SynPick: A Dataset for Dynamic Bin Picking Scene Understanding

Arul Selvam Periyasamy*, Max Schwarz*, and Sven Behnke

Abstract-We present SynPick, a synthetic dataset for dynamic scene understanding in bin-picking scenarios. In contrast to existing datasets, our dataset is both situated in a realistic industrial application domain-inspired by the well-known Amazon Robotics Challenge (ARC)-and features dynamic scenes with authentic picking actions as chosen by our picking heuristic developed for the ARC 2017. The dataset is compatible with the popular BOP dataset format. We describe the dataset generation process in detail, including object arrangement generation and manipulation simulation using the NVIDIA PhysX physics engine. To cover a large action space, we perform untargeted and targeted picking actions, as well as random moving actions. To establish a baseline for object perception, a state-of-the-art pose estimation approach is evaluated on the dataset. We demonstrate the usefulness of tracking poses during manipulation instead of single-shot estimation even with a naive filtering approach. The generator source code and dataset are publicly available.

I. INTRODUCTION

6D pose estimation is an important and effective perceptual tool in many robotic applications, such as grasping (both in industrial as well as service robotics contexts), state estimation, and prediction. It explains scene geometry using few parameters.

Learning 6D pose estimation is tricky, however. There are issues with the problem definition itself, such as how to deal with object symmetries and other ambiguities, or with the chosen parametrization of the SE(3) group. A more practical issue is how to gather training data for pose estimation. Whereas manual annotations for classification, object detection, and even semantic segmentation can be done in reasonable time, pose annotations are both time-consuming and prone to annotation errors.

To address these issues, it has become common practice to augment smaller-scale real datasets with larger synthetic datasets [1]. We present such a synthetic dataset specifically focused on bin picking. Our immediate inspiration is the setting of the Amazon Robotics Challenge 2017, which required participants to do both *targeted* picking of desired objects, and *untargeted* emptying of totes.

In contrast to existing datasets, our SynPick dataset does not only have static scenes, but consists of fully dynamic picking sequences (see Fig. 1). It is our hope that having such a dataset (and the accompanying generator) will enable the bin picking community to advance from analysis of static scenes to live object tracking during manipulation.



Fig. 1. SynPick contains typical dynamic bin picking sequences with pose & semantic segmentation annotations.

In short, our contributions include:

- 1) A dynamic scene generator, capable of producing realistic picking and moving sequences, both online and offline,
- 2) a larger-scale multi-view dataset produced with said generator and detailed analysis of its properties, and
- 3) a baseline evaluation of a state-of-the-art object pose estimation method and a filtering method for tracking on the dataset.

Both generator source code and the SynPick dataset itself are publicly available¹.

II. RELATED WORK

In recent years, with the advent of deep learning models for 6D object pose estimation, the datasets for training and benchmarking these models have also grown in size. The generation of large-scale datasets for 6D object pose estimation remains work-intensive, though. Unlike obtaining ground truth annotations for 2D computer vision tasks like object detection or object classification, annotating 6D ground truth poses is time-consuming, making the manual annotation of 6D poses on a large scale prohibitive. Thus, most of the datasets for benchmarking 6D object pose estimation rely on semi-automated pipelines. These pipelines involve capturing short video sequences of tabletop scenes, annotating the first frame of the sequence manually, and propagating the annotated poses to the rest of the frames by computing the camera trajectory using visual odometry techniques. During the training of deep learning methods for 6D pose estimation, real annotated training images are often supplemented with synthetic images. Synthetic image

¹http://ais.uni-bonn.de/datasets/synpick/

^{*} Equal contribution.

