

# Amodal 3D Reconstruction for Robotic Manipulation via Stability and Connectivity

William Agnew, Christopher Xie, Aaron Walsman, Octavian Murad, Caelen Wang, Pedro Domingos, Siddhartha Srinivasa

University of Washington

{wagnew3, chrisxie, awalsman, ovmurad, wangc21, pedrod, siddh}@cs.washington.edu

**Abstract:** Learning-based 3D object reconstruction enables single- or few-shot estimation of 3D object models. For robotics, this holds the potential to allow model-based methods to rapidly adapt to novel objects and scenes. Existing 3D reconstruction techniques optimize for visual reconstruction fidelity, typically measured by chamfer distance or voxel IOU. We find that when applied to realistic, cluttered robotics environments, these systems produce reconstructions with low physical realism, resulting in poor task performance when used for model-based control. We propose ARM, an amodal 3D reconstruction system that introduces (1) a stability prior over object shapes, (2) a connectivity prior, and (3) a multi-channel input representation that allows for reasoning over relationships between groups of objects. By using these priors over the physical properties of objects, our system improves reconstruction quality not just by standard visual metrics, but also performance of model-based control on a variety of robotics manipulation tasks in challenging, cluttered environments. Code is available at [github.com/wagnew3/ARM](https://github.com/wagnew3/ARM).

**Keywords:** 3D Reconstruction, 3D Vision, Model-Based

## 1 Introduction

Manipulating previously unseen objects is a critical functionality for robots to ubiquitously function in unstructured environments. One solution to this problem is to use methods that do not rely on explicit 3D object models, such as model-free reinforcement learning [1, 2]. However, quickly generalizing learned policies across wide ranges of tasks and objects remains an open problem. On the other hand, obtaining detailed 3D object models can enable robots to physically reason about interactions with them to accomplish robotic tasks. For example, CAD models [3] have extensively been used to detect the 6D pose of objects [4, 5, 6], facilitating many different kinds of manipulation tasks. Such 3D models can also be integrated with high-fidelity physics simulators [7, 8] to provide accurate simulations for planning and learning, enabling model-based methods to generate high-level and/or low-level plans in order to accomplish long-horizon tasks [9, 10, 11]. Unfortunately, these techniques can not be extended to unseen objects without building new models on the fly.

Generalizing interactive robotics problems to previously unseen objects using robust 3D reconstruction is the primary focus of this paper. Rather than rely on a large database of models that have been laboriously hand-crafted or captured using a 3D scanner, we instead focus on techniques that can reconstruct meshes using observations of unseen objects in the robot’s environment. While SLAM methods can reconstruct highly accurate models given many views [12], it can be challenging for these methods to separate objects in clutter and generate faithful reconstructions from a small number of observations. The computer vision community has recently made significant progress in addressing these limitations using neural networks to estimate 3D models from single or few images [13, 14, 15, 16]. In this work, we investigate the use of such methods to solve robotic manipulation tasks involving previously unseen objects (instances and classes).

Unfortunately, we find that directly applying state-of-the-art unseen object reconstruction techniques [14] to cluttered environments frequently fails to reconstruct objects in regions occluded by distract-

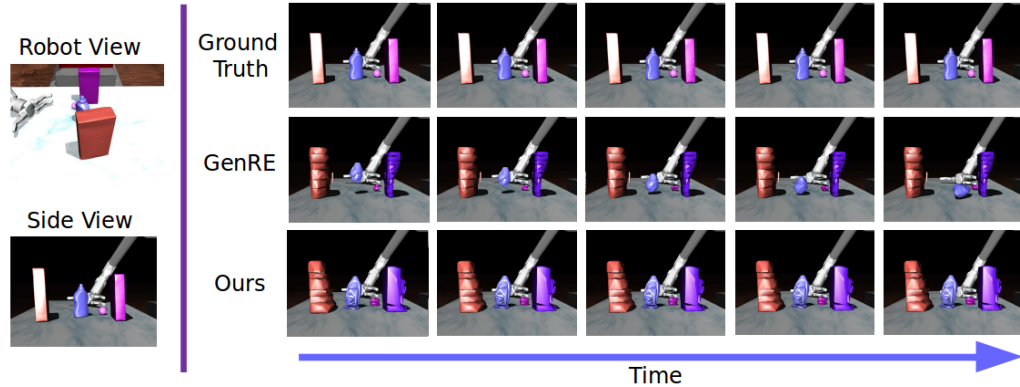


Figure 1: Comparison of physical behaviors of reconstructions from different algorithms. The baseline reconstruction of the light purple occluded mustard bottle is unstable and topples over, while our reconstruction is stable.

tor objects, leading to physically unstable models and poor performance in downstream manipulation problems. This is due to these methods optimizing reconstruction metrics such as Chamfer distance, which are not necessarily relevant for manipulation tasks that utilize the resulting 3D model. Thus, our key insight is to adapt such systems to produce high *physical* fidelity, improving manipulation success rates for unseen objects.

We accomplish this by encouraging the reconstruction network to provide physically realistic outputs. First, we assume that the scene and its objects are stable prior to manipulation, which motivates us to design a novel loss function that penalizes unstable reconstructions. This encourages the network to reconstruct stable scenes (see Figure 1 for an example). Second, as mentioned above, current reconstruction methods struggle to adequately predict occluded portions of objects. This leads to disconnected objects which are not physically realistic. Thus, we design another loss function to penalize disconnectedness of predicted 3D models. Furthermore, both of our novel loss functions are differentiable which allows for end-to-end training. To our knowledge, we are the first to add physical priors on 3D reconstruction. Finally, we introduce a multi-channel voxel representation that allows reasoning over the spatial extent of other objects during the reconstruction phase, and we empirically show that this benefits performance.

