

SCENEREPLICA: Benchmarking Real-World Robot Manipulation by Creating Replicable Scenes

Ninad Khargonkar*, Sai Haneesh Allu*, Yangxiao Lu, Jishnu Jaykumar P
Balakrishnan Prabhakaran, Yu Xiang

Abstract—We present a new reproducible benchmark for evaluating robot manipulation in the real world, specifically focusing on a pick-and-place task. Our benchmark uses the YCB object set, a commonly used dataset in the robotics community, to ensure that our results are comparable to other studies. Additionally, the benchmark is designed to be easily reproducible in the real world, making it accessible to researchers and practitioners. We also provide our experimental results and analyzes for model-based and model-free 6D robotic grasping on the benchmark, where representative algorithms are evaluated for object perception, grasping planning, and motion planning. We believe that our benchmark will be a valuable tool for advancing the field of robot manipulation. By providing a standardized evaluation framework, researchers can more easily compare different techniques and algorithms, leading to faster progress in developing robot manipulation methods.¹

I. INTRODUCTION

Benchmarking has been crucial for successfully boosting the rate of research in machine learning-related communities. With a common benchmark for experiments, different algorithms can be compared in a fair and standardized manner, and a relative improvement between them can be established. A well-known benchmark in AI research is the ImageNet dataset [1] and its associated challenge [2]. It has been used to train and test image classification methods and can be attributed to the rise of network architecture revolution in deep learning. Similarly, the KITTI dataset [3] has been widely used to study problems related to autonomous driving such as car detection and tracking, lane segmentation, and odometry estimation. In NLP too, a wide variety of benchmarks were introduced with different levels of difficulty, such as GLUE [4], SuperGLUE [5] and MMLU [6]. Recently, researchers have proposed several benchmarks for different robotics tasks, such as the NIST assembly boards [7] for assembly and the GRASPA [8] for grasping. However, most of them have not seen widespread adoption and usage, thereby limiting the rate of research progress as it becomes harder to compare methods and investigate their true applicability to a problem.

The key difficulty in robotics benchmarking is that robot tasks in the real world involve a complex pipeline compared to running experiments on fixed test datasets. Real-world tasks don't usually have a common test dataset, and it is difficult to ensure replicability of the test scenarios across different

research labs. Despite this, the robotics community has made progress in several interesting directions for benchmarking, especially robot manipulation. First, for assembly tasks, the NIST assembly boards [7] provide a reproducible test bed. Second, the YCB Object and Model Set [9] provides a way for easily available objects with their 3D mesh models. Although these objects are limited to dozens of rigid objects, they are useful for testing manipulation in an accessible manner. Third, several benchmarking protocols for real-world manipulation are introduced using YCB objects. These protocols show how to arrange objects in the real world to create reproducible scenes. For example, [10] introduces a grasp planning protocol that uses drawing on the robot workspace for the placement of objects. GRASPA [8] uses printable boards with AR tags for reproducible placement of YCB objects. However, [10] cannot guarantee accurate reproducible object placement, while [8] has to use AR tags which create unnatural scenes for testing.

In this work, we introduce a new benchmark, SceneReplica, for real-world robot manipulation to overcome the limitations of existing benchmarks. We use 16 YCB objects which can be easily purchased either online or in-person. Second, we created a dataset of 20 scenes with 5 YCB objects in each scene as a test bed for pick-and-place. We utilize a procedural scene generation for dataset design [11] in simulation to create these 20 scenes. The design considers the reachability space of a robot and the diversity of poses of objects in building the real-world test bed. Finally, we provide an easy-to-use tool to guide users in reproducing these 20 scenes in the real world without using any AR tag or external equipment. To the best of our knowledge, SceneReplica is the only benchmark for real-world reproducible scenes without using AR markers or external cutouts to guide object placement.