All authors are with the Autonomous Intelligent Systems group of University of Bonn, Germany; periyasa@ais.uni-bonn.de

 TABLE I

 6D Pose Estimation and Tracking Datasets

Name	Туре	Objects	#Frames	Annotation	Diverse Lighting	Dynamics	Multi-View
YCB-Video [2]	Real videos	21	133,827	Semi-auto	Yes	Static	Moving cam
Linemod-Occluded [3]	Real videos	8	1,214	Semi-auto	No	Static	Moving cam
TUD-L [1]	Real videos	3	23,914	Semi-auto	Yes	Static	Moving cam
TYO-L [1]	Real images	21	1,670	Manual	Yes	Static	No
HomebrewedDB [4]	Real videos	33	17,420	Semi-auto	No	Static	Moving cam
BlenderProc4BOP [5]	Synthetic images	flexible	50,000	Automated	Yes	Static	25 views
FAT [6]	Synthetic videos	21	61,500	Automated	Yes	Falling	Stereo
ObjectSynth [7]	Synthetic images	39	600,000	Automated	Yes	Static	200 views
SynPick (ours)	Synthetic videos	21	503232	Automated	Yes	Pick/Move	3 views

generation is done based on two approaches: "render & paste" and physics-based rendering (PBR). The "render & paste" technique involves rendering objects onto random backgrounds using standard rasterization. This technique is simple and fast but generates physically unrealistic, poor-quality training images. PBR is often implemented using ray tracing, which although compute-intensive and thus time-consuming, generates high-quality training images.

Existing datasets for 6D object pose estimation include YCB-Video [2], Linemod-Occluded [3], HomebrewedDB [4], TUD-Light (TUD-L) [1], Toyota Light (TYO-L) [1], and Rutgers APC [8]. See Table I for an overview. Each of the datasets focuses on specific real-world scenarios. Hodaň et al. [5] unified the existing datasets to a common BOP Dataset. Furthermore, they supplemented each dataset with PBR training images as a part of the BOP Challenge for benchmarking the progress in 6D object pose estimation research. One of the key findings from the BOP Challenge 2020 is that 50K PBR images yield better results than 1M "render & paste" images. This finding provides a compelling motivation to develop efficient pipelines for physics-based training image generation. Despite the advantages, the timeconsuming nature of the ray tracing methods limits their applicability. Schwarz and Behnke [9] introduced Stillleben, an efficient data generation pipeline that can generate PBR training images using OpenGL rasterization on the fly, completely eliminating the need for an offline data generation pipeline and demonstrated the advantages of synthetic data generation for semantic segmentation tasks.

An interesting orthogonal approach to supervised pose estimation is self-supervised 6D object pose estimation proposed by Deng *et al.* [10]. In this approach, a 6D pose estimation module trained in a supervised manner is used to initialize the 6D poses of objects in the scene. The robot changes the configuration of the objects in the scene by random pick and place actions. By capturing the images of objects in new configurations and propagating the initial pose estimate to these using forward kinematics, the authors generate new training data and refine the pose estimator actively. While this method provides a scalable approach to train pose estimators, it is limited by the simple object manipulation actions the robot can perform without breaking the object pose tracking module. Our proposed method does not suffer from this limitation and can model complex object interactions accurately. Furthermore, our data generator runs faster than real-time and can be easily parallelized.

Nearly all datasets geared towards 6D pose estimation feature only static scenes. While the camera often travels around the arrangement, object configurations remain fixed. In contrast, our dataset features dynamic scenes, suitable for training and evaluating not only pose estimation, but also pose tracking methods.

Most related to our work is the Falling Things Dataset (FAT) [6]. It consists of renderings of randomly sampled subsets of YCB objects placed in different indoor and outdoor scenes. Unlike most other datasets, FAT does not include tabletop scenes where the objects are nicely arranged without significant occlusions. Instead, the objects are dropped from a height onto the scenery, using a physics engine to model object interactions. However, both the context (kitchen/temple/forest) as well as the dynamics (objects falling on the background geometry) do not really fit a bin picking application. In comparison, our dataset features a standard picking tote and demonstrates pick and move dynamics.

III. DATASET GENERATION

We use Stillleben [9] as the base rendering engine for generating the SynPick dataset, as it has built-in physics simulation for table-top arrangements and is capable of online generation. However, Stillleben was developed for fast scene generation of different scenes. In this section, we describe the extensions to Stillleben we developed to adapt it for generation of continuous bin picking sequences and detail the dataset generation process.

