

# SceneComplete: Open-World 3D Scene Completion in Cluttered Real World Environments for Robot Manipulation

Aditya Agarwal<sup>1</sup>, Gaurav Singh<sup>2</sup>, Bipasha Sen, Tomás Lozano-Pérez<sup>3</sup>, and Leslie Pack Kaelbling<sup>4</sup>

**Abstract**—Careful robot manipulation in every-day cluttered environments requires an accurate understanding of the 3D scene, in order to grasp and place objects stably and reliably and to avoid colliding with other objects. In general, we must construct such a 3D interpretation of a complex scene based on limited input, such as a single RGB-D image. We describe SceneComplete, a system for constructing a complete, segmented, 3D model of a scene from a single view. SceneComplete is a novel pipeline for composing general-purpose pretrained perception modules (vision-language, segmentation, image-inpainting, image-to-3D, visual-descriptors and pose-estimation) to obtain highly accurate results. We demonstrate its accuracy and effectiveness with respect to ground-truth models in a large benchmark dataset and show that its accurate whole-object reconstruction enables robust grasp proposal generation, including for a dexterous hand. We release the code and additional results on our website.

**Index Terms**—Perception for grasping and manipulation, RGB-D perception, manipulation planning.

## I. INTRODUCTION

AS MANIPULATION robots move from constrained environments such as factories and workshops to open-world environments such as homes and hospitals, they must be able to construct representations of their environment that enable robust, careful manipulation. Such representations need to individuate objects and characterize their shapes, so that the robot can reliably select stable grasps and placements for individual objects and manipulate them without unwanted collisions. These representations must generally be constructed from limited input, such as a single RGB-D image. This problem is fundamentally ill-posed, but we are now in a position to address it using strong priors that have been learned by vision foundation models.

Received 14 June 2025; accepted 14 October 2025. Date of publication 10 November 2025; date of current version 24 November 2025. This article was recommended for publication by Associate Editor Y. Shirai and Editor J. Borràs Sol upon evaluation of the reviewers' comments. This work was supported in part by NSF under Grant 2214177, in part by AFOSR under Grant FA9550-22-1-0249, in part by ONR MURI under Grant N00014-22-1-2740 and GrantN00014-24-1-2603, in part by the MIT Quest for Intelligence, and in part by Robotics and Artificial Intelligence Institute. (Corresponding author: Aditya Agarwal.)

Aditya Agarwal, Bipasha Sen, Tomás Lozano-Pérez, and Leslie Pack Kaelbling are with the Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139 USA (e-mail: adityaag@mit.edu; bise@mit.edu; tlp@mit.edu; lpk@mit.edu).

Gaurav Singh is with the Department of Computer Science, Brown University, Providence, RI 02912 USA (e-mail: gaurav@brown.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2025.3630884>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2025.3630884

In this letter, we propose a solution to this open-world scene completion problem in the form of a perception pipeline, SceneComplete (cite Fig. 1). It combines multiple large pre-trained vision models into a system that takes a single RGB-D image as input and predicts as output a completed scene, consisting of a set of meshes for all the visible objects, including those that are partially occluded. Crucially, it makes no assumptions about the categories of the objects, their arrangement, or the camera viewpoint. It is constructed from multiple highly capable pre-trained perception components: a vision-language model (VLM) for identifying and generating short descriptions of the objects in a scene, a text-grounded image-segmentation model for localizing objects in the image, a 2D image-inpainting model for predicting the appearance of occluded parts of objects, an image-to-3D model for generating complete object meshes, and visual descriptor and pose-estimation modules to aid in composing individual predicted meshes into a final scene. None of these components can individually solve the problem, but in combination they provide robust object-centric interpretation of complex images, producing a segmented set of object meshes that are suitable for robot planning and manipulation.

To demonstrate the effectiveness of SceneComplete, we conduct extensive quantitative and qualitative evaluations on real-world tabletop scenes. Our quantitative evaluations are on the Graspnet-1B [1] and YCB-Video [2] datasets, which consist of cluttered tabletop scenes. In these scenarios, accurately predicting the full scene—including partially occluded objects—is crucial for stable grasping, collision-free motion with an object in the hand, and reliable placing.

We further illustrate the utility of our shape-reconstruction methods by using them as input to parallel-jaw [3], [4] and dexterous grasping [5] methods, the latter being especially sensitive to the detailed shape of the entire object.

## II. RELATED WORK

**Feed-forward Scene Reconstruction:** Feed-forward multi-object scene completion methods such as [6], [7], [8], [9], [10] learn end-to-end mappings from single-view RGB-D input to completed object meshes or occupancy grids. These methods are attractive because they are computationally fast at inference time; however, suffer from key limitations: ShAPO [7], FSD [6], and CRISP [10] are closed-set methods, trained extensively on fixed benchmark datasets with predefined object categories, and fail to generalize to novel objects. OctMAE [8] addresses open-set scenes by performing scene-level reconstruction. However, its predictions are surface-level and do not individuate objects, limiting their utility for robotic manipulation. ZeroGrasp [9]



Fig. 1. (a) takes as input a *single* RGB-D image of a given scene, visualized here as a point cloud; (b) produces high-quality fully completed, accurately segmented object meshes in scenes with substantial occlusion and novel objects; and (c) enables downstream dexterous manipulation that requires accurate complete shape information.

performs scene reconstruction with individuated objects for grasp generation, but the reconstruction quality—crucial for collision avoidance—is subpar since the focus is primarily on grasping. In our results, we compare and outperform OctMAE and ZeroGrasp in reconstructing scenes and generating grasps.