Another contribution of our work is the benchmarking of two main paradigms for 6D robotic grasping using SceneReplica: model-based grasping and model-free grasping. Model-based grasping leverages 3D models of objects for perception and planning. It consists of a pipeline using 6D object pose estimation for perception [19], [20], [21], 3D model-based grasp planning [22], [23] and motion planning and motion control [24] to reach the planned grasps. Model-free grasping does not assume the availability of 3D object models. It consists of a pipeline of unseen object instance segmentation [25], [26], [27], object point cloud-based grasp planning [28], [29] and motion planning and motion control [24] to reach the planned grasps. Different object perception, grasping planning, motion planning, and control methods can be evaluated within the two 6D robotic grasping

* Equal contribution, the authors are with the Department of Computer Science, University of Texas at Dallas, Richardson, TX 75080, USA {firstname.lastname}@utdallas.edu

¹Appendix, code and videos for the project are available at <https://irvlutd.github.io/SceneReplica>

Benchmark	Type	Task	Objects	AR Tag-Free	Scene Reproducibility
Meta-World [12]	Simulation	50 tasks	Synthetic	✓	✓
RLBench [13]	Simulation	100 Tasks	Synthetic	✓	✓
robosuite [14]	Simulation	9 Tasks	Synthetic	✓	✓
Grasp Planning Protocol [10]	Real	Grasp Planning	YCB (single)	✓	✗
NIST Assembly [7]	Real	Assembly	Task Boards	✓	✓
FurnitureBench [15]	Real	Assembly	3D Printing	✗	✓
GRASPA [8]	Real	Grasping	YCB (clutter)	✗	✓
OCRTOC [16]	Real	Rearrangement	YCB + Others	✓	✗
RB2 [17]	Real	Pouring, Scooping, Zipping, Insertion	Others	✓	✗
Box and Blocks Test [18]	Real	Pick-and-Place	Blocks	✓	✗
SceneReplica (Ours)	Real	Pick-and-Place	YCB (clutter)	✓	✓

TABLE I: Comparison between different robotic manipulation benchmarks.

paradigms. This high-level view of robot manipulation benchmarking is a departure from existing benchmarks which are specific to some tasks or do not give details about the complete pipeline. We also perform experiments and provide our analyses on a selected set of representative methods. To the best of our knowledge, it is the first effort in which model-based and model-free grasping pipelines are systematically evaluated and fairly compared to each other using the same real-world testing scenes.

By creating SceneReplica, we provide the robotics community with an easy-to-setup, reproducible real-world benchmark for robot manipulation. The core method for SceneReplica involving scene generation in simulation and replication in real-world with rendered images is applicable to *any robotic platform*, which allows for testing with different robots. We plan to maintain the benchmark continuously and evaluate additional and upcoming perception, planning, and control methods on it. It is our hope that these results, which combine effects from the entire pipeline, can serve as a reference for robot manipulation research.

II. RELATED WORK

Simulation-based Robotic Benchmarks. Since reproducibility in the real world is challenging, the setup of testing environments in physics simulators provides a workaround for benchmarking in robotics. For example, the Meta-World [12] benchmark introduces 50 robotic manipulation tasks, which can be used to study meta-reinforcement learning and multitask learning. The robosuite [14] based on the MuJoCo physics engine [30] supports 9 manipulation tasks for different robots. The RLBench benchmark [13] increases the number of tasks to 100. The advantage of using simulation is that reproducibility is guaranteed for evaluation, and a large number of tasks can be used. However, due to the sim-to-real gap, the performance in simulation cannot simply be transferred to the real world. Although there is concurrent research effort on improving the physics simulators, nuances in the real world such as lighting, friction, external perturbation, etc., are extremely difficult to model in simulation. This motivates our need to have real-world benchmarks for robotics.

Real-world Robotic Benchmarks. We summarize a list of representative efforts to benchmark robot manipulation in the real world in Table I. According to the evaluation tasks, they can be classified into assembly [7], [15], grasping [10], [18], [8], [31], rearrangement [16], and others (pouring, scooping, zipping, insertion in [17]). In terms of scene reproducibility

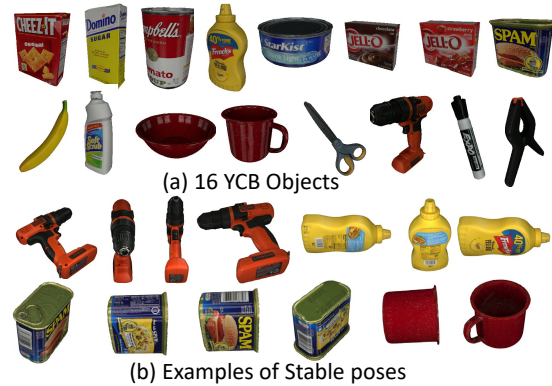


Fig. 1: 16 YCB objects are used in SceneReplica.

in evaluation, NIST Assembly [7], FurnitureBench [15] and GRASPA [8] are the only three benchmarks whose test scenes are reproducible. NIST Assembly uses a fixed set of assembly task boards, while FurnitureBench and GRASPA use AR markers to guide the scene creation process. For grasping and pick-and-place, SceneReplica is the only benchmark that can create locally reproducible scenes without using AR markers. Another flavor of benchmarks includes TOTO [32] and RRC [33] which provide remote access to shared hardware, to benchmark common tasks. SceneReplica can be used in conjunction with such benchmarks in a ranking system as proposed by RB2 [17] where researchers have the option of local and remote evaluation of their proposed methods.

