

# RIC: Rotate-Inpaint-Complete for Generalizable Scene Reconstruction

Isaac Kasahara, Shubham Agrawal, Selim Engin, Nikhil Chavan-Dafle, Shuran Song, and Volkan Isler  
Samsung AI Center, New York

**Abstract**—General scene reconstruction refers to the task of estimating the full 3D geometry and texture of a scene containing previously unseen objects. In many practical applications such as AR/VR, autonomous navigation, and robotics, only a single view of the scene may be available, making the scene reconstruction task challenging. In this paper, we present a method for scene reconstruction by structurally breaking the problem into two steps: rendering novel views via inpainting and 2D to 3D scene lifting. Specifically, we leverage the generalization capability of large visual language models (DALL·E 2) to inpaint the missing areas of scene color images rendered from different views. Next, we lift these inpainted images to 3D by predicting normals of the inpainted image and solving for the missing depth values. By predicting for normals instead of depth directly, our method allows for robustness to changes in depth distributions and scale. With rigorous quantitative evaluation, we show that our method outperforms multiple baselines while providing generalization to novel objects and scenes. Code and data is available at <https://samsunglabs.github.io/RIC-project-page/>.

## I. INTRODUCTION

An understanding of 3D scene geometry is essential for many downstream applications. In robotics, it allows for accurate manipulation and motion planning. In the field of augmented reality, it allows for better mapping and rendering to bridge the virtual world to the real world. With smartphones and robots equipped with high-quality depth sensors, 3D scene reconstruction task is becoming feasible.

These depth sensors allow for accurate reconstruction of the observed parts of the scene. However, to reconstruct the unseen parts, we must use prior information conditioned on the observed information. The missing information in the input image combined with the diversity in shapes, sizes, and depth distribution of the household objects presents a major challenge for scene reconstruction in-the-wild. In this paper, we study this problem in a general setting, where the goal is to reconstruct a complex scene with multiple novel objects, given only one RGB-D image of the scene.

We present our method Rotate-Inpaint-Complete (RIC), which predicts both the 3D geometry and the texture of the unseen parts of the scene in the input image by leveraging the inpainting capabilities of large visual-language models. Given an RGB-D image of a scene, first we generate novel views (RGB and depth images) by rotating and then projecting the input scene. Then we use a surface-aware masking method to select regions in the image to allow us to inpaint utilizing the powerful 2D inpainting capabilities of DALL·E 2 [1] for exposing the potential object geometry not visible in the input image. Finally, we optimize the depth images using the input depth values, as well as the occlusion boundaries




Fig. 1: RIC inputs a single RGB-D image and generates complete 3D scene reconstruction with texture (bottom-left). RIC first **R**otates the input RGB-D image to a novel viewpoint, **I**npaints the missing regions using generalizable visual language models, and finally **C**ompletes the depth via normal prediction and optimization. For comparison, input as a point cloud is also shown on the bottom-left.

and normals estimated from the inpainted images. These inpainted and completed novel RGB-D views provide the reconstructed scene geometry as a fused point cloud with associated textures. To mitigate the object hallucination and spatial inconsistency of predictions from DALL·E 2, we use a consistency filtering method to enforce consistency across viewpoints which plays a crucial role for generalizable, yet accurate and robust scene reconstruction.

In short, our main contributions are: *i)* We present an integrated approach for scene completion of unseen objects under occlusion and clutter, by solving the problem through novel view inpainting and 2D to 3D scene lifting. *ii)* We develop a method for selectively inpainting regions in the novel views of the input scene that enables synthesis of consistent 2D geometry. *iii)* We train a 2D to 3D lifting method on the YCB-V [2] dataset and demonstrate the generalization capability to cluttered scenes containing novel household objects and categories.