**Compositional Scene Reconstruction:** This line of work composes multiple open-set models with zero-shot generalization into a pipeline for scene reconstruction. While such methods are typically slower than feedforward methods, they can directly handle unseen objects with little or no retraining [11], [12], [13]. CAST [11] and Gen3DSR [12] both reconstruct full scenes from a single RGB image by first predicting depth maps from monocular depth estimation models. CAST then generates meshes for individual objects and applies a physics-aware correction step to enforce physically consistent placements, while Gen3DSR integrates components such as DreamGaussian to produce full 3D meshes. However, because both methods rely on predicted depth, their pipelines are tightly coupled to synthetic outputs and fail to adapt to raw RGB-D input, where sensor noise is unavoidable. These approaches are primarily designed for asset generation, and their inability to incorporate ground-truth depth limits their applicability in robotics. By contrast, SceneComplete leverages observed sensor depth (often noisy), enabling a more flexible pipeline tailored for manipulation tasks.

A different line of work, Open6DOR [13], introduces a benchmark for language-driven 6-DoF object rearrangement, along with a baseline pipeline for grasp generation. Their approach composes a set of modules similar to those in SceneComplete, but focuses exclusively on predicting object poses. As a result, the pipeline does not perform full scene reconstruction, omits critical steps such as registration, and restricts predictions to objects that are almost fully visible.

### III. METHOD

Fig. 2 illustrates the overall design of SceneComplete. It takes a single RGB-D image as input and produces a set of object meshes that are registered with the input 3D scan. The objective is to provide an accurate 3D reconstruction of the scene, in terms of segmentation into rigid components, and the shape of each component, expressed as a mesh. Importantly, each step in the pipeline makes use of existing pre-trained open-world

visual-processing models, with almost no additional training (we do a small amount of low-rank adaptation of the inpainting model). This means that, as improved models become available for each of these tasks, as they inevitably will, we will be able to immediately profit from these improvements. We describe each process in detail in the following subsections.

#### A. Prompting and Segmentation

We begin by using a vision-language model to determine the number and basic description of objects in the scene. In our implementation, we pass the RGB image  $I$  into ChatGPT-4o<sup>1</sup> with the prompt “describe the objects in the image with their generic name and color as prompts in a list.” It produces a text response as a list  $t_1, \dots, t_n$  of text descriptions of objects. E.g., in the example shown in Fig. 2, it returns “Blue Bowl, Tape, Banana” and so on.

Next, we obtain an image mask for each object. For each text description  $t_i$ , we prompt a grounded segmentation model (in our implementation, GroundedSam2 [14]) using  $t_i$  on image  $I$  and obtain a candidate set of pairs of masks and confidence values. For example, the prompt “a pear” might return multiple useful masks in a scene with two pears; in other cases, some of the masks will be unhelpful (and hopefully low-confidence). We use the confidence values to greedily select a set of non-overlapping masks, and associate each mask with the text prompt that generated it, producing the set  $(R_1, T_1), \dots, (R_N, T_N)$ ,  $R$  and  $T$  denoting the masks and prompts, respectively, for  $N$  objects in the scene.

#### B. Image Inpainting

As illustrated in the Fig. 3, the objects represented by the masks in  $(R_1, T_1), \dots, (R_N, T_N)$  could be partially occluded by other objects in the scene. In this step, we use an image inpainting algorithm to fill in the occluded parts of each object’s image. This is important for the next step of predicting a 3D mesh, because the image-to-3D models only performs reliably with a complete view of the object. Fig. 4(a) illustrates the poor results of attempting reconstruction from an incomplete image.