III. THE SCENEREPLICA BENCHMARK

Our goal is to introduce reproducible scenes of objects for robot manipulation, where researchers can set up the same set of scenes in different environments for evaluation. Currently, we focus on tabletop cluttered scenes for pick-and-place where the clutter is in terms of a packed scene with some level of separation between objects. This is in contrast with other classes of cluttered scenes where objects are piled on top of each other, making it harder to replicate exactly in the real-world or simulation. Overall, the scene generation pipeline happens in Gazebo with access to ground truth object data for constraint checking over the scenes.

A. Creating Scenes in Simulation

Object Set. 16 YCB objects [9] are used in our benchmark to create scenes. Their 3D mesh models are shown in Fig. 1. We select these objects because i) they can be easily purchased

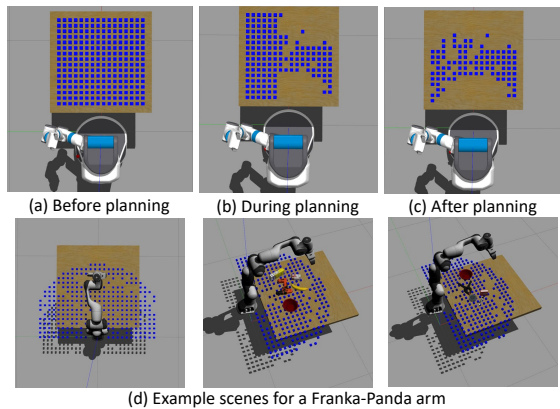


Fig. 2: (a,b) Illustration of motion planning check to filter out reachable locations. (c) Blue cubes remaining on table after planning indicate reachable locations. (d) The algorithm can be extended to different robots.

online or in grocery stores; ii) they are a subset of the YCB Video dataset [19] which has been widely used to study 6D object pose estimation; (iii) they can reliably fit within a parallel-jaw gripper. Therefore, we can leverage existing 6D object pose estimation methods for experiments, for example, [19], [21], [34].

Stable Poses of Objects. We create scenes in a Gazebo simulation, where stable poses of objects are used. We do not create random clutter via a piling procedure for object spawning as it requires careful tuning and is not easily reproducible in real world. Therefore, we opt to compute stable poses of objects, which are the object orientations when they rest statically on a planar surface. Such stable poses are also safe to use in real world scenes as the chance of an object falling over and disturbing the scene is low. We used the `compute_stable_poses()` function in `trimesh` to obtain such poses, and Fig. 1 shows examples of stable poses of some YCB objects.

Determining the Reachable Space of a Robot. When placing these objects on a table, we need to place them within the reach of the robot to avoid motion planning failures. We employ a simple algorithm as depicted in Fig. 2 to compute the reachable regions of a tabletop. In Step I, the robot model is loaded in Gazebo with its arm stowed. Next, a table of dimension $1m \times 1m \times 0.745m$ is placed at an offset of $(0.8, 0, 0)$ with respect to the robot model, where the x -axis is the forward direction of the robot. In Step II, the surface of the table is partitioned into grids 16×16 and a block of size $3cm \times 3cm$ is placed in the center of each cell of the grid. In Step III, we determine whether there is a feasible motion plan for the stand-off pose to each cube. Blocks that do not have a feasible plan will be removed from the scene and earlier steps are repeated. The blocks surviving the removal denote reachable locations which are then used to spawn objects for scene creation.

Scene Generation. Using the stable poses of the 16 YCB objects, we can generate random scenes in Gazebo by placing them in the reachable locations of the Fetch robot. We limit the maximum number of objects in a scene to five and

sequentially add objects to the scene one by one. Initially, an object is placed randomly on the table. When adding the next object, we randomly sample a nearby reachable location with an existing object that is collision-free between these objects. Stable poses of objects are also randomly sampled during placement. This process repeats until five objects are added to the scene. By sampling nearby reachable locations, we tend to generate cluttered scenes where objects occlude each other, thereby creating a strong benchmark.