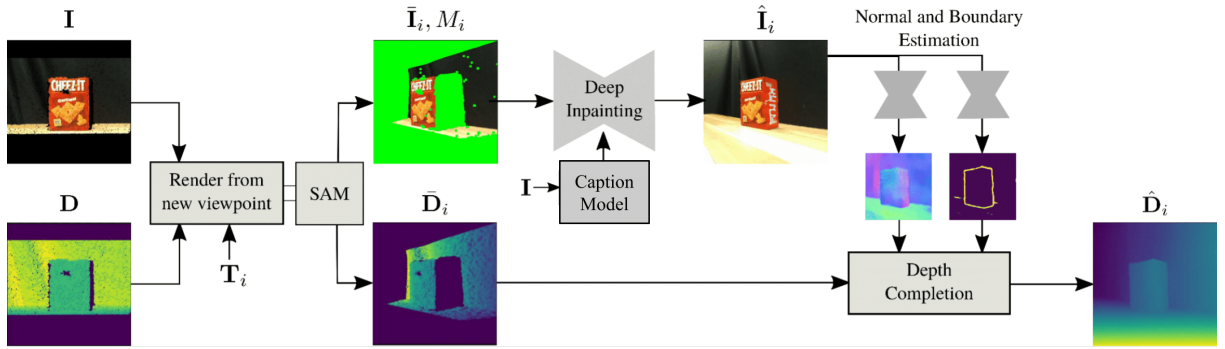


Fig. 2: **Method Overview:** RIC takes as input an RGB-D image and starts by rendering incomplete RGB-D images  $\bar{\mathcal{I}}_i$  and  $\bar{\mathcal{D}}_i$  from a new viewpoint  $\mathbf{T}_i$ . The missing RGB values of  $\bar{\mathcal{I}}_i$  are inpainted using a diffusion-based VLM given a generated prompt, such as “a photo of household objects on a table”, where the pixels to be inpainted are determined by our Surface-Aware Masking (SAM) technique. The inpainted image is used to predict surface normals and occlusion boundaries at the new viewpoint  $\mathbf{T}_i$ , which are then used for completing the missing depth values along with the incomplete depth image  $\bar{\mathcal{D}}_i$ . After repeating this process for  $V$  viewpoints, the final output of RIC is a merge of deprojected depth predictions.

## II. RELATED WORK

**Scene Reconstruction:** While single-object reconstruction is a well-studied problem, full-scene reconstruction is explored in limited settings. Previously works in scene reconstruction were focused either on room scale [3], [4] or in autonomous driving settings [5]–[7] where the scene geometries are usually more structured. In this work, we focus on an object-level scale, specifically tabletop environments, where objects can be in cluttered configurations. While methods like [8]–[10] show an accurate reconstruction of objects at the scene level, they do not generalize to novel category objects. Recently, [11] introduced a method for reconstructing 3D geometries of objects and scenes of unseen categories, and demonstrated generalization capability to objects in-the-wild. However, different from our setting, they mostly focus on isolated objects and scenes with little to no clutter. In contrast, our method can reconstruct geometries and textures of complex scenes with objects from novel categories under heavy occlusions, as we show in our experiments.

**Inpainting:** While traditional inpainting methods made use of hand-crafted image priors to fill small gaps for tasks like image restoration [12], deep generative methods like Generative Adversarial Networks (GAN) [13] have shown remarkable success for tasks like image denoising [14], super-resolution [15], and inpainting [16]–[18]. However, GAN models are known for potentially unstable training for large datasets [19]. More recently, resulting from the growth of visual language diffusion models, which can be efficiently trained on internet-scale datasets, inpainting through image diffusion has shown great generalization capabilities to many different objects and scenes [1]. In this paper, we develop a process to use a visual language diffusion model for inpainting cluttered scenes involving heavy occlusions.

**Text-to-3D Synthesis:** Recent papers such as [20]–[22] have demonstrated the ability to generate 3D models of very diverse objects from merely a text description. Despite the realistic appearance, these generated objects are not grounded to any real-world geometry. To overcome this limitation, [23], [24] extended these methods to reconstruct based on

a ground truth reference image. These papers demonstrate high accuracy on individual objects, but do not demonstrate the ability to reconstruct multiple objects in cluttered scenes. Moreover, their runtime is a concern, as optimizing neural radiance fields [25] can take up to an hour. [26] attempts to produce faster results by optimization without NeRF but are still limited to single object reconstruction.

## III. METHOD

In this section, we present our method Rotate-Inpaint-Complete, or RIC, for generalizable reconstruction of a 3D scene containing multiple objects, given a single RGB-D image of the scene.

RIC takes in as input an RGB-D image  $\mathcal{I} = (\mathbf{I}, \mathbf{D}) \in \mathbb{R}^{H \times W \times 4}$  and outputs a color point cloud  $S \in \mathbb{R}^{N \times (3+3)}$ , where  $N$  is the number of predicted points in the scene. Our method consists of three main components: 1) An inpainting step that takes in an RGB-D image  $\mathcal{I}$  and outputs an inpainted RGB image  $\hat{\mathcal{I}}_i$  from a novel viewpoint  $\mathbf{T}_i \in SE(3)$ . 2) A depth completion component that takes in the inpainted RGB image  $\hat{\mathcal{I}}_i$  as well as an incomplete depth image  $\bar{\mathcal{D}}_i$  rendered from the viewpoint  $\mathbf{T}_i$ , and outputs a completed depth  $\hat{\mathcal{D}}_i$  at that viewpoint. 3) A viewpoint selection and consistency filtering method that utilizes the above two components to generate completed RGB-D images at rotated novel views and uses them to reconstruct the scene. We explain each of these components in detail next.