<sup>1</sup><https://openai.com/index/hello-gpt-4o/>

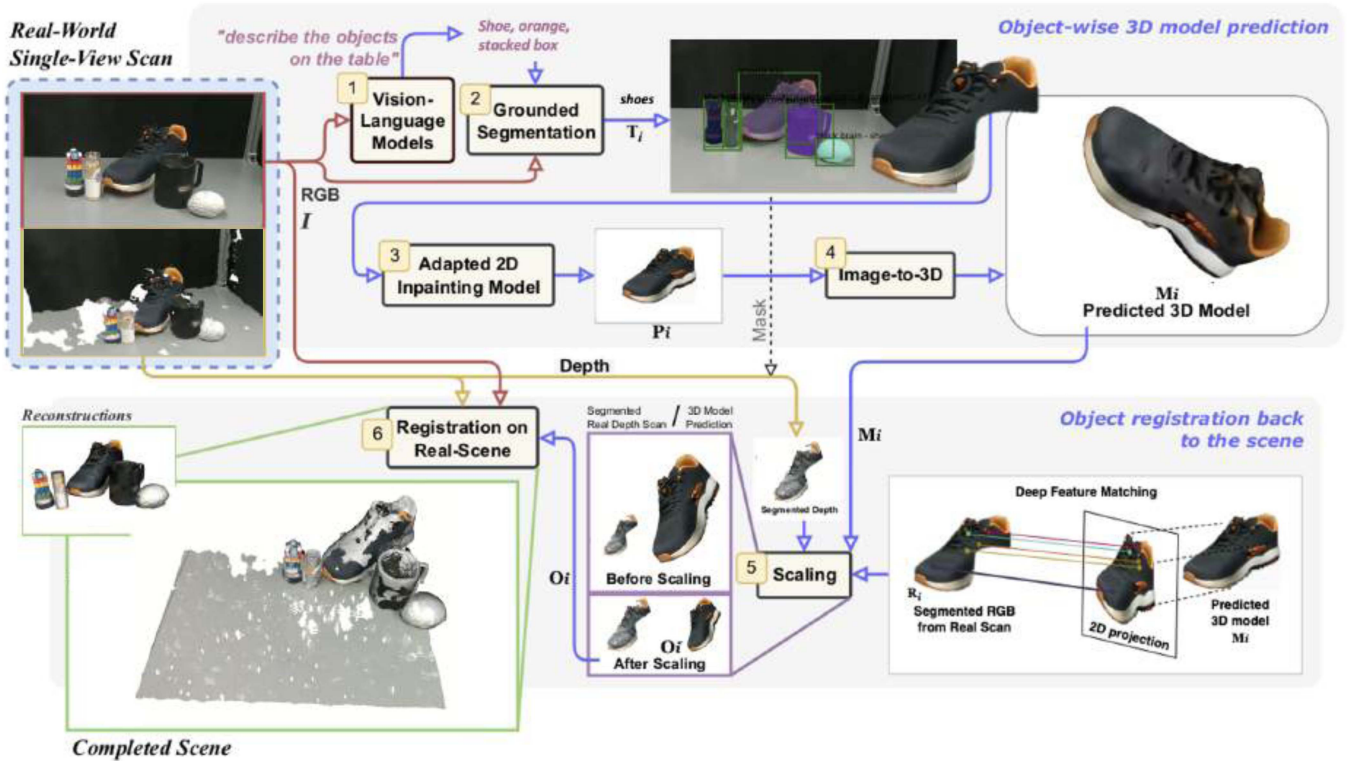


Fig. 2. **Overview of the SceneComplete pipeline.** Starting from a single RGB-D input, the system produces a set of object meshes registered with the input 3D scan, yielding a complete 3D scene reconstruction. The pipeline consists of six key phases: (1) An RGB image is fed into a VLM to enumerate and describe objects, (2) object descriptions and the RGB image are processed by a grounded segmentation model to generate object masks, (3) occluded regions are completed via image inpainting model adapted to output single fully observable objects on a white background, (4) the inpainted 2D images are passed into an image-to-3D model to produce object meshes, (5) object meshes are scaled according to the segmented partial point cloud, and (6) mesh poses are adjusted within the 3D coordinate frame of the original scan using 6DOF pose estimation. Each step leverages pre-trained open world large vision models, enabling scalability and benefiting from future model improvements.



Fig. 3. In the image inpainting module, occluded objects (blue borders) are transformed into single fully observable objects.

In our implementation we begin with BrushNet [15], which takes an image with explicitly masked out regions and a text prompt, and produces a completed image, with the masked portions filled in. So for each  $(R_i, T_i)$  pair, we begin by constructing an image  $I_i$  with just the segmented object region  $R_i$  on a white background. Then we construct an inpainting mask consisting of the union of the regions  $R_j$  for  $j \neq i$ . Finally we query Brushnet with  $I_i, \cup_{j \neq i} R_j$ , and  $T_i$  and obtain an inpainted image  $P_i$  of the completed object.

However, we observed that, out of the box, BrushNet occasionally synthesized additional objects in the occluded areas, as shown in Fig. 4(b), possibly due to the “artistic” data set on which it was trained. To improve this behavior, we adapt BrushNet to full, single objects. It is important to note that, even though we want to adapt BrushNet to **single fully observed** objects, we want to retain the open world capabilities of the model. To achieve this, we use Low-Rank Adaptation (LoRA) [16] on its learnable layers with the PEFT method,

targeting domain-specific improvements on the tabletop YCB dataset [17]. LoRA is an effective method for adapting pretrained models to domain-specific outputs while retaining the inherent generalization of the models. To perform this adaptation, we project the different 3D YCB object meshes from arbitrary poses on a white background and add random brush masks for inpainting as suggested by BrushNet. The adapted model allows us to reliably inpaint individual occluded objects, including those from categories not in the adaptation dataset.

### C. Image-to-3D Models for Object Reconstruction

At this point, we have, for each object, a fully observed object image  $P_i$  on a white background. The next step is to generate a 3D mesh model for the object. Although methods exist for operating directly on the point-cloud generated from the depth channel of the RGB-D image, the depth information is often very low-resolution and noisy, so these depth-only models tend to be highly tuned to specific categories [7], [18], [19], [20] and viewpoints, or operate over idealized sensory input. In recent years, substantial improvements have been made in RGB-only methods, which take advantage of the high quality of the RGB signal and the enormous amounts of available training data [21], [22], [23].

For RGB image as inputs, one option is to use methods that optimize a 3D mesh through differentiable rendering [24], [25], [26]. Although these models have impressive open world results, they have substantial computational cost. On the other hand,




Fig. 4. (a) The impact of inpainting on image-to-3D reconstruction. Without inpainting (top), the image-to-3D model generates incomplete meshes. Inpainting (bottom) fills in occluded parts, producing accurate 3D reconstructions. (b) Comparison of inpainting models. Unadapted BrushNet (middle) introduces artifacts, while the adapted version (right) inpaints occluded parts correctly producing a fully observed object.

feed-forward methods such as InstantMesh [23] that directly map a single-view RGB image into a 3D mesh are highly accurate, open world, and computationally inexpensive at the time of inference. For these reasons, we use InstantMesh, providing each image  $P_i$  as input and obtaining a complete textured 3D mesh  $M_i$  as output. The resulting meshes are produced in an arbitrary orientation, at an arbitrary scale, so more work remains to be done, as explained in the subsequent subsections.