Scene Selection. After generating a set of scenes, we target to select a set of 20 scenes with the following properties. i) object count distribution across scenes should be balanced. We define a range $[C_{\min}, C_{\max}]$ for the object counts. ii) The selected scenes should adequately cover the stable poses of each object. iii) All the objects in a selected scene should be reachable from the robot. To check the reachability of an objects in a scene, we check for existence of valid motion plans to (offline-generated) grasps around the object.

With these three properties, our algorithm for scene selection works as follows. We generate a large number of scene sets. Each scene set contains 20 scenes, and it satisfies properties i) and iii). In addition, each set of scenes is given a score on how many stable poses of these 16 YCB objects are covered by the 20 scenes in it. This acts as a measure of the diversity of a set. Finally, we select the set of scenes with the highest pose coverage score. Fig. 3 shows the final 20 scenes selected for SceneReplica.

Transfer to other robots. Generating in simulation enables access to complete state information, allowing for a seamless `sim2real` transfer for reproducible real-world experiments across different robotic platforms (e.g. 2) For use with a new robot, users just need to re-render the reference images (e.g. 3) for the scenes in simulation via the robot camera’s parameters. Users then recreate the scenes in real world using the rendered images which act as the new reference images. Further details are presented in the supplementary material.

B. Reproducing Scenes in the Real World

To use the 20 selected scenes for experiments, we need to replicate these scenes in the real world. Previous benchmarks rely on AR markers to create reproducible scenes [8], [15]. In contrast, we show that directly using the scene images in Fig. 3 as reference images, we can place objects in the real world as in the simulator. The process of replicating a scene in the real world is illustrated in Fig. 4. First, the Fetch camera pose in the real world is set to be the same as in the simulation. Therefore, the camera extrinsic, that is, the $SE(3)$ transformation from the robot base frame to the RGB camera frame, is always the same. This can be done by controlling the Fetch head pose. Second, maintain uniform table height at $0.745m$ for both real-world and simulated environments, achievable with a height-adjustable table. Third, with the same camera extrinsic and table height, we can overlay the half-transparent version of an RGB image of a reference scene to the real RGB image from the Fetch camera for object placement guidance. Finally, the user can place objects by interactively checking the overlap between




Fig. 3: 20 scenes in our SceneReplica benchmark with 5 YCB objects in each scene.

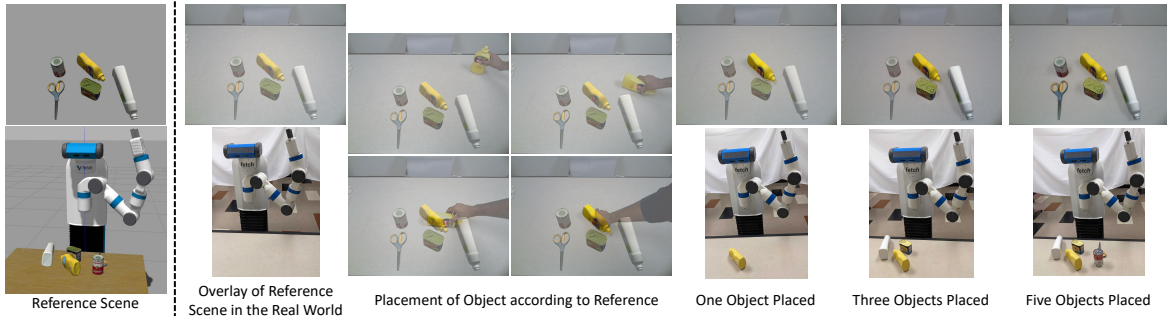


Fig. 4: The process of replicating a scene in the real world. The image of the reference scene is overlaid on the image of the real camera to guide how to place objects into the real-world scene.

the reference image and the real-world image for accurate alignment (Fig. 4).

By matching camera extrinsic between simulation and reality, object poses in the robot base frame replicate those in the camera frame for manipulation. This procedure is transferable to different robot platforms using SceneReplica, as only the reference images need to be re-generated by the new camera.




Fig. 5: Illustration of the coordinate frames.

Fig. 5 illustrates these coordinate frames used for our experiments.