### A. Inpainting

This section describes the inpainting process, as well as the intermediate steps taken before and after to go from the input image  $\mathcal{I}$  to  $\hat{\mathcal{I}}_i$  at a novel viewpoint  $\mathbf{T}_i$ .

1) *Rotate and Project RGB-D Image:* Given an RGB-D image of a scene and the camera intrinsics, we deproject the image into a point cloud in the camera frame. This point cloud is then projected onto a novel viewpoint  $\mathbf{T}_i$  and the resulting image is masked using our Surface-Aware Masking method (SAM), which we describe in detail in the following section. The projection from this new viewpoint

creates a new RGB-D image  $\bar{I}_i$  with missing RGB and depth information as seen in Figure 2. Small holes of the missing RGB values are filled with a naive inpainting algorithm [27] by inpainting pixels that are covered after a morphological closing operation of kernel size 5 is applied to the mask. The larger missing areas are left for the deep inpainting module.

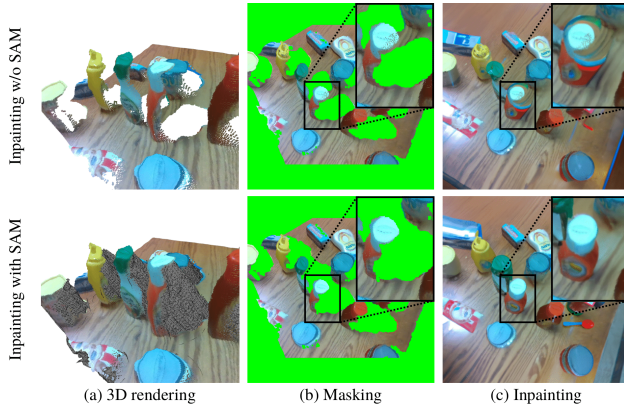


Fig. 3: Surface-Aware Masking (SAM) is a necessary step to obtain realistic inpaintings. Naively rotating the input point cloud moves the background pixels next to the foreground object pixels (b-top) which results in poor inpainting (c-top). Using SAM, we correctly mask out the background pixels which results in good inpainting results (b-bottom).

2) *Surface-Aware Masking*: In order for inpainting to work properly, a mask covering the areas to inpaint needs to be generated. After projecting to the new camera frame, any 3D space possible to be reconstructed needs to be represented as an inpainting mask in the 2D image. This issue can be seen in Figure 3 as the table takes up pixels we may want to fill in with the bottle. To solve this problem, a 3D frustum is generated from the original camera and depth image. For every pixel in the original camera frame, a ray is cast from the camera through each point in the projected point cloud from  $\mathcal{I}$ . Once the ray has passed through its respective point, it is used to generate a list of points along the ray from that depth onward with  $m$  points of equal spacing  $c$ . This is done for every ray, and from this process results a point cloud covering the potential space that the 3D scene could possibly fill. This point cloud is then converted to a mesh, and when the point cloud from the RGB-D image  $\mathcal{I}$  is rotated to novel views, the mesh is rotated with it. Finally, when projecting back to the camera frame after rotation, points that are occluded by the mesh are discarded. Any blank pixels are then used as the 2D inpainting mask to be filled when passed to the inpainting step. This procedure of generating the final image and mask is detailed in our technical report and its outputs are shown in Figure 2, with the green pixels representing the inpainting mask.

3) *Diffusion-based Inpainting*: Once these preprocessing steps have been completed, we pass the processed image and a mask of areas to be filled in to the inpainting algorithm. We use DALL-E 2 [1] for image inpainting since it demonstrates the ability to produce the most realistic results. This model takes in the incomplete image  $\bar{I}_i$ , the mask generated in the

previous step  $M$ , and an input prompt  $P$  that describes the context of the image in words. For prompt, we pass the RGB image  $I$  to a deep captioning model [28] and prefix the generated caption with “A photo of”. We also explore using a more specific and generic prompt in our ablation experiments (Table II). The output from this inpainting method is an image  $\hat{I}_i$  that now contains estimated areas from the diffusion model. Figure 2 shows an example before and after inpainting with DALL-E 2.

## B. Depth Completion

We use a method proposed in [29] for generating a complete depth image  $\hat{D}_i$  from an incomplete depth image  $\bar{D}_i$  and its corresponding RGB image. This method estimates the normals and occlusion boundaries from the RGB image, and optimizes for the complete depth by utilizing the estimated normals, occlusion boundaries, and incomplete depth.