#### D. Mesh Scaling Using Dense Correspondence Matching

The next step is to rescale the meshes  $M_i$ , using a point-cloud constructed from the region  $R_i$  to determine the scale factor as follows: (i) Using the default viewpoint  $v$  from InstantMesh, generate an image  $V_i$  as a 2D projection of mesh  $M_i$ . (ii) Following [27], find dense visual descriptors of the segmented original image,  $R_i$  and the projected image,  $V_i$ , using a pre-trained vision transformer. In our observation,  $v$  is usually close to the viewpoint from which  $R_i$  was rendered. This enables us to obtain matching visual descriptors across  $R_i$  and  $V_i$ . (iii) Generate dense pixel-wise correspondences between the descriptor images, so that we have a set of pairs  $(R_i^j, V_i^j)$  where  $R_i^j$  is a pixel in the original image of the object and  $V_i^j$  is a pixel in the synthesized view. (iv) Map these correspondences into 3D, to obtain a pair  $(R_i^j, V_i^j)$  of points in 3D. (Note that these point sets may not yet be aligned in 3D—we address that problem in the next step.) (v) Center each resulting 3D cloud, compute the average Euclidean distance from the points to the centroid, and compute the ratio of these values as the scale factor. (vi) Apply the scale factor to the mesh  $M_i$  to put it in the same scale as the original point cloud to obtain  $O_i$  as shown in Fig. 2-(5).

#### E. 6D Object Pose Estimation and Registration

Now, we have, for each object, an appropriately scaled mesh  $O_i$ , and we need to reconstruct the entire scene. To do this, we need to find a 6DOF transform for each mesh that causes it to register well with the observed point cloud. For this task, we use FoundationPose [28], a robust object-pose estimation method designed to operate without being limited to specific object categories. We use its model-based mode, which takes as input a partial point-cloud derived from region  $R_i$  of the input RGB-D image and an appropriately scaled textured object mesh  $O_i$ , and returns a 6DOF transform  $\tau_i$  mapping  $O_i$  into the coordinate frame of the point cloud.

As a result of this process, we have a set of pairs  $(O_i, \tau_i)$ , where each  $O_i$  is a textured complete mesh for object  $i$  and  $\tau_i$  is

a pose for that mesh in the camera coordinate frame. This scene reconstruction provides a highly general representation for a wide variety of downstream object-manipulation tasks. Accurate reconstructions of unobserved parts of objects enables a wide variety of manipulation operations, including many types of robust grasping, moving safely in cluttered but unobserved parts of the scene, moving safely when holding a grasped object, etc.

## IV. EXPERIMENTS

We evaluate SceneComplete through three main experimental regimes:

- **Scene Reconstruction and Grasping:** We compare SceneComplete against an existing single-view scene reconstruction method on a large-scale dataset of tabletop scenes, GraspNet-1B [1], captured using a RealSense D435 camera. We also evaluate how the reconstructed scenes contribute in generating collision-free grasps.
- **Object-grasping and Dexterous Manipulation:** We evaluate the effectiveness of SceneComplete in generating grasps for successfully picking up objects inside a simulation environment using a parallel-jaw gripper. We also demonstrate that the reconstructed object models are of sufficient fidelity to enable dexterous grasp proposals for a multi-fingered hand, which depends on having a good estimate of the entire object shape.
- **Real-world Evaluations:** To assess the real-world applicability of SceneComplete, we conduct pick-and-place experiments on a physical robot on a smaller-scale dataset of tabletop scenes collected in our lab (also using a RealSense D435 camera), that include everyday objects which are less likely to have appeared in the training distribution of any of the models used. We evaluate across 15 scenes, each with 4 to 6 objects.

We observe that the runtime of the current implementation of SceneComplete is about 20s per object on a single NVIDIA RTX 4090 GPU; we expect this to improve with advances in models and GPU architectures.

#### A. Scene Reconstruction and Grasping

We first evaluate SceneComplete on the task of reconstructing tabletop scenes from just a single RGB-D image.

**Dataset:** We use the GraspNet-1B dataset [1] for our evaluation. This dataset is particularly suitable for our setting, due to its large collection of 190 cluttered tabletop scenes, featuring




Fig. 5. **Qualitative comparisons** of scene reconstructions on the GraspNet-1B dataset. For each scene we show, the input RGB-D image, OctMAE reconstruction (rendered as normal maps as it predicts scene-level occupancy values), ZeroGrasp reconstruction (rendered as normal maps), our reconstruction (visualized as individually reconstructed object meshes color-matched to the ground truth), and ground-truth object meshes. Highlighted regions indicate missing area (black) or spurious region connecting distinct objects (red).

88 unique objects in various configurations. It includes ground-truth 3D object models and poses for each scene, as well as real RGB-D images.

**Baseline:** Our primary baselines for comparison are OctMAE [8] and ZeroGrasp [9]. OctMAE performs reconstruction by combining octree-based representation with a 3D Masked AutoEncoder (MAE). ZeroGrasp performs simultaneous 3D reconstruction and 6D grasp pose prediction using an octree-based CVAE. Both methods expect a single-view RGB-D image along with a corresponding foreground mask as input, and output the reconstructed scene. For our experiments, we provide OctMAE and ZeroGrasp with the ground-truth segmentation masks. We also report numbers against the partial point cloud (PartialDecomp).

