pose error (Middle) and detection rate (Right) with different number of maximum episode length (T).




Fig. 5. (Left) An example of the learned trajectory for the coffee cup dataset. (Right) Every column represents one time step (Top and middle row) Image from the viewpoint before and after taking the action. (Bottom row) 3D visualization of the action for each time step.

in the bin in each scenario. Distance traveled is the total distance the camera moved where it is measured along the surface of the hemisphere as it is unrealistic for the sensor to move through the objects.

Synthetic dataset. Table. I shows the mean pose error and the correct detection rate on the synthetic dataset. To obtain the results, episode lengths T are held constant at 5 steps. For Coffee Cup object, our approach consistently outperforms all baselines in both $e_{ADD/I}$ and detection rate metric by a significant margin. In most cases, *Unidirectional* policy results in low distance d as the agent only moves horizontally. For Bunny object, the proposed method outperforms all baselines in all metrics with a narrower margin due to the smaller number of instances in the scene. Fig.5 shows how the agent behaves at each time step of the trajectory. The agent selects the object with the lowest confidence, which

is expressed as the verification score, but also not too far away to reduce the distance traveled. The attention shifts to different objects in different time step and it can be noted the number of the correctly detected object (highlighted in green) increases every step.

Real world dataset. Fig. 6 shows a sample trajectory generated by the learned policy network and quantitative comparison with other baselines. Unidirectional baseline is not included since the distribution of viewpoints in the real dataset is not in a grid form or uniform. Since most of the viewpoints are densely populated on the top part of the scene, the learned policy tends to select views with lower elevation resulting in a behavior similar to maximum distance baselines. The proposed method outperforms the other baselines including the entropy-based method [10].

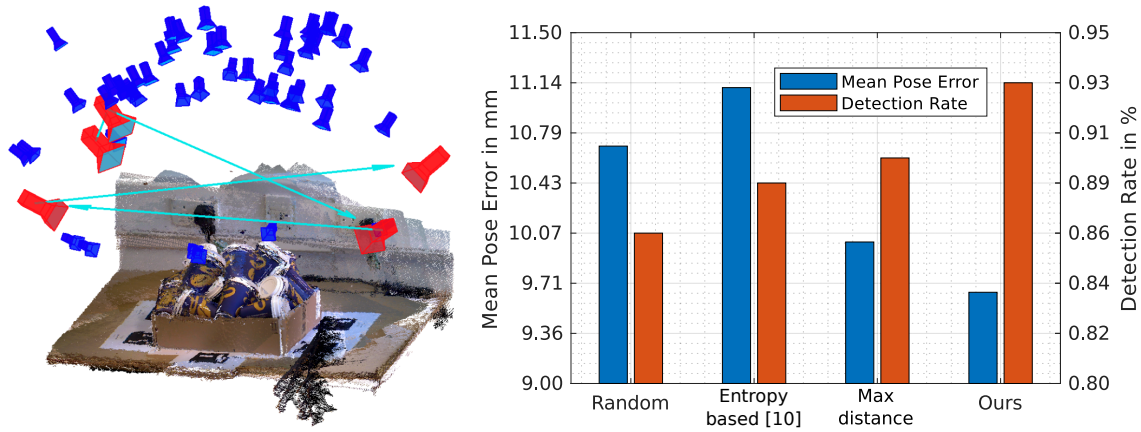


Fig. 6. Evaluation on real dataset. (Left) Visualization of a sample trajectory. Blue camera represents a set of allowable viewpoints and red cameras show the viewpoints selected by the proposed system. The arrow shows the order of view selection. (Right) Performance of different baselines on the real dataset.

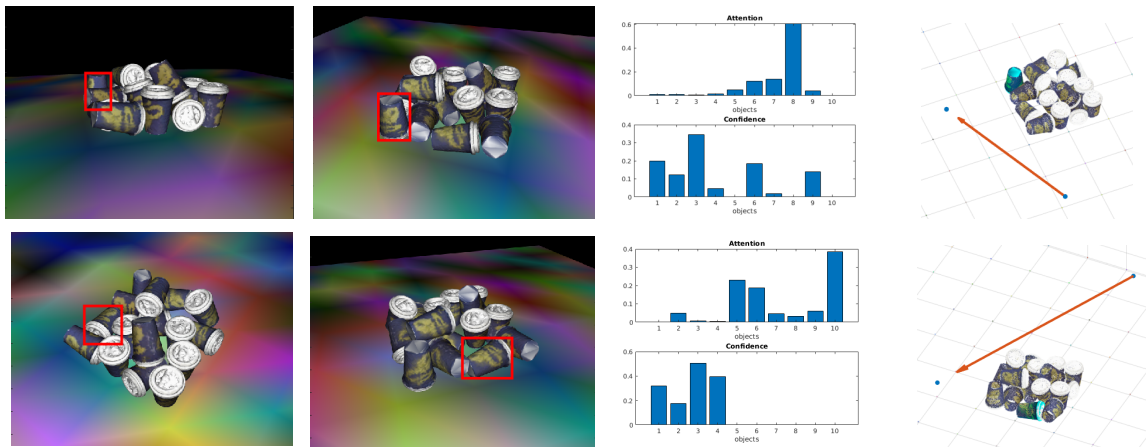


Fig. 7. Relationship between attention, verification score and action. (First column) Images of the scene before taking action. Attended object is highlighted in red box. (Second column) Images of the scene after taking action. (Third column) Bar graph showing the object score (\hat{w}^2 in the manuscript) and verification score (shown as confidence). (Fourth column) 3D visualization of camera trajectory.