A. Scene Generation

We rely on Stilleben's scene arrangement engine for physically realistic scene generation. We start with adding the tote object to the scene. Then, we randomly sample a set of objects such that the total volume of the objects does not exceed a threshold of 71, which generates scenes that contain many occlusions. We let the objects fall into the tote from a fixed height and simulate the effect of gravity and collision with other objects and the tote. Stilleben internally uses the PhysX engine by NVIDIA² for simulation. We use meshes provided by the YCB-Video dataset [2] for

²PhysX: https://developer.nvidia.com/physx-sdk



Fig. 3. Lighting variants. We show the same object arrangement in different sIBL light maps.

rendering. Simulating collisions with these high-resolution meshes would be prohibitively slow. For most objects, we compute convex hulls of low complexity. For three highly concave objects (cup, bowl, and power drill), we use the V-HACD [12] convex decomposition algorithm to find a number of smaller convex meshes describing the geometry accurately.

After the scene arrangement has been generated, we start the generation of training frames. During generation, we run PhysX simulation with a very small step size of 2 ms to ensure realistic behavior. With a rate of 15 Hz (relative to simulated time), we generate output frames from three different camera poses above the tote. The simulated cameras have Full HD resolution (1920×1080) and are mounted 2 m above the tote, similar to our ARC 2017 system [11].

For RGB rendering, the standard Stillleben pipeline is used. A physics-based rendering shader pipeline combined with randomly selected image-based lighting (IBL) maps from the sIBL archive³ ensure interesting and realistic lighting effects (see Fig. 3).

In the following sections, we describe the simulation of pick and move actions in detail.

B. Untargeted Picking

In the pick task simulation, we simulate the robot emptying the tote by picking the objects out of the tote, similar to

³http://www.hdrlabs.com/sibl/archive.html



Fig. 4. Articulated gripper. The cone-shaped finger ends in an actuated tiltable suction cup (joint axis marked in red).

the *Stow* task of the Amazon Robotics Challenge 2017. Starting with the tote filled with objects in random configurations generated by the arrangement engine, we follow the pick planning and grasp heuristics selection strategy proposed by Schwarz *et al.* [11]. See Fig. 2 for a visualization. We substitute the semantic segmentation network with the ground-truth segmentation masks produced by Stillleben. After extracting object contours, ideal suction points are found inside the contour. Depending on object weight, either the *pole of inaccessibility* [13], i.e. the point with maximum distance to the contour is found to minimize the chance of catching other objects, or the center of mass is computed from the contour to ensure good mass distribution.

The system introduced by Schwarz *et al.* [11] also computes a *clutter graph*, identifying which object is resting on top of which object. The final decision on which object to pick is based on this graph, ensuring that we do not attempt to grasp objects caught beneath others.

Once the target object and the corresponding suction point are determined, we find an inverse kinematics solution for the articulated gripper (see Fig. 4), which places the suction cup orthogonally on the local surface, as estimated by a local average of the pixel-wise normals. A gripper trajectory is then computed to bring the gripper to the target pose, apply suction, and lift the object out of the tote.

The gripper is moved along the trajectory using Cartesian impedance control, with a stiffness of $2500 \frac{N}{m}$ and a spring damping of $200 \frac{Ns}{m}$. The force exerted by the impedance control is limited to 200 N. The high stiffness simulates an industrial robotic arm holding the gripper in place.

Once the target position is reached, switching on the suction is simulated by performing ten raycasts in gripper direction, emitting from the gripper perimeter. Every object hit by the raycast (with a maximum distance of 3 cm) is considered caught. For every caught object, a PhysX joint is created between it and the suction cup, simulating a strong force pulling the object against the gripper as well as limiting its orientation relative to the gripper. The joints have a force



Fig. 2. Picking heuristic. Starting from an RGB image and the ground truth segmentation, our ARC 2017 pipeline [11] generates object contours with suction grasp points, as well as a clutter graph describing the scene layout. Each edge in the graph points from the object on top to the object below.