IV. MODEL-BASED AND MODEL-FREE 6D GRASPING

In the SceneReplica benchmark, we assess tabletop pick-and-place tasks involving grasping, lifting, gripper rotation, and object drop-off, in sequence. Metrics are collected for each stage, with success categorized as (1) grasping and (2) pick-place success. Grasping-success indicates successful grasp but failure in later steps, while pick-and-place success denotes completion of the entire pipeline. This task necessitates 6D grasping, controlling both gripper orientation and translation, distinct from top-down grasping.

A. Model-based 6D Grasping

Object Perception. When 3D models of objects are available, we deal with 6D object pose estimation, an active research area in robotics and computer vision [35]. It, determines 3D orientation and translation with respect to the camera. Though pose accuracy can be assessed on datasets, its impact on grasping remains unclear. SceneReplica integrates pose estimation with planning and control for 6D grasping, enabling analysis of pose accuracy’s effect on success rates

Grasp Planning. Model-based grasp planning methods such as GraspIt! [22] and OpenRAVE [23] use 3D models of objects in physics simulators to generate grasps. We use GraspIt! in our experiments. We initially generated numerous grasps for each YCB object in GraspIt! then sampled a fixed amount using farthest point sampling for grasp pose diversity.

Model-based Motion Planning. After estimating the 6D pose, planned grasps on the object model are transformed to the robot base frame. Each grasp is checked for a feasible trajectory to reach it using motion planning algorithms from OMPL [24]. Upon obtaining a trajectory, it is executed to grasp the target. All Object pose estimation enables utilization of object 3D models as obstacles in the planning scene.

B. Model-free 6D Grasping

Object Perception. When there is no 3D model of objects available, we can tackle the *unseen object instance segmentation (UOIS)* problem [25], [26], [27] to segment arbitrary objects. Once the objects are segmented, model-free




Fig. 6: Model-based grasping pipeline used for benchmarking




Fig. 7: Model-free grasping pipeline used for benchmarking

grasp planning methods can be used to generate grasps of the segmented objects. UOIS methods are usually trained using synthetic images of a large number of objects. They learn the concepts of objects and are able to segment novel objects in the real world.

Grasp Planning. Grasp planning methods such as 6D GraspNet [28], [29] take point clouds of objects or whole scenes as input and output grasps of objects in the scenes. Neural networks trained for grasp synthesis, like those for unseen object instance segmentation, learn from large datasets of object-grasp pairs, allowing them to plan grasps for new objects using partial point clouds.

Motion Planning. Once the target object is segmented and grasps are synthesized, a motion planning algorithm generates a trajectory towards a planned grasp. Instead of using 3D object models, oriented 3D bounding boxes for the segmented objects are computed for obstacle avoidance during planning.

V. EXPERIMENTS

In pick-and-place tasks, the primary evaluation metric is the success rate of transporting objects between the locations. This success relies on the performance of all the framework components. Even with accurate perception, errors in grasp planning can prevent the robot from picking up objects effectively. Therefore, we evaluated the success rate, perception evaluation and failure analysis of pick-and-place of 100 trials [20 scenes \times 5 objects per scene] across diverse grasping frameworks presented in Table II.

Failure Analysis. Any error during the experiment is classified into 3 categories: (i) Perception error (ii) Planning error (iii) Execution error. We use the average distance for symmetric objects (ADD-S) metric mentioned in [19] to evaluate 6D object pose estimation. For unseen object instance segmentation, we follow [25], [26] to compute precision, recall, F-measure for region & boundary overlaps.

(i) A *Perception Error* occurs when grasping fails due to either pose estimation (in model-based grasping) or segmentation (in model-free grasping) inaccuracies. We denote a pose estimation error when the grasp fails and $ADD-S > 0.1 * d$ ($d = 3D$ object model’s diameter). For segmentation methods, we consider the prediction a failure if the overlap F-measure is $< 75\%$. (ii) Despite no obvious perception error, if the

manipulator fails to grasp the object, it indicates issues with either the grasp strategy or trajectory planning. In such a scenario we record a *Planning Error* for the experiment. (iii) When the robot successfully grasps the object but fails to drop it off due to slipping or any other reason, it is classified as an *Execution Error*.

Grasping Orders. For grasping in cluttered scenes, the order of grasping objects matters. We evaluate two orders in SceneReplica when the grasping order can be programmed: the near-to-far order based on object perception and the fixed order we defined for each scene. Intuitively, the near-to-far order is simpler, since the robot can clear up closer objects first, those which potentially occlude other objects in the scene.