**Metrics:** To measure 3D reconstruction quality and fidelity, we report well-known 3D metrics such as *Chamfer distance* (*CD*) and *Earth Mover’s distance-maximum mean discrepancy* (*MMD-EMD*) metric, similar to existing works [8], [19], [29]. Both *CD* and *MMD-EMD* metrics expect pointclouds as input. For SceneComplete, we sample points uniformly from the reconstructed object meshes. Since OctMAE predicts occupancy values, we use the reconstructed point cloud produced by its occupancy values, normal vectors, and SDF. We also report the *Mesh Intersection-over-Union* (*MIoU*) metric, which is based on comparing the ground-truth meshes with the reconstructed meshes. Specifically, let  $U^*$  be the union of volumes enclosed by the ground truth object meshes, and let  $\hat{U}$  be the union of the volumes enclosed by the meshes produced by a reconstruction algorithm. Then the intersection-over-union metric between the meshes is

$$\text{MIoU}(U^*, \hat{U}) = \frac{U^* \cap \hat{U}}{U^* \cup \hat{U}}.$$

*MIoU* explicitly penalizes both under-reconstructions (missing parts), over-reconstructions (excess geometry), and registration errors. To compute *MIoU*, we make the reconstructed meshes produced by SceneComplete, ZeroGrasp, and OctMAE watertight using ManifoldPlus [30].

**Collision-free Grasping:** We perform a basic test of the utility of SceneComplete on an important downstream task of grasping, by using an antipodal grasp generation method [3] to generate

TABLE I  
COMPARISON OF SHAPE RECONSTRUCTION METHODS ON GRASPNET-1B. HIGHER *MIoU* INDICATES BETTER SHAPE FIDELITY. LOWER *CD*, *MMD-EMD*, AND *GC* INDICATE MORE ACCURATE AND FEASIBLE RECONSTRUCTIONS FOR DOWNSTREAM GRASPING. *CD* AND *MMD-EMD* ARE SCALED BY  $10^4$  AND  $10^2$  RESPECTIVELY

	Reconstruction			Grasping
	<i>MIoU</i> ↑	<i>CD</i> ↓	<i>MMD-EMD</i> ↓	<i>GC</i> ↓
PartialDecomp	0.166	3.16	3.32	53.5
OctMAE [8]	0.445	1.73	3.11	20.3
ZeroGrasp [9]	0.440	1.86	3.07	18.9
SceneComplete	<b>0.478</b>	<b>1.54</b>	<b>3.06</b>	<b>16.4</b>

grasps on the objects in the reconstructed scene. To illustrate the importance of whole-scene reconstruction on grasping, we ask the question: of these potential grasps, which ones are *in collision with the ground-truth scene*? Concretely, for a reconstructed scene, we (a) sample a set  $G$  of collision-free grasps using antipodal sampling, (b) evaluate grasps in  $G$  that would cause a collision in the ground truth scene to obtain a subset  $G'$ , and compute the *Grasp Collision metric*  $GC$  as  $(|G'|/|G|)$  which is the percentage of grasps that the reconstruction would allow, that in fact collide, similar to the metric adopted by [19], [31]. In our case,  $G$  is set to 40.

**Results:** As shown in Table I, our method outperforms the baseline methods in all metrics. We also visualize the comparisons in Fig. 5. While OctMAE produces visually plausible reconstructions, it sometimes fails to recover parts of objects that are occluded by other objects, or are partially observable due to the viewpoint. This also results in higher grasp collisions as shown in Table I, while our method recovers such missing regions. Moreover, since OctMAE directly predicts a scene-level reconstruction, it often hallucinates geometry connecting distinct objects (highlighted in Fig. 5), which results in a lower *MIoU*. In contrast, our object-centric approach reconstructs each object individually, preserving a clear separation between them. We find that on average, only **16%** of grasps generated by SceneComplete are in collision with real objects, as opposed to 20% of those generated by OctMAE, 19% of those generated by ZeroGrasp, and 53% of those generated by PartialDecomp.

## B. Object Grasping and Dexterous Manipulation

We next evaluate the utility of SceneComplete for object grasping and dexterous manipulation in simulation.

**Dataset:** We use a subset of the YCB-Video [2] dataset consisting of 30 cluttered tabletop scenes with 4 to 5 objects per scene from the YCB [17] dataset for simulated grasping experiments using a parallel jaw gripper. For dexterous manipulation, we evaluate on 20 object instances chosen arbitrarily from the GraspNet-1B [1] dataset.

**Baseline:** We compare against the raw input partial point cloud (PartialDecomp) as the baseline. Since this experiment is object-centric, it's not compatible with OctMAE, which produces scene-level reconstruction without per-object separation. While instance segmentation could in principle be used to individuate objects, they do not reliably recover the extent or shape of each reconstructed object, making fair comparison infeasible for object-centric manipulation tasks. We show comparisons with OctMAE on a scene-level grasping task in the next section.

**Metrics:** To assess the effectiveness of SceneComplete for object manipulation, we evaluate the *Grasp Success Rate (GSR)* using simulated grasp attempts (a common metric adopted by prior works [32], [33]) in Isaac Gym. Our scene consists of a Franka Emika Panda arm with a parallel jaw gripper, and scenes from the YCB-V dataset [2]. We compare the GSR achieved by SceneComplete against PartialDecomp and adopt two established and distinct grasping methods for generating grasp proposals:

- Antipodal grasping using [3]: This method samples antipodal points to generate collision-free grasps on the object meshes. To make the objects compatible with simulation and to improve efficiency, we perform an approximate convex decomposition of the watertight object meshes into convex hulls using CoACD [34].
- Contact-GraspNet [4]: This method generates grasp proposals directly from point clouds and was trained to predict grasps from partial observations. We generate grasp proposals on both the reconstructed point cloud (from SceneComplete) and the raw partial point cloud, and evaluate these grasps on the ground-truth scenes.