TABLE II

DATASETS STATISTICS							
Mode	Split	Frames	Object visibility				
Untargeted pick	train	137,544	0.77				
Move	train	99,786	0.67				
Targeted pick	train	164,991	0.77				
Untargeted pick	test	31,119	0.77				
Move	test	23,910	0.69				
Targeted pick	test	45,882	0.75				
Total		503232	0.74				

Number of frames in each SynPick dataset split along with the corresponding mean visibility fraction.

limit of 40 N or 2 N, depending on whether all rays found a target (indicating a good vacuum seal). If the force required to keep the object at the suction cup exceeds this limit, the joint breaks and the object is dropped back into the tote.

The picking process is repeated until the tote is empty. If three picking attempts have failed, we also stop the sequence to prevent infinite loops. An example sequence is shown in Fig. 5a).

C. Targeted Picking

In a second mode, we simulate *targeted* picking, where a specific object needs to be extracted. We run the same pipeline as above, with one key difference: We choose the object which is occluded *most*, i.e. is at the bottom of the clutter graph. This choice should lead to more complex object interactions during picking.

D. Moving

As a third possible action, we perform a non-picking manipulation sequence. The goal is to disturb the object arrangement so that other objects become visible.

We simulate the move action by moving the gripper, starting from one corner inside the tote, to all four corners in random order. The gripper is moved at a fixed velocity of $0.1 \frac{m}{s}$. In this mode, the gripper is operated with a lower stiffness of $1000 \frac{N}{m}$ and a force limit of 30 N. This ensures we do not squeeze objects against the immovable tote too much, which could result in instability of the physics simulation. An exemplary scene for the move action is shown in Fig. 5b).

IV. DATASET STATISTICS



Fig. 6. Number of 6D pose annotations for each object category present in SynPick dataset splits.

1) Number of Frames: SynPick consists of 240 scenes for training and 60 scenes for testing for each of the targeted pick, untargeted pick, and move tasks. Each scene consists of a varying number of frames. In Table II, we present the number of frames present in each split.

2) Mean Visibility Fraction: Like in any computer vision task, occlusions present a significant hindrance for 6D pose estimation. A prerequisite for training robust 6D pose estimation models is a dataset that captures real-world occlusion scenarios. Physically realistic simulation of dynamic bin picking scenes captures more realistic occlusion scenarios that are not captured in a static scene. To analyze the degree of occlusions present in the SynPick dataset, we present in Table II, the mean visibility fraction for the different SynPick dataset splits.

3) Object Distribution: As discussed in Section III-A, the arrangement engine picks objects at random until the total volume exceeds a threshold. In Section IV, we present the distribution of object categories across various splits of the SynPick dataset.

V. BASELINE OBJECT PERCEPTION METHODS

A. Single-View RGB(D) 6D Pose Estimation

a) State-of-the-Art Methods: The state-of-the-art methods for single-view RGB(D) pose estimation are predominantly deep-learning-based. Some example methods include direct-regression [2], [14], [15], keypoint-based methods



Fig. 5. Exemplary scenes from the dataset demonstrating the evolution of the scene while the gripper is performing picking (a) and moving (b) actions. Objects in the tote are covered/uncovered as the scene evolves.

[15]–[18], render-and-compare methods [19]–[21], and augmented autoencoders [22], [23]. The BOP Challenge is organized to benchmark the progress of 6D pose estimation methods [1], [5].