1) *Normals and Occlusion Boundaries Prediction*: In order to obtain estimations for the normals and occlusion boundaries, we train Deeplabv3+ with DRN-D-54 in the same manner as in [30]. The ground truth normals and occlusion boundaries are obtained using the depth images from the YCB-V training dataset [2], the YCB-V synthetic dataset [31], [32], and the HomebrewedDB synthetic dataset [33].

2) *Optimize for Depth*: Given the incomplete depth, the estimated normals from the image, and estimated occlusion boundaries, we solve for the completed depth. The main idea behind this method in [29] is that the areas with missing depth can be computed by tracing along the estimated normals from areas of known depth with the occlusion boundaries acting as barriers where normals should not be traced across. Formally we solve a system of equations to minimize an error  $E$ , where  $E$  is defined as  $E = \lambda_D E_D + \lambda_S E_S + \lambda_N E_N B$ . Here,  $E_D$  is the distance between the ground truth and estimated depth,  $E_S$  influences nearby pixels to have similar depths,  $E_N$  measures the consistency of estimated depth and estimated normal values, and  $B$  weights the normal values based on the probability that it is a boundary. We use the same  $\lambda_D, \lambda_S, \lambda_N$  values as in [30].

## C. Scene Completion

This section describes the complete process we follow to reconstruct a 3D scene from a single RGB-D image.

1) *Viewpoint Selection*: For diffusion-based inpainting, “known” pixels, i.e., the non-masked areas, guide the prediction of the unknown masked areas. We refer to the known pixels as context pixels and define the context ratio  $C$  for any given image as  $C = (\#context\ pixels) / (\#all\ pixels)$ . This ratio gives us some indication about how accurately the inpainting model will be able to fill in the missing areas. With a low  $C$ , many areas are unknown and inpainting will struggle, and with a high  $C$  inpainting will do well but only fill in minimal information. An example of different context ratio values can be seen in Figure 5.

We then design our viewpoint selection process to search for a context ratio that will allow for accurate inpainting. To do this, we define a sphere with a center as the mean of

the input point cloud, and the radius as the distance between the center and the initial camera location. Then from the starting viewing angle, we rotate in various directions along this sphere away from the starting position. At each step in this rotation, we project the input point cloud onto the new camera location. Using this newly projected image, we compute the context  $C$  of the image. If the  $C$  is closest to our chosen context threshold  $C^*$ , we use that viewpoint  $\mathbf{T}_i$  as next to inpaint. We repeat this process for  $V$  evenly spaced directions we traverse along the sphere as visualized in Figure 5, where both  $C^*$  and  $V$  are chosen using the experiment described in Section IV.

2) *Enforcing Consistency Across Viewpoints*: The final step in our method involves combining these generated viewpoints while enforcing consistency across them. One drawback of utilizing DALL-E 2 for inpainting real objects, is its inconsistent completion of objects as well as the hallucination of objects that are not originally in the scene. To combat this issue, we filter for consistent predictions across viewpoints. The final prediction is achieved by first deprojecting the RGB-D images from each viewpoint  $\mathbf{T}_i$  back into the original camera frame as point clouds. We then apply the following *consistency rule* across all the generated points: If a predicted point from one viewpoint has a predicted point within a 1cm radius from at least two other viewpoints we keep that point, otherwise we remove that point from our final prediction. This rule allows us to keep points that only multiple viewpoints predict. We then combine all filtered points to obtain our final output point cloud of the completed scene  $S$ , which contains more accurate geometry and color than without filtering as seen in Table III and Figure 4.

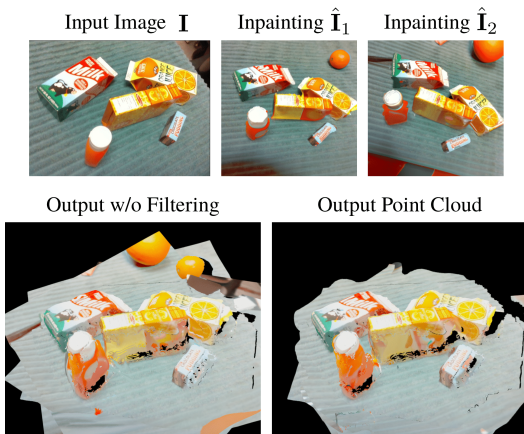


Fig. 4: The consistency filtering step is used to remove the hallucinated objects in 3D, e.g., the oranges in the top right of the images get filtered out. For simplicity we visualize our consistency filtering step for only two viewpoints, while we filter using all viewpoints in our main method.