For each grasping method, we first generate candidate grasps on both SceneComplete's reconstructed scene and PartialDecomp. For each grasp, we simulate a pick attempt on the reconstructed meshes in Isaac Gym and retain only those grasps for which the object remains securely held in the gripper. From these filtered grasps, we randomly select up to 40 grasps per object and evaluate them on the ground-truth object meshes in Isaac Gym. Each evaluation consists of successfully picking up the object from its computed grasp and holding it in the air without dropping, as shown in Fig. 6. GSR is computed as the proportion of successful grasp attempts out of the 40 evaluated grasps. Objects that do not yield valid grasps due to constraints such as exceeding the gripper width are excluded from our evaluation.

**Dexterous Grasping and Stability:** A significant test of the utility of our approach is whether the object reconstructions support the computation of good dexterous grasps for a multi-fingered hand. We evaluate object reconstruction as:

- Pass a *reconstructed* object mesh  $O_i$  into DexGraspNet [5], configured for a dexterous hand (we used both Shadow and Allegro hands), to obtain a grasp  $g_i$ .
- Instantiate Isaac Gym with the selected hand and the *ground truth* object mesh.




Fig. 6. Evaluating GSR on the YCB-V dataset in Isaac Gym. We show a close-up of 3 distinct objects being picked up.



Fig. 7. We demonstrate dexterous grasps using both Shadow Hands and Allegro Hands [5] on objects from the GraspNet-1B dataset, highlighting improved manipulation of complex objects with complete 3D reconstructions.

- Similar to other methods that evaluate dexterous grasping [5], [35], we lift the hand and rotate it within the simulation, and detect whether the object is dropped using PhysX as the physics engine. We visualize dexterous grasps on representative objects using both Shadow and Allegro hands in Fig. 7.

We calculate the percentage of such tests that succeed. For evaluation, we selected 20 objects from the GraspNet-1B dataset and evaluated SceneComplete against PartialDecomp on ground-truth objects—an important upper bound illustrating the reliability of DexGraspNet for selecting such grasps.

**Results:** As shown in Table II, reconstructing scenes with SceneComplete significantly improves grasp success rates in simulation using a parallel jaw gripper. Across both grasping methods, SceneComplete achieves over **twice** the number of successful grasps compared to those from partial input alone. For dexterous grasping, we observe that the number of stable grasps sampled by DexGraspNet varies with object geometry. On average, we evaluate up to 100 random grasps per object and find that SceneComplete enables **twice** as many valid dexterous grasps compared to PartialDecomp. These improvements show that reconstructing scenes enables more reliable grasping and manipulation in cluttered environments.

## C. Real Robot Experiments

We validate the real-world applicability of SceneComplete through experiments on a physical robot. Our experimental setup includes a Franka Emika Panda arm equipped with a wrist-mounted Intel RealSense D435 camera.

**Dataset:** We evaluate our method on 15 distinct tabletop scenes collected in our lab, each containing 4 to 6 everyday objects, specifically chosen to minimize their likelihood of appearing in the training distribution of the methods.

**Baseline:** We compare against PartialDecomp and OctMAE in our evaluation.

TABLE II  
OVERALL GRASP SUCCESS RATE (GSR $\uparrow$ ) ON THE YCB-V DATASET FOR TWO GRASPING METHODS

	Contact-GraspNet GSR	Antipodal Grasping GSR	Overall GSR $\uparrow$
PartialDecomp	$0.46 \pm 0.34$	$0.17 \pm 0.13$	0.32
SceneComplete	<b><math>0.81 \pm 0.2</math></b>	<b><math>0.73 \pm 0.18</math></b>	<b>0.77</b>