Model-based Grasping Frameworks. For model-based grasping, we have evaluated three 6D object pose estimation methods: PoseCNN [19] and PoseRBPF [21] and GDRNPP [34], [36]. GraspIt! [22] is used for grasp planning in model-based grasping. The OMPL library [24] in MoveIt is used to plan the trajectory of the robot arm and control to follow the planned trajectory. In cases when there is no motion plan to reach all the grasps from GraspIt!, we simply apply top-down grasping to grasp the target object.

From Table II we can see that the GDRNPP-based framework achieves a higher success rate in both grasping and pick-and-place. GDRNPP is the overall best method in the BOP Challenge 2022 [38]. The accuracy of pose estimation in terms of ADD-S is presented in Table III, which is calculated using 200 images in the real world, that is, 20 scenes \times 5 grasps per scene \times 2 orders. A better pose estimation model improves the success rate.

Model-free Grasping Frameworks. We have evaluated two unseen object segmentation methods: UCN [26] and MSMFormer [27], two grasp planning methods: 6D GraspNet [28] and Contact-GraspNet [29]. The OMPL library [24] in MoveIt used for motion planning and control. Similarly, top-down grasping is used if there is no plan from the GraspNets. From the results in Table II, we can see that the Contact-GraspNet-based methods achieve a higher success rate than the 6D GraspNet-based methods. Contact-GraspNet grasps usually have higher quality by modeling the contact

Method #	Perception	Grasp Planning	Motion Planning	Control	Ordering	Pick-and-Place Success	Grasping Success
Model-based Grasping							
1	PoseRBPF [21]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Near-to-far	58 / 100	64 / 100
1	PoseRBPF [21]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Fixed	59 / 100	59 / 100
2	PoseCNN [19]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Near-to-far	47 / 100	48 / 100
2	PoseCNN [19]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Fixed	40 / 100	45 / 100
3	GDRNPP [34], [36]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Near-to-far	66 / 100	69 / 100
3	GDRNPP [34], [36]	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Fixed	62 / 100	64 / 100
Model-free Grasping							
4	UCN [26]	GraspNet [28] + Top-down	OMPL [24]	MoveIt	Near-to-far	43 / 100	46 / 100
4	UCN [26]	GraspNet [28] + Top-down	OMPL [24]	MoveIt	Fixed	37 / 100	40 / 100
5	UCN [26]	Contact-graspnet [29] + Top-down	OMPL [24]	MoveIt	Near-to-far	60 / 100	63 / 100
5	UCN [26]	Contact-graspnet [29] + Top-down	OMPL [24]	MoveIt	Fixed	60 / 100	64 / 100
6	MSMFormer [27]	GraspNet [28] + Top-down	OMPL [24]	MoveIt	Near-to-far	38 / 100	41 / 100
6	MSMFormer [27]	GraspNet [28] + Top-down	OMPL [24]	MoveIt	Fixed	36 / 100	41 / 100
7	MSMFormer [27]	Contact-graspnet [29] + Top-down	OMPL [24]	MoveIt	Near-to-far	57 / 100	65 / 100
7	MSMFormer [27]	Contact-graspnet [29] + Top-down	OMPL [24]	MoveIt	Fixed	61 / 100	70 / 100
8	MSMFormer [27]	Top-down	OMPL [24]	MoveIt	Fixed	56 / 100	59 / 100
End-to-end Learning-based Grasping							
9	Dex-Net 2.0 [37] (Top-Down Grasping)		OMPL [24]	MoveIt	Algorithmic	43 / 100	51 / 100
Ground truth pose-based Grasping							
10	Ground truth object pose	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Near-to-far	78 / 100	82 / 100
11	Ground truth object pose	GraspIt! [22] + Top-down	OMPL [24]	MoveIt	Fixed	78 / 100	87 / 100

TABLE II: Different grasping frameworks evaluated on SceneReplica using a Fetch mobile manipulator.




Fig. 8: Example failures in our experiment pipelines. Green boxes indicate the targets for grasping.

Pose Method	ADD-S	Segmentation Method	Overlap			Boundary			%75
			P	R	F	P	R	F	
PoseCNN [19]	66.01	UCN+ [26] MSMFormer+ [27]	70.0	81.2	71.9	48.3	57.2	50.3	77.6
PoseRBPF [21]	68.73		76.1	83.3	76.3	56.0	61.4	57.0	80.2
GDRNPP [34]	71.26								

TABLE III: Perception evaluation of pose estimation and unseen object instance segmentation in the real world.

between the gripper with the target. When comparing the two segmentation methods, MSMFormer achieves higher segmentation accuracy, as shown in Table III. The combination of MSMFormer and Contact-GraspNet achieves the best performance among model-free grasping pipelines. MSMFormer with top-down grasping performs as a strong baseline thanks to the object segmentation network.