A. Ablation experiments

α and T parameters. Experiment with different trajectory is presented to show how the performance changes within the spectrum between single-view pose estimation ($T = 0$) and all-view pose estimation ($T = 99$ in our case). Fig. 4 shows the performance indeed increases with more views but the benefit saturates after $T = 10$. For α , excessively emphasizing the r^{dist} term results in the agent to give attention to the furthest object and move the camera towards the object to maximize the reward, increasing the travel distance. However it is difficult for the agent to learn a policy without r^{dist} term since during training often there are multiple views around the object of interest which are equally favorable.

Attention module. To verify whether the attention module properly learns to automatically select object, a baseline which assigns attention to the object with the lowest verification score is tested. Compared to the verification score-based selection, our baseline improved the detection rate by 3% and mean pose error by 9%.

Qualitative evaluations. More qualitative evaluation results including the attention value for detected object for better insights are shown in Fig. 7. It can be noted the selector network tends to choose the object with the most uncertainty.

V. CONCLUSION

We presented a framework to deal with the active 6D object pose estimation in cluttered scenarios problem. We used recent reinforcement learning advances to formulate our framework and carefully designed all the different components to achieve superior performance compared to different baselines. We believe there is a margin for improvement in terms of performance and real-world framework evaluation due to the lack of proper datasets. As future work, we would like to explore the use of the proposed framework on a real robot, its impact on the graspability of objects, its generalizability to different scenarios/objects and the use hand-object pose priors to accelerate the policy training [47], [48]. Also, using a real robot could ease the real world dataset generation.