b) Baseline Method: We evaluate CosyPose [24], the state-of-the-art winning entry of BOP Challenge 2020 [5], on the SynPick dataset to establish a baseline for our dataset. CosyPose consists of an object detection module, a coarse refinement module, and a fine refinement module. The coarse and fine refinement modules, based on DeepIM [19], formulate 6D pose estimation as an iterative refinement process. The coarse refinement module estimates an initial 6D pose given the canonical pose of the object as the input. It is trained completely on a large synthetically generated dataset. In contrast, the fine refinement module that estimates an accurate 6D pose given the coarse 6D pose as an input is trained using the YCB-Video dataset, disturbing the ground truth pose annotations to form inputs to be corrected by the network. Since the coarse refinement module is trained only on a domain-agnostic synthetic dataset, fine-tuning it on the SynPick dataset is not necessary. Thus, we directly use the coarse refinement module weights provided by the authors, whereas we fine-tune the object detection and fine pose refinement module on the SynPick dataset. Evaluation results on our test splits are presented in Table III.

c) Evaluation Metrics: We use the area under the accuracy-threshold curve (AUC) of ADD and ADD-S metrics for varying thresholds between 0 and 0.1m [2]. ADD metric is the average distance between the model points in ground truth and estimated poses. Formally,

$$ADD = \frac{1}{m} \sum_{x \in \mathcal{M}} \| (\mathbf{R}\mathbf{x} + \mathbf{T}) - (\widetilde{\mathbf{R}}\mathbf{x} + \widetilde{\mathbf{T}}) \|$$
(1)

where \mathcal{M} is the set of 3D model points with m number of points, \mathbf{R} and \mathbf{T} are orientation and translation components of ground truth 6D pose, and $\widetilde{\mathbf{R}}$ and $\widetilde{\mathbf{T}}$ are estimated orientation and translation, respectively.

Objects that exhibit symmetries perform poorly on the ADD metric. ADD-S is a variant of ADD metric that takes symmetries into account by formulating an ICP-like metric that selects closes points:

$$ADD-S = \frac{1}{m} \sum_{x_1 \in \mathcal{M}} \min_{x_2 \in \mathcal{M}} \| (\mathbf{R}\mathbf{x_1} + \mathbf{T}) - (\widetilde{\mathbf{R}}\mathbf{x_2} + \widetilde{\mathbf{T}}) \|.$$
(2)

B. 6D Pose Tracking

The 6D pose estimation baseline established in Section V-A is useful, but does not really capture a real bin picking situation. In our picking scenario, which is typical for industrial applications, a tote of objects is emptied completely, object by object. It is certainly beneficial to monitor the object poses over time, to make use of dependencies between frames. This way, not only temporary effects such as occlusions by the gripper or other objects can be mitigated, but noisy predictions can also be smoothed to obtain a more precise estimate than from a single frame alone.

TABLE III EVALUATION METRICS.

	CosyPose		Filter	$\alpha = 0.05$	Filter α =0.2	
Object	ADD	ADD-S	ADD	ADD-S	ADD	ADD-S
master_chef_can	41.5	79.0	48.8	81.8	47.4	80.4
cracker_box	81.5	88.3	68.5	76.4	64.8	73.9
sugar_box	71.1	76.4	72.8	79.0	70.5	77.2
tomato_soup_can	42.4	51.4	48.1	61.5	44.4	57.8
mustard_bottle	69.9	80.4	65.9	76.8	64.7	76.0
tuna_fish_can	57.2	75.8	62.6	76.9	62.2	76.5
pudding_box	60.2	68.3	68.1	76.9	65.1	74.7
gelatin_box	60.4	64.7	69.1	74.2	66.8	72.9
potted_meat_can	70.0	78.9	65.5	77.5	63.6	76.4
banana	37.3	50.7	38.4	50.2	37.9	49.7
pitcher_base	79.2	84.9	72.0	80.2	70.0	78.5
bleach_cleanser	68.7	74.0	70.4	78.0	68.5	76.9
bowl	75.7	74.2	75.1	75.1	74.1	74.1
mug	64.7	77.2	64.2	77.7	61.7	76.3
power_drill	69.6	76.6	76.2	83.2	73.8	82.0
wood_block	73.4	72.3	74.7	74.7	73.2	73.2
scissors	44.7	49.8	41.2	47.0	38.7	44.6
large_marker	31.3	37.9	36.9	45.2	37.2	45.4
large_clamp	56.8	68.5	60.2	60.2	59.1	59.1
extra_large_clamp	56.7	72.5	59.4	59.4	57.5	57.5
foam_brick	54.6	75.3	76.5	76.5	74.8	74.8
Mean	60.3	70.3	62.6	70.9	60.8	69.4