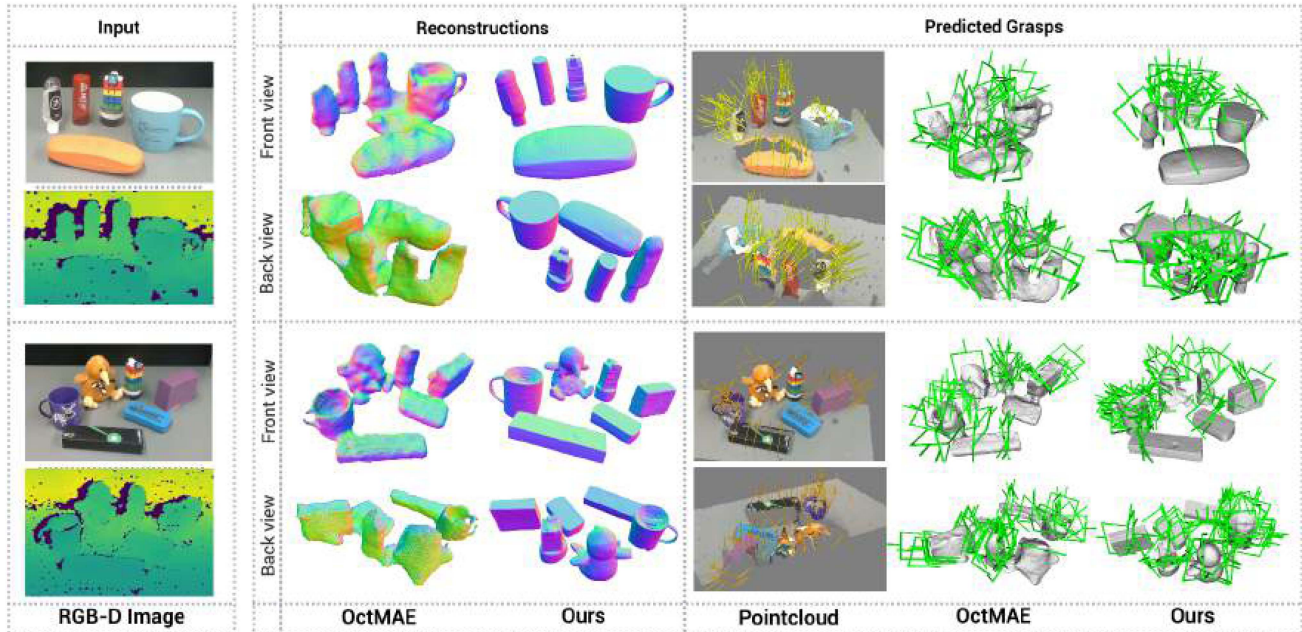


Fig. 8. **Qualitative comparisons** of scene reconstructions on the scans collected in our lab. For each scene we show, the input RGB-D image, OctMAE reconstructions, our reconstructions, grasp proposals on input partial point cloud, grasp proposals on OctMAE’s reconstructions, and grasp proposals on our reconstructions. We show the scene from a front and back viewpoint.

TABLE III  
REAL ROBOT SUCCESS RATE (%) MEASURED AS THE PERCENTAGE OF OBJECTS PICKED SUCCESSFULLY BY THE REAL ROBOT FOR EACH METHOD

	Partial	OctMAE [8]	SceneComplete
<b>Success Rate<math>\uparrow</math></b>	$36.7 \pm 9.9$	$59.6 \pm 15.3$	<b><math>73.3 \pm 15.2</math></b>

**Metrics:** For each scene, we capture an initial RGB-D image using the wrist-mounted camera and randomly select and execute kinematically feasible and collision-free grasps on each object in the scene. We measure the success rate as the percentage of objects that were successfully picked up and report the results in Table III.

**Results:** Our method achieves a success rate of  $\sim 73\%$ , significantly outperforming OctMAE and grasps generated directly from the input partial point cloud alone. We show qualitative comparisons with OctMAE in Fig. 8 and additional results in the supplementary. OctMAE struggles to produce plausible reconstructions on these everyday objects, often hallucinating geometry between distinct objects, leading to no valid grasps being generated for some objects. SceneComplete on the other hand, produces object reconstructions that align closely with the ground truth scene. We note that for SceneComplete, most grasping failures occur due to inaccuracies in the estimated object size, which subsequently leads to errors in registration. In

general, SceneComplete allows robust manipulation of objects in cluttered real-world settings.

## V. DISCUSSION

*a) Limitations* Although our results are very promising, there are of course many failure modes in a composition of so many modules, which can have cascading effects on overall system performance. We outline some failure modes and opportunities for improvement. **Prompting and segmentation:** Occasionally the VLM fails to detect some of the object(s) in the image. Tuning the prompt mitigates the problem, as can prompting multiple ways for multiple hypotheses. **Segmentation:** Grounded-SAM occasionally segments parts of an object along with the full object which leads to multiple reconstruction hypotheses for the same object. We mitigate this partially using IoU-based de-duplication. **Inpainting:** Our current inpainting strategy operates on a relatively isolated object, which removes some important context. We mitigate this by slightly increasing the bounding box and adapting the model, but there is room for improvement. **Image-to-3D:** Although remarkable, these models can sometimes fail to generate plausible reconstructions when given images from highly unusual viewpoints. **Scaling and registration:** Our scaling method is naive and would be improved in some cases by making it non-isotropic. Registration sometimes fails on uniformly-textured objects, where it is difficult to find distinctive features. **Seed value:** The performance

of the image-to-3D and inpainting models may sometimes vary depending on the seed value selected.

*b) Conclusion & Future Work* We have presented a system that solves full-scene reconstruction from a single real-world RGB-D input in cluttered, occluded scenes with no assumptions about object categories. We have built on an incredibly strong foundation of existing general-purpose open-domain perception models and believe our approach will be able to adapt to and profit from future advances in such models. One goal of this letter is to emphasize the importance of this problem for robot manipulation in real open-world environments and to encourage others to propose alternative solution strategies. We hope that overall advances in scene understanding in realistic manipulation settings will enable much more robust and capable robot manipulation systems. One important strategy for making the system less error-prone is to move to a more generative setting with quantified uncertainty. If each module could generate multiple hypotheses, conditioned on its inputs, it would be possible to search for an interpretation that is collectively high-probability for all the modules.