End-to-end Learning-based Grasping. End-to-end learning approaches for grasping can also be evaluated using SceneReplica. These approaches take an input image and then output a grasping action using convolutional neural networks [37] (Dex-Net 2.0), [39], [40] or policies learned with reinforcement learning [41], [42], [43] (QT-Opt). In particular, these methods focus mainly on top-down grasping. Until now, we have evaluated Dex-Net 2.0 [37] on SceneReplica as shown in Table II, where we synthesize a top-down view depth image using the point cloud from the Fetch RGB-D camera. Its performance is inferior compared to a segmentation-based top-down grasping pipeline due to its requirement of a top-down camera and not relying on a perception module.

Failure Analysis. Fig. 8 shows some examples of different failure types. Detailed success and failure analysis for each object and each method can be found in the supplementary material on our project website. We can see that perception errors and planning errors are dominant and that perception errors are relatively more frequent. Small and thin objects, such as scissors and markers, are more difficult to grasp.

VI. DISCUSSION AND FUTURE WORK

Limitations. Our benchmark is limited by the 16 YCB objects that are used for testing. Using a more diverse set will improve the manipulation benchmarking. Second, only results with the Fetch mobile manipulator are included. Extending the benchmark to robots and external cameras would require setting up a third-person view and conducting hand-eye calibration. Once the camera is set up, we can re-render the scene reference images using the new camera pose, and these images can be used for scene recreation. Our experiments on SceneReplica hint at several directions to improve robot pick-and-place. First, object perception and grasp planning accuracy can be further improved. Second, one important aspect that is not considered in the model-based and model-free grasping methods is force feedback and force closure. Slippage failures can be reduced by introducing force feedback and can also improve the grasping of fragile objects.

Future Work. We plan to continuously evaluate more methods on SceneReplica and release these results to the community. One line of work that is missing from our current evaluation is reinforcement learning-based approaches for 6D grasping. We plan to evaluate additional end-to-end grasping approaches on SceneReplica. We hope the community can use SceneReplica to benchmark their proposed robot manipulation methods. It can also be extended to other manipulation tasks such as object rearrangement in the future.

Acknowledgement. This work was supported in part by the DARPA Perceptually-enabled Task Guidance (PTG) Program under contract number HR00112220005 and the Sony Research Award Program.