#### IV. EXPERIMENTS

In this section, we evaluate the performance of RIC for single view RGB-D scene reconstruction task. We also report ablation studies for understanding the dependence of our

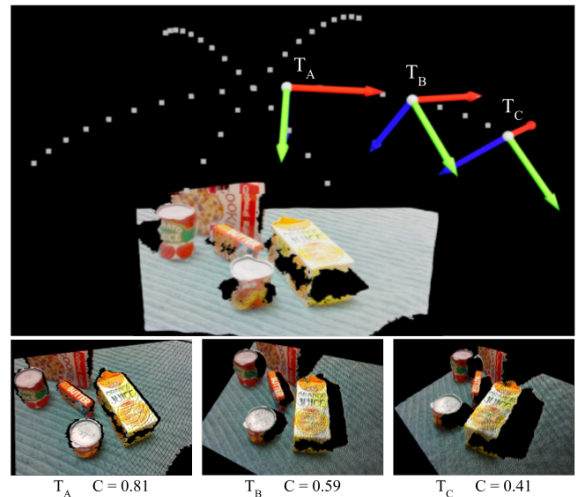


Fig. 5: RIC samples viewpoints on evenly spaced directions along the viewing sphere (white dots). Several viewpoints (top) as well as their corresponding rendered images are shown together with the context ratio of each image (bottom).

method on (1) prompt specificity, (2) inpainting model, and (3) consistency filtering.

**Implementation Details:** The inpainting step of our algorithm is based on OpenAI’s DALL-E 2 API. For our implementation of SAM, we choose a spacing value  $c$  of 0.01 meters with  $m = 100$  points for generating our rays. For choosing the number of viewpoints/viewpoint directions  $V$  as well as the context ratio  $C^*$  described in our method section, we perform a parameter search using 4 held out validation scenes from the YCB-V test set. We test using 6, 8, 10, and 12 viewpoints as well as a value of 0.3, 0.4, 0.5, 0.6, and 0.7 for our context threshold. We found that 10 views and 0.4 as a context threshold gave us the best accuracy on the validation set. 12 views and 0.4 as a context threshold performed similarly, but in the interest of runtime we use 10 for the final method. The full parameter grid search are provided in our technical report. For our module that enforces consistency between synthesized views, we choose a threshold of 0.01 meters when computing the intersection between points in the viewpoints point clouds.

**Datasets:** We trained our depth completion model using the YCB-V training dataset [2]. For testing, we test on 8 unseen scenes from the YCB-V test set, and select 5 RGB-D images from each of the scenes, i.e., we test on a total of 40 RGB-D images in total. For ground truth point clouds, we deproject the RGB-D frames of the scene and concatenate them together. We also place the ground truth meshes in the scene for the objects and convert those to point clouds before concatenating them as well. Finally, we crop this point cloud around the ground truth meshes with a 10cm buffer as the RGB-D frames may contain floors and walls far away that we are not interested in reconstructing. This creates our final ground truth point cloud covering the majority of the scene with full geometry of the objects in the scene.

To demonstrate our model’s capabilities of generalizing to unseen objects and to entirely new datasets, we also

Method	IoU $\uparrow$	F-Score $\uparrow$	CD( $S^*, S$ ) $\downarrow$	CD( $S, S^*$ ) $\downarrow$	CD $\downarrow$
YCB-V [2]					
CON [35]	0.087	0.354	0.036	0.014	0.050
ShellNet [37]	0.224	0.607	0.019	0.012	0.031
CenterSnap [10]	0.225	0.622	0.019	<b>0.009</b>	<b>0.028</b>
RIC (Ours)	<b>0.294</b>	<b>0.661</b>	<b>0.018</b>	0.010	<b>0.028</b>
HOPE [39]					
CON [35]	0.086	0.279	0.094	0.035	0.128
ShellNet [37]	0.185	0.523	0.035	0.013	0.047
CenterSnap [10]	0.180	0.526	0.037	0.006	0.042
RIC (Ours)	<b>0.290</b>	<b>0.649</b>	<b>0.031</b>	<b>0.005</b>	<b>0.036</b>

TABLE I: Comparison of methods for the task of 3D scene completion on YCB-V [2] and HOPE [39]. Higher numbers for the IoU and F-score metrics, and lower numbers for the Chamfer Distances (CD) indicate better performance.

compare our method on the HOPE dataset [34]. HOPE test set only contains individual RGB-D images and is unusable for generating full scene point cloud. Instead, we use HOPE training dataset for evaluation, as the train set contains RGB-D video and cluttered tabletop scenes with novel objects. The dataset has 10 scenes, and we again sample 5 frames per scene. Ground truth point clouds are obtained similarly to the YCB-V dataset.