#### REFERENCES

- [1] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "GraspNet-1billion: A large-scale benchmark for general object grasping," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11444–11453.
- [2] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes," *Robot.: Sci. Syst. (RSS)*, 2018.
- [3] J. Mahler et al., "Learning ambidextrous robot grasping policies," *Sci. Robot.*, vol. 4, no. 26, 2019, Art. no. eaau4984.
- [4] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-DoF grasp generation in cluttered scenes," in *Proc. 2021 IEEE Int. Conf. Robot. Automat.*, 2021, pp. 13438–13444.
- [5] R. Wang et al., "DexGraspNet: A large-scale robotic dexterous grasp dataset for general objects based on simulation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 11359–11366.
- [6] M. Lunayach, S. Zakharov, D. Chen, R. Ambrus, Z. Kira, and M. Z. Irshad, "FSD: Fast self-supervised single RGB-D to categorical 3D objects," in *Proc. 2024 IEEE Int. Conf. Robot. Automat.*, 2024, pp. 14630–14637.
- [7] M. Z. Irshad, S. Zakharov, R. Ambrus, T. Kollar, Z. Kira, and A. Gaidon, "ShAPO: Implicit representations for multi-object shape, appearance, and pose optimization," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 275–292.
- [8] S. Iwase et al., "Zero-shot multi-object scene completion," in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 96–113.
- [9] S. Iwase et al., "ZeroGrasp: Zero-shot shape reconstruction enabled robotic grasping," in *Proc. Comput. Vis. Pattern Recognit. Conf.*, 2025, pp. 17405–17415.
- [10] J. Shi, R. Talak, H. Zhang, D. Jin, and L. Carlone, "CRISP: Object pose and shape estimation with test-time adaptation," in *Proc. Comput. Vis. Pattern Recognit. Conf.*, 2025, pp. 11644–11653.
- [11] K. Yao et al., "Cast: Component-aligned 3D scene reconstruction from an rgb image," *ACM Trans. Graph.*, vol. 44, no. 4, pp. 1–19, 2025.
- [12] A. Ardelean, M. Özer, and B. Egger, "Gen3DSR: Generalizable 3D scene reconstruction via divide and conquer from a single view," in *Proc. Int. Conf. 3D Vis. (3DV)*, 2025.
- [13] Y. Ding et al., "Open6DOR: Benchmarking open-instruction 6-DoF object rearrangement and a VLM-based approach," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024, pp. 7359–7366.
- [14] T. Ren et al., "Grounded SAM: Assembling open-world models for diverse visual tasks," 2024, *arXiv:2401.14159*.
- [15] X. Ju, X. Liu, X. Wang, Y. Bian, Y. Shan, and Q. Xu, "BrushNet: A plug-and-play image inpainting model with decomposed dual-branch diffusion," 2024, *arXiv:2403.06976*.
- [16] E. J. Hu et al., "Lora: Low-rank adaptation of large language models," in *Proc. Int. Conf. Learn. Representations*, 2022.
- [17] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The YCB object and model set and benchmarking protocols," 2015, *arXiv:1502.03143*.
- [18] J. Zhang et al., "Unsupervised 3D shape completion through gan inversion," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 1768–1777.
- [19] B. Sen, A. Agarwal, G. Singh, B. Brojeshwar, S. Sridhar, and M. Krishna, "Scarp: 3D shape completion in arbitrary poses for improved grasping," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 3838–3845.
- [20] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, "Convolutional occupancy networks," in *Proc. 16th Eur. Conf. Comput. Vis.*, Glasgow, U.K., 2020, pp. 523–540.
- [21] R. Liu, R. Wu, B. V. Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick, "Zero-1-to-3: Zero-shot one image to 3D object," in *Proc. 2023 IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 9298–9309.
- [22] M. Liu et al., "One-2-3-4-5: Any single image to 3D mesh in 45 seconds without per-shape optimization," in *Proc. Adv. Neural Informat. Process. Syst.*, 2024, pp. 22226–22246.
- [23] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan, "InstantMesh: Efficient 3D mesh generation from a single image with sparse-view large reconstruction models," 2024, *arXiv:2404.07191*.
- [24] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, "DreamFusion: Text-to-3D using 2D diffusion," 2022, *arXiv:2209.14988*.
- [25] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich, "Score jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation," in *Proc. 2023 IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 12619–12629.
- [26] J. Seo et al., "Let 2D diffusion model know 3D-consistency for robust text-to-3D generation," in *Proc. 12th Int. Conf. Learn. Representations*, 2024.
- [27] S. Amir, Y. Gandelsman, S. Bagon, and T. Dekel, "Deep vit features as dense visual descriptors," *ECCVW What is Motion For?*, 2022.
- [28] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "FoundationPose: Unified 6D pose estimation and tracking of novel objects," in *Proc. 2024 IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 17868–17879.
- [29] D. W. Shu, S. W. Park, and J. Kwon, "3D point cloud generative adversarial network based on tree structured graph convolutions," in *Proc. 2019 IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 3859–3868.
- [30] J. Huang, Y. Zhou, and L. Guibas, "ManifoldPlus: A robust and scalable watertight manifold surface generation method for triangle soups," 2020, *arXiv:2005.11621*.
- [31] J. Carvalho et al., "Grasp diffusion network: Learning grasp generators from partial point clouds with diffusion models in so (3) XR3," 2024, *arXiv:2412.08398*.
- [32] H. Zhang, S. Christen, Z. Fan, O. Hilliges, and J. Song, "GraspXL: Generating grasping motions for diverse objects at scale," in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 386–403.
- [33] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, "Synergies between affordance and geometry: 6-DoF grasp detection via implicit representations," *Robot.: Sci. Syst.*, 2021.
- [34] X. Wei, M. Liu, Z. Ling, and H. Su, "Approximate convex decomposition for 3D meshes with collision-aware concavity and tree search," *ACM Trans. Graph.*, vol. 41, no. 4, pp. 1–18, 2022.
- [35] Y. Shao and C. Xiao, "Bimanual grasp synthesis for dexterous robot hands," *IEEE Robot. Automat. Lett.*, vol. 9, no. 12, pp. 11377–11384, Dec. 2024.