REFERENCES

- [1] J. Sock, K. I. Kim, C. Sahin, and T.-K. Kim, "Multi-task deep networks for depth-based 6d object pose and joint registration in crowd scenarios," in *BMVC*, 2018.
- [2] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *JMLR*, 2016.
- [3] C. Eppner, S. Höfer, R. Jonschkowski, R. Martin, A. Sieverling, V. Wall, and O. Brock, "Lessons from the amazon picking challenge: Four aspects of building robotic systems," in *RSS*, 2016.
- [4] A. Zeng, K.-T. Yu, S. Song, D. Suo, E. Walker, A. Rodriguez, and J. Xiao, "Multi-view self-supervised deep learning for 6D pose estimation in the amazon picking challenge," in *ICRA*, 2017.
- [5] S. Hinterstoisser, S. Holzer, C. Cagniard, S. Ilic, K. Konolige, N. Navab, and V. Lepetit, "Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes," in *ICCV*, 2011.
- [6] E. Brachmann, A. Krull, F. Michel, S. Gumhold, J. Shotton, and C. Rother, "Learning 6d object pose estimation using 3D object coordinates," in *ECCV*, 2014.
- [7] A. Tejani, D. Tang, R. Kouskouridas, and T.-K. Kim, "Latent-class hough forests for 3D object detection and pose estimation," in *ECCV*, 2014.
- [8] T. Hodan, F. Michel, E. Brachmann, W. Kehl, A. GlentBuch, D. Kraft, B. Drost, J. Vidal, S. Ihrke, X. Zabulis *et al.*, "Bop: benchmark for 6d object pose estimation," in *ECCV*, 2018.
- [9] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6D object pose and predicting next-best-view in the crowd," in *CVPR*, 2016.
- [10] J. Sock, S. Hamidreza Kasaei, L. Seabra Lopes, and T.-K. Kim, "Multi-view 6d object pose estimation and camera motion planning using rgbd images," in *ICCVW*, 2017.
- [11] C. Sahin, G. Garcia-Hernando, J. Sock, and T.-K. Kim, "Instance- and category-level 6d object pose estimation," *arXiv preprint arXiv:1903.04229*, 2019.
- [12] V. Balntas, A. Doumanoglou, C. Sahin, J. Sock, R. Kouskouridas, and T.-K. Kim, "Pose guided rgbd feature learning for 3d object pose estimation," in *ICCV*, 2017.
- [13] C. Sahin, G. Garcia-Hernando, J. Sock, and T.-K. Kim, "A review on object pose recovery: from 3d bounding box detectors to full 6d pose estimators," *Image and Vision Computing*, 2020.
- [14] P. Castro, A. Armagan, and T.-K. Kim, "Accurate 6d object pose estimation by pose conditioned mesh reconstruction," in *ICASSP*, 2020.
- [15] B. Drost, M. Ulrich, N. Navab, and S. Ilic, "Model globally, match locally: Efficient and robust 3D object recognition," in *CVPR*, 2010.
- [16] M. Rad and V. Lepetit, "BB8: A scalable, accurate, robust to partial occlusion method for predicting the 3D poses of challenging objects without using depth," in *ICCV*, 2017.
- [17] B. Tekin, S. Sinha, and P. Fua, "Real-time seamless single shot 6D object pose prediction," *arXiv preprint arXiv:1711.08848*, 2017.
- [18] S. Peng, Y. Liu, Q. Huang, H. Bao, and X. Zhou, "Pvnet: Pixel-wise voting network for 6dof pose estimation," in *CVPR*, 2019.
- [19] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "SSD-6D: Making rgb-based 3D detection and 6D pose estimation great again," in *ICCV*, 2017.
- [20] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," in *Robotics: Science and Systems (RSS)*, 2018.
- [21] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," *CVPR*, 2019.
- [22] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in *ECCV*, 2018.
- [23] F. Viksten, R. Soderberg, K. Nordberg, and C. Perwass, "Increasing pose estimation performance using multi-cue integration," in *ICRA*, 2006.
- [24] A. Collet and S. S. Srinivasa, "Efficient multi-view object recognition and full pose estimation," in *ICRA*, 2010.
- [25] Ö. Erkent, D. Shukla, and J. Piater, "Integration of probabilistic pose estimates from multiple views," in *ECCV*, 2016.
- [26] A. Kanazaki, Y. Matsushita, and Y. Nishida, "Rotationnet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints," in *CVPR*, 2018.
- [27] C. Li, J. Bai, and G. D. Hager, "A unified framework for multi-view multi-class object pose estimation," in *ECCV*, 2018.
- [28] S. Mathe, A. Pirinen, and C. Sminchisescu, "Reinforcement learning for visual object detection," in *CVPR*, 2016.
- [29] A. Gonzalez-Garcia, A. Vezhnevets, and V. Ferrari, "An active search strategy for efficient object class detection," in *CVPR*, 2015.
- [30] J. C. Caicedo and S. Lazebnik, "Active object localization with deep reinforcement learning," in *CVPR*, 2015.
- [31] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *ICRA*, 2017.
- [32] P. Mirowski *et al.*, "Learning to navigate in complex environments," *arXiv*, 2016.
- [33] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," in *CVPR*, 2017.
- [34] B. Calli, W. Caarls, M. Wisse, and P. P. Jonker, "Active vision via extremum seeking for robots in unstructured environments: Applications in object recognition and manipulation," *IEEE Transactions on Automation Science and Engineering*, 2018.
- [35] P. Ammirato, P. Poirson, E. Park, J. Košecká, and A. C. Berg, "A dataset for developing and benchmarking active vision," in *ICRA*, 2017.
- [36] R. Cheng, A. Agarwal, and K. Fragkiadaki, "Reinforcement learning of active vision formantipulating objects under occlusions," *arXiv preprint arXiv:1811.08067*, 2018.
- [37] R. Cheng, Z. Wang, and K. Fragkiadaki, "Geometry-aware recurrent neural networks for active visual recognition," in *NIPS*, 2018.
- [38] X. Yang, Y. Wang, Y. Wang, B. Yin, Q. Zhang, X. Wei, and H. Fu, "Active object reconstruction using a guided view planner," *arXiv preprint arXiv:1805.03081*, 2018.
- [39] K. Xu, Y. Shi, L. Zheng, J. Zhang, M. Liu, H. Huang, H. Su, D. Cohen-Or, and B. Chen, "3d attention-driven depth acquisition for object identification," *TOG*, 2016.
- [40] J. Kim and J. Canny, "Interpretable learning for self-driving cars by visualizing causal attention," in *ICCV*, 2017.
- [41] C. Devin, P. Abbeel, T. Darrell, and S. Levine, "Deep object-centric representations for generalizable robot learning," in *ICRA*, 2018.
- [42] A. Aldoma, F. Tombari, L. Di Stefano, and M. Vincze, "A global hypotheses verification method for 3d object recognition," in *ECCV*, 2012.
- [43] D. Wang, C. Devin, Q.-Z. Cai, F. Yu, and T. Darrell, "Deep object centric policies for autonomous driving," *arXiv preprint arXiv:1811.05432*, 2018.
- [44] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes," in *ACCV*, 2012.
- [45] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [46] R. Brégier, F. Devernay, L. Leyrit, and J. Crowley, "Symmetry aware evaluation of 3D object detection and pose estimation in scenes of many parts in bulk," in *ICCVW*, 2017.
- [47] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim, "First-person hand action benchmark with rgb-d videos and 3d hand pose annotations," in *CVPR*, 2018.
- [48] D. Antotsiou, G. Garcia-Hernando, and T.-K. Kim, "Task-oriented hand motion retargeting for dexterous manipulation imitation," in *ECCVW*, 2018.