a) State of the Art: While there are many works on single-view pose estimation, the research field in 6D pose tracking is narrower, but just as diverse. Wang *et al.* [25] present a 6D pose estimation and tracking framework based on self-supervised sparse keypoints. Deng *et al.* [23] formulate the pose tracking problem in the particle filter framework. In contrast, Wen *et al.* [26] follow a render-andcompare approach to perform very fast 6D pose tracking with 90 fps.

b) Baseline: We present a naive tracking baseline based on the CosyPose approach evaluated in Section V-A. This baseline is intended to demonstrate the low-hanging fruit which can be reached by temporal filtering. We implement an exponentially moving average (EMA) with recursive filter coefficient α :



Fig. 7. Position trajectory (z axis) of the *scissors* object in the first test scene. We show raw CosyPose predictions and an exponentially moving average with different recursive filter coefficients α . The time of each pick attempt has been marked with a red circle.

$$\hat{t}_n = (1 - \alpha)\hat{t}_{n-1} + \alpha t_n,$$
 (3)

where t_n is the CosyPose translation estimate at frame n of the sequence and \hat{t}_n is the filtered output. The object orientations are filtered very similarly, but in quaternion space in order to interpolate the orientations correctly:

$$\hat{q}_n = \operatorname{slerp}(\hat{q}_{n-1}, q_n, \alpha), \tag{4}$$

where slerp is the spherical linear interpolation function, which interpolates with $0 < \alpha < 1$ between the two given rotations.

Figure 7 shows a sequence of raw CosyPose translation predictions in the z axis (into the image) for one exemplary object. It can be seen that CosyPose exhibits both stationary noise as well as large deviations, which are mostly caused by temporary occlusions—either by the gripper or other objects. While our naive filtering cannot address steady-state errors, it does smoothen the stationary noise and softens the large jumps caused by occlusions. Table III corroborates these results: Both ADD and ADD-S scores benefit from filtering.

VI. DISCUSSION & CONCLUSION

We presented the SynPick dynamic bin-picking dataset together with its online generator. Our baseline experiments demonstrate that state-of-the-art can achieve good per-frame accuracy, but also that temporal filtering should be employed to correct both estimator noise and effects of temporary occlusions. Further research in the applicability of recent tracking approaches to industrial robotics applications is definitely warranted. We hope that our dataset can serve as an inspiration to researchers in the field of bin picking and facilitates 6D object tracking for industrial automation.

ACKNOWLEDGMENT

This research has been supported by an Amazon Research Award and by the Competence Center for Machine Learning Rhine Ruhr (ML2R), which is funded by the Federal Ministry of Education and Research of Germany (grant no. 01—S18038A).

REFERENCES

- 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 *European Conference on Computer Vision (ECCV)*, 2018, pp. 19–34.
- [2] 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.
- [3] E. Brachmann, A. Krull, F. Michel, S. Gumhold, J. Shotton, and C. Rother, "Learning 6D object pose estimation using 3D object coordinates," in *European Conference on Computer Vision (ECCV)*, 2014, pp. 536–551.
- [4] R. Kaskman, S. Zakharov, I. Shugurov, and S. Ilic, "HomebrewedDB: RGB-D dataset for 6D pose estimation of 3D objects," in *IEEE International Conference on Computer Vision (ICCV) Work*shops, 2019.
- [5] T. Hodaň, M. Sundermeyer, B. Drost, Y. Labbé, E. Brachmann, F. Michel, C. Rother, and J. Matas, "BOP challenge 2020 on 6D object localization," in *European Conf. on Computer Vision (ECCV)*, 2020.