**Baselines:** We compare RIC against four baselines: Convolutional Occupancy Networks (CON) [35] is a 3D scene reconstruction method that inputs a sparse point cloud. We use their pre-trained model for *Synthetic Indoor Scene dataset* where similar to our YCB-V and HOPE datasets, they place multiple ShapeNet [36] objects in indoor scenes. CoReNet [9] is a multi-object shape estimator that inputs an RGB image and estimates a mesh. We compare against CoReNet’s pre-trained model qualitatively since its predictions lack scale information. ShellNet [37] is trained for single object reconstruction. Given a scene depth image and object instance mask, ShellNet produces reconstruction for the object instance. We re-implemented ShellNet’s architecture and trained it with Mask R-CNN [38] as the segmentation network on YCB-V dataset. Finally, we compare against CenterSnap [10], a multi-object point cloud prediction method. CenterSnap inputs an RGB-D image and predicts point clouds for each object in the scene. Similar to CenterSnap’s original training, we first train it on YCB-V synthetic dataset [31], [32], then fine-tune it on the YCB-V real training dataset [2]. Since CenterSnap and ShellNet only predict the point clouds for objects and not the rest of the scene, for a fair evaluation, we concatenate their outputs with deprojected point cloud from the input RGB-D image.

Table I shows quantitative evaluations on within-training-distribution YCB-V dataset [2] and out-of-training-distribution HOPE dataset [34]. On YCB-V dataset, RIC is able to outperform CON and ShellNet on all 3D scene reconstruction metrics. CON takes as input a sparse point-cloud of the scene. When major parts of the input point clouds are missing, as the common case for single-view RGB-D point clouds, CON fails to infer those regions. ShellNet is trained to predict back-side depth image for the detected object. We notice that with varying viewing directions, ShellNet backside depths are either too thin or too thick resulting

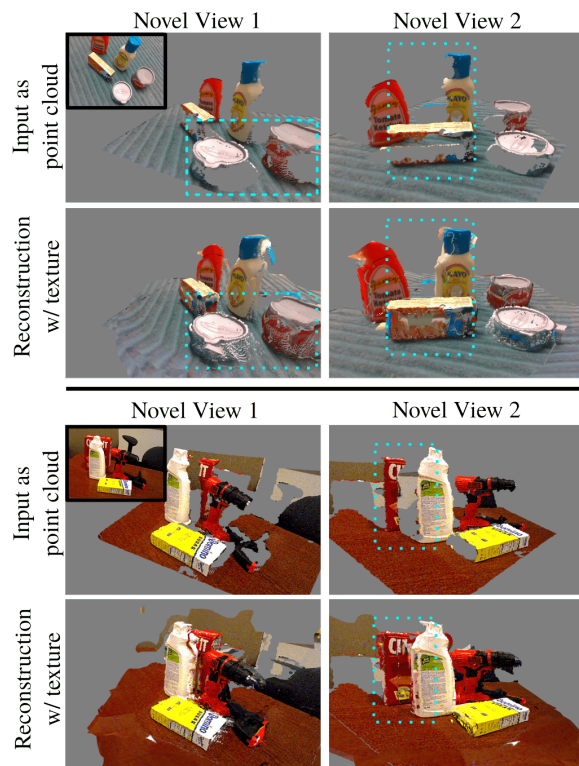


Fig. 6: **Qualitative Results:** We show our scene completion results given a single RGB-D image, as color point clouds from two viewpoints. Top two rows are from the HOPE dataset [39], and the bottom two are from YCB-V [2].

in low performance. MaskRCNN’s failure to detect objects also directly contributed to lower performance for ShellNet. CenterSnap inputs RGB-D image and predicts object shapes via a multi-step procedure allowing CenterSnap to learn strong shape and pose priors for objects within training distribution. This allowed CenterSnap to perform strongly on YCB-V objects as it was trained on them, but we noticed it struggles with cases of occluded objects. RIC which is trained without ground truth object pose or shape supervision is able to match or outperform the baselines in all metrics. Fig. 7 shows a qualitative comparison with baselines.

On the out-of-distribution HOPE dataset, RIC is able to outperform all baselines by an even larger margin. This shows that our normal and occlusion boundary-based depth completion method generalizes well to unseen novel scenes. Figure 6 shows qualitative results on these datasets.

As a byproduct, our method also produces novel views of unseen multi-object scenes from a single RGB-D image. Figure 8 shows our method compared to the ground truth. We show that by combining our masking method with DALL-E 2’s inpainting capability, realistic novels views can be generated for multiple unseen objects.