REFERENCES

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [2] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, pp. 211–252, 2015.
- [3] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [4] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," *arXiv preprint arXiv:1804.07461*, 2018.
- [5] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman, "Superglue: A stickier benchmark for general-purpose language understanding systems," *Advances in neural information processing systems*, vol. 32, 2019.
- [6] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt, "Measuring massive multitask language understanding," *arXiv preprint arXiv:2009.03300*, 2020.
- [7] K. Kimble, K. Van Wyk, J. Falco, E. Messina, Y. Sun, M. Shibata, W. Uemura, and Y. Yokokohji, "Benchmarking protocols for evaluating small parts robotic assembly systems," *IEEE robotics and automation letters*, vol. 5, no. 2, pp. 883–889, 2020.
- [8] F. Bottarel, G. Vezzani, U. Pattacini, and L. Natale, "Graspa 1.0: Graspa is a robot arm grasping performance benchmark," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 836–843, 2020.
- [9] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.
- [10] Y. Bekiroglu, N. Marturi, M. A. Roa, K. J. M. Adjigble, T. Pardi, C. Grimm, R. Balasubramanian, K. Hang, and R. Stolkin, "Benchmarking protocol for grasp planning algorithms," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 315–322, 2019.
- [11] R. Newbury, M. Gu, L. Chumbley, A. Mousavian, C. Eppner, J. Leitner, J. Bohg, A. Morales, T. Asfour, D. Kragic *et al.*, "Deep learning approaches to grasp synthesis: A review," *IEEE Transactions on Robotics*, 2023.
- [12] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning," in *Conference on robot learning*. PMLR, 2020, pp. 1094–1100.
- [13] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3019–3026, 2020.
- [14] Y. Zhu, J. Wong, A. Mandlke, R. Martín-Martín, A. Joshi, S. Nasiriany, and Y. Zhu, "robosuite: A modular simulation framework and benchmark for robot learning," *arXiv preprint arXiv:2009.12293*, 2020.
- [15] M. Heo, Y. Lee, D. Lee, and J. J. Lim, "Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation," *arXiv preprint arXiv:2305.12821*, 2023.
- [16] Z. Liu, W. Liu, Y. Qin, F. Xiang, M. Gou, S. Xin, M. A. Roa, B. Calli, H. Su, Y. Sun *et al.*, "Oortco: A cloud-based competition and benchmark for robotic grasping and manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 486–493, 2021.
- [17] S. Dasari, J. Wang, J. Hong, S. Bahl, Y. Lin, A. Wang, A. Thankaraj, K. Chahal, B. Calli, S. Gupta *et al.*, "Rb2: Robotic manipulation benchmarking with a twist," *arXiv preprint arXiv:2203.08098*, 2022.
- [18] A. S. Morgan, K. Hang, W. G. Bircher, F. M. Alladkani, A. Gandhi, B. Calli, and A. M. Dollar, "Benchmarking cluttered robot pick-and-place manipulation with the box and blocks test," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 454–461, 2019.
- [19] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *arXiv preprint arXiv:1711.00199*, 2017.
- [20] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *arXiv preprint arXiv:1809.10790*, 2018.
- [21] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "Poserbpf: A rao-blackwellized particle filter for 6-d object pose tracking," *IEEE Transactions on Robotics*, vol. 37, no. 5, 2021.
- [22] A. T. Miller and P. K. Allen, "Grasplit! a versatile simulator for robotic grasping," *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [23] R. Diankov and J. Kuffner, "Openrave: A planning architecture for autonomous robotics," *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34*, vol. 79, 2008.
- [24] I. A. Sucas, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [25] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, "The best of both modes: Separately leveraging rgb and depth for unseen object instance segmentation," in *Conference on robot learning*, 2020, pp. 1369–1378.
- [26] Y. Xiang, C. Xie, A. Mousavian, and D. Fox, "Learning rgb-d feature embeddings for unseen object instance segmentation," in *Conference on Robot Learning*. PMLR, 2021, pp. 461–470.
- [27] Y. Lu, Y. Chen, N. Ruozzi, and Y. Xiang, "Mean shift mask transformer for unseen object instance segmentation," *arXiv preprint arXiv:2211.11679*, 2022.
- [28] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [29] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 438–13 444.
- [30] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.
- [31] M. Rudorfer, M. Suchi, M. Sridharan, M. Vincze, and A. Leonardis, "Burg-toolkit: Robot grasping experiments in simulation and the real world," *arXiv preprint arXiv:2205.14099*, 2022.
- [32] G. Zhou, V. Dean, M. K. Srirama, A. Rajeswaran, J. Pari, K. Hatch, A. Jain, T. Yu, P. Abbeel, L. Pinto *et al.*, "Train offline, test online: A real robot learning benchmark," *arXiv preprint arXiv:2306.00942*, 2023.
- [33] S. Bauer, M. Wüthrich, F. Widmaier, A. Buchholz, S. Stark, A. Goyal, T. Steinbrenner, J. Akpo, S. Joshi, V. Berenz *et al.*, "Real robot challenge: A robotics competition in the cloud," in *NeurIPS 2021 Competitions and Demonstrations Track*. PMLR, 2022, pp. 190–204.
- [34] G. Wang, F. Manhardt, F. Tombari, and X. Ji, "GDR-Net: Geometry-guided direct regression network for monocular 6d object pose estimation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 16611–16621.
- [35] "International workshop on recovering 6d object pose," http://cmp.felk.cvut.cz/sixd/workshop_2022/, 2022.
- [36] X. Liu, R. Zhang, C. Zhang, B. Fu, J. Tang, X. Liang, J. Tang, X. Cheng, Y. Zhang, G. Wang, and X. Ji, "Gdrnpp," https://github.com/shanice-l/gdrnpp_bop2022, 2022.
- [37] J. Mahler, F. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [38] 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 *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 19–34.
- [39] D. Morrison, P. Corke, and J. Leitner, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," *arXiv preprint arXiv:1804.05172*, 2018.
- [40] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9626–9633.
- [41] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [42] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4238–4245.
- [43] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke *et al.*, "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation," *arXiv preprint arXiv:1806.10293*, 2018.