- [6] J. Tremblay, T. To, and S. Birchfield, "Falling things: A synthetic dataset for 3D object detection and pose estimation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2018, pp. 2038–2041.
- [7] T. Hodaň, V. Vineet, R. Gal, E. Shalev, J. Hanzelka, T. Connell, P. Urbina, S. N. Sinha, and B. Guenter, "Photorealistic image synthesis for object instance detection," in *IEEE International Conference on Image Processing (ICIP)*, 2019, pp. 66–70.
- [8] C. Rennie, R. Shome, K. E. Bekris, and A. F. De Souza, "A dataset for improved RGBD-based object detection and pose estimation for warehouse pick-and-place," *IEEE Robotics and Automation Letters* (*RA-L*), pp. 1179–1185, 2016.
- [9] M. Schwarz and S. Behnke, "Stillleben: Realistic scene synthesis for deep learning in robotics," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10502–10508.
- [10] X. Deng, Y. Xiang, A. Mousavian, C. Eppner, T. Bretl, and D. Fox, "Self-supervised 6D object pose estimation for robot manipulation," in *Int. Conference on Robotics and Automation (ICRA)*, 2020.
- [11] M. Schwarz, C. Lenz, G. M. García, S. Koo, A. S. Periyasamy, M. Schreiber, and S. Behnke, "Fast object learning and dualarm coordination for cluttered stowing, picking, and packing," in *International Conference on Robotics and Automation (ICRA)*, 2018.
- [12] K. Mamou, E Lengyel, and A. Peters, "Volumetric hierarchical approximate convex decomposition," in *Game Engine Gems 3*, AK Peters, 2016, pp. 141–158.
- [13] D. Garcia-Castellanos and U. Lombardo, "Poles of inaccessibility: A calculation algorithm for the remotest places on earth," *Scottish Geographical Journal*, vol. 123, no. 3, pp. 227–233, 2007.
- [14] C. Wang, D. Xu, Y. Zhu, R. Martín-Martín, C. Lu, L. Fei-Fei, and S. Savarese, "DenseFusion: 6D object pose estimation by iterative dense fusion," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 3343–3352.
- [15] B. Tekin, S. N. Sinha, and P. Fua, "Real-time seamless single shot 6D object pose prediction," in *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2018.
- [16] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao, "PVNet: Pixel-wise voting network for 6DOF pose estimation," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [17] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3828–3836.
- [18] M. Oberweger, M. Rad, and V. Lepetit, "Making deep heatmaps robust to partial occlusions for 3D object pose estimation," in *European Conference on Computer Vision (ECCV)*, 2018.
- [19] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox, "DeepIM: Deep iterative matching for 6D pose estimation," in *European Conference* on Computer Vision (ECCV), 2018, pp. 683–698.
- [20] J. Shao, Y. Jiang, G. Wang, Z. Li, and X. Ji, "PFRL: Pose-free reinforcement learning for 6D pose estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [21] A. S. Periyasamy, M. Schwarz, and S. Behnke, "Refining 6D object pose predictions using abstract render-and-compare," in *International Conf. on Humanoid Robots (Humanoids)*, 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 *European Conference on Computer Vision* (ECCV), 2018.
- [23] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "PoseRBPF: A Rao-Blackwellized particle filter for 6D object pose tracking," in *Robotics: Science and Systems (RSS)*, 2019.
- [24] Y. Labbe, J. Carpentier, M. Aubry, and J. Sivic, "CosyPose: Consistent multi-view multi-object 6D pose estimation," in *European Conference on Computer Vision (ECCV)*, 2020.
- [25] C. Wang, R. Martín-Martín, D. Xu, J. Lv, C. Lu, L. Fei-Fei, S. Savarese, and Y. Zhu, "6-PACK: Category-level 6D pose tracker with anchor-based keypoints," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 059–10 066.
- [26] B Wen, C Mitash, B Ren, and K. Bekris, "SE(3)-TrackNet: Datadriven 6D pose tracking by calibrating image residuals in synthetic domains," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.