#### A. Ablation Studies

**Prompt Reliance:** Image diffusion models tend to heavily rely on the input prompt. To test our methods robustness, we performed an experiment using a general prompt (G), “a photo of household objects on a table”, for every scene to

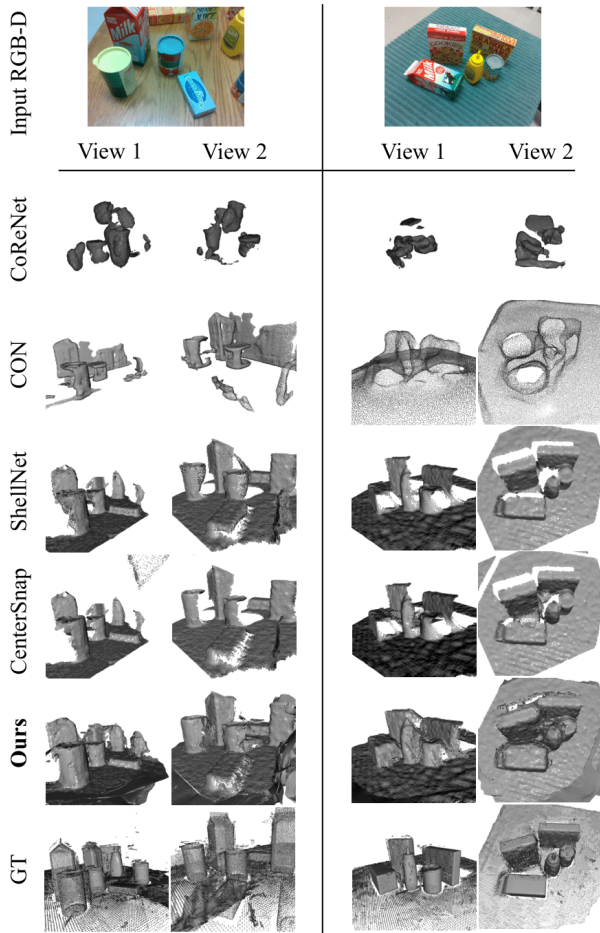


Fig. 7: **Qualitative Comparison:** We compare our method against baselines for completing scene geometry given a single RGB-D image. Views 1 and 2 show novel viewpoints of the predicted point clouds from each method. Our method provides a denser and more complete reconstruction.

Method	IoU $\uparrow$	F-Score $\uparrow$	CD $\downarrow$
RIC (S)	0.262	0.613	0.038
RIC (G)	0.261	0.613	0.038
RIC (Ours)	<b>0.290</b>	<b>0.649</b>	<b>0.036</b>

TABLE II: Prompt specificity results on HOPE dataset [34]: RIC (S) denotes our model with scene specific prompt, RIC (G) uses “household objects on a table” as the prompt for all scenes, and RIC (Ours) uses an image caption generator [28].

see how much performance degrades. We also use a specific prompt (S) where using the ground truth list of objects we list out every object on the table as the prompt. Table II shows that our method does not largely depend on the type of prompt. We hypothesize that our view selection method retains enough surrounding context information required for the inpainting model to inpaint successfully.

**Inpainting Model:** We substitute in Stable Diffusion 2’s [40] inpainting model as an open-source alternative to DALL-E 2 in Table III. We find that while accuracy decreases, it is still a viable inpainting substitute for our method.




Fig. 8: **Qualitative Novel View Results:** As a byproduct of our method we also show qualitative results for generating novel views of scenes from the the YCB-V dataset.

Method	IoU $\uparrow$	F-Score $\uparrow$	CD( $S^*, S$ ) $\downarrow$	CD( $S, S^*$ ) $\downarrow$	CD $\downarrow$
RIC (SD-2)	0.265	0.620	0.033	<b>0.005</b>	0.037
RIC (No Filter)	0.271	0.574	<b>0.017</b>	0.030	0.047
RIC (Ours)	<b>0.290</b>	<b>0.649</b>	0.031	<b>0.005</b>	<b>0.036</b>

TABLE III: Result of swapping out various parts of our method shown on the HOPE [39] dataset. (Ours) utilizes OpenAI’s DALL-E 2 model and our consistency filtering method, (SD-2) uses Stable Diffusion 2’s inpainting model, and (No Filter) refers to our method without filtering.

**Consistency Filtering:** We test our method without applying the consistency filtering step by just combining all predicted viewpoints in Table III. This caused a substantial decrease in accuracy as any hallucinated object is kept in.

## V. DISCUSSION

We presented RIC, a novel method for 3D scene reconstruction. RIC solves the problem of 3D reconstruction of a cluttered scene of novel objects by leveraging the generalization capabilities of large visual language models. More specifically, our method utilizes the 2D inpainting capabilities of DALL-E 2 and generates a coherent set of inpainted views of the scene. It then lifts that information into 3D through a novel geometric multi-step method to finally output the point cloud of the reconstructed scene.

While we demonstrate the effectiveness of our method for generalizable scene completion, we also note that DALL-E 2 can generate unrealistic objects and parts of objects in the inpainted images. We mitigate this issue through various ways described in our method section, but these irregularities can adversely affect the reconstruction quality in a few cases. While RIC shows the ability to complete the front and sides of objects in the scene, the backside of objects is often left incomplete as enough context to accurately reconstruct the scene is not available. At large angles away from the original viewpoint, the inpainting quality degrades due to the large areas of missing information - an exciting yet challenging problem for future work.

## REFERENCES

- [1] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [2] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [3] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017.
- [4] Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2018.
- [5] Ran Cheng, Christopher Agia, Yuan Ren, Xinhai Li, and Liu Bingbing. S3cnet: A sparse semantic scene completion network for lidar point clouds. In *Conference on Robot Learning*, pages 2148–2161. PMLR, 2021.
- [6] Christoph B Rist, David Emmerichs, Markus Enzweiler, and Dariu M Gavrilă. Semantic scene completion using local deep implicit functions on lidar data. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):7205–7218, 2021.
- [7] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3991–4001, 2022.
- [8] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9785–9795, 2019.
- [9] Stefan Popov, Pablo Bauszat, and Vittorio Ferrari. Corenet: Coherent 3d scene reconstruction from a single rgb image. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 366–383. Springer, 2020.
- [10] Muhammad Zubair Irshad, Thomas Kollar, Michael Laskey, Kevin Stone, and Zsolt Kira. Centersnap: Single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10632–10640. IEEE, 2022.
- [11] Chao-Yuan Wu, Justin Johnson, Jitendra Malik, Christoph Feichtenhofer, and Georgia Gkioxari. Multiview compressive coding for 3d reconstruction. *arXiv preprint arXiv:2301.08247*, 2023.
- [12] Omar Elharrouss, Noor Almaadeed, Somaya Al-Maadeed, and Younes Akbari. Image inpainting: A review. *Neural Processing Letters*, 51:2007–2028, 2020.
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [14] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3155–3164, 2018.
- [15] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4681–4690, 2017.
- [16] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [17] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017.
- [18] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. *arXiv preprint arXiv:2103.10428*, 2021.
- [19] Lilian Weng. What are diffusion models? [lilianweng.github.io](https://lilianweng.github.io), Jul 2021.
- [20] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022.
- [21] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [22] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. *arXiv preprint arXiv:2211.10440*, 2022.
- [23] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurlift-360: Lifting an in-the-wild 2d photo to a 3d object with 360° views. *arXiv preprint arXiv:2211.16431*, 2022.
- [24] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Realfusion: 360° reconstruction of any object from a single image. *arXiv preprint arXiv:2302.10663*, 2023.
- [25] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [26] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.
- [27] Alexandru Telea. An image inpainting technique based on the fast marching method. *Journal of Graphics Tools*, 9(1):23–34, 2004.
- [28] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022.
- [29] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [30] Shreeyak Sajjan, Matthew Moore, Mike Pan, Ganesh Nagaraja, Johnny Lee, Andy Zeng, and Shuran Song. Clear grasp: 3d shape estimation of transparent objects for manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3634–3642. IEEE, 2020.
- [31] Maximilian Denninger, Martin Sundermeyer, Dominik Winkelbauer, Dmitry Olefir, Tomas Hodan, Youssef Zidan, Mohamad Elbadrawy, Markus Knauer, Harinandan Katam, and Ahsan Lodhi. BlenderProc: reducing the reality gap with photorealistic rendering. *Robotics: Science and Systems (RSS) Workshops*, 2020.
- [32] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. BOP challenge 2020 on 6D object localization. *European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [33] Roman Kaskman, Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Homebreweddb: Rgb-d dataset for 6d pose estimation of 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- [34] Stephen Tyree, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Jeffrey Smith, and Stan Birchfield. 6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark. In *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [35] Songyu Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 523–540. Springer, 2020.
- [36] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [37] Nikhil Chavan-Dafle, Sergiy Popovych, Shubham Agrawal, Daniel D Lee, and Volkan Isler. Simultaneous object reconstruction and grasp prediction using a camera-centric object shell representation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1396–1403. IEEE, 2022.
- [38] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017.
- [39] Yunzhi Lin, Jonathan Tremblay, Stephen Tyree, Patricio A. Vela, and Stan Birchfield. Multi-view fusion for multi-level robotic scene understanding. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6817–6824, 2021.

- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022.