We integrate our proposed loss functions into a modular framework to provide Amodal 3D Reconstructions for Robotic Manipulation (ARM). We use the state-of-the-art method, GenRE [14], as our reconstruction method, however we are free to choose any method in place of GenRE as our framework is modular. To evaluate our method, we introduce a challenging cluttered 3D reconstruction benchmark. We empirically demonstrate that ARM improves the reconstruction quality by 28% on this task, and manipulation success rates on unseen objects on a range of challenging tasks including grasping, pushing, and rearrangement by 42% over GenRE.

## 2 Related Work

**3D Reconstruction.** 3D reconstruction is a challenging problem that has been studied for decades. Recently, learning-based methods have provided significant progress when focusing on reconstructing single objects in isolation [17, 18]. Recently, Gkioxari et al. [13] introduces graph neural networks to refine mesh predictions. Kulkarni et al. [19] introduces pairwise object relations and a refinement procedure to improve object pose and shape estimation. Additionally, reconstructing previously unseen classes compounds the difficulty of the problem [20, 21, 14].

Amodal 3D reconstruction is the problem of reconstructing partially visible objects, which is still a relatively unexplored task [22]. Kundu et al. [23] approaches this problem by learning class-specific shape priors from large datasets of CAD models. Kulkarni et al. [19] study amodal reconstruction of scenes from single RGB images, while Suvar et al. [24] handles occlusion by using multiple RGBD views. However, because robot manipulation settings are our desired environment, we re-

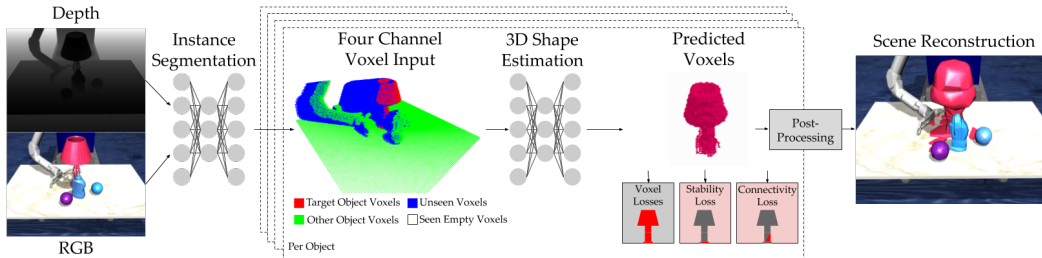


Figure 2: A visual representation of our modular framework.

quire not only amodal reconstruction of objects, but also the ability to deliver physically realistic reconstructions which warrants more informed loss functions including stability and connectivity.

**Exploiting Physics for Scene Understanding** Some works have investigated the use of physics to better inform reconstructions by encouraging physical plausibility. In particular, [25, 26] use a stability prior and Chen et al. [27] use collision and support priors to fit 3D bounding boxes around objects. Our work introduces a differentiable and efficiently computable stability prior to allow generation of stable 3D meshes, rather than just 3D bounding boxes. Additionally, our connectivity prior promotes better reconstruction in occluded regions.

**3D Reconstruction in Robotics** While applying 3D reconstruction to robotics provides an appealing solution to manipulation, few works have investigated this. Such reconstructions can be used to synthesize grasps for single objects using analytic and/or learning-based solutions [28, 29, 30]. Sucar et al. [24] considers grasp synthesis in tabletop scenes with multiple objects of known classes, but does not consider highly cluttered scenes. Most similar to our work, [31, 32] compute grasps for reconstructed objects in cluttered scenes. However, they do not take advantage of physics, which reduces the physical realism. Our work attempts to solve a wider range of manipulation tasks while incorporating physical notions of stability and connectivity to improve performance.

### 3 Amodal 3D Reconstruction

#### 3.1 System Architecture

In this section we describe the architecture of our ARM framework, which consists of four stages. 1) We first apply an instance segmentation network to the input RGB-D image. 2) For each object we detect, we pre-process its point cloud to compute its *four channel voxel input representation*, defined below. 3) ARM uses this representation to perform 3D shape estimation with a deep network, followed by post-processing. 4) Lastly, we obtain mesh representations which we employ for manipulation planning. Our framework is visually summarized in Figure 2.

**Instance Segmentation** ARM takes as input a RGB image,  $I \in \mathbb{R}^{h \times w \times 3}$ , and an organized point cloud,  $P \in \mathbb{R}^{h \times w \times 3}$  computed by backprojecting a depth image with camera intrinsics. This is passed to an instance segmentation network  $\mathcal{S}$  which outputs instance masks  $L = \mathcal{S}(I, D) \in \mathcal{L}^{h \times w}$ , where  $\mathcal{L} = \{0, \dots, K\}$  and  $K$  is the number of detected object instances. We use UOIS-Net [33] as  $\mathcal{S}$  which produces high quality segmentations for unseen objects.

**Four Channel Voxel Input Computation** We introduce a four-channel voxel representation to enable ARM to reason about the spatial extent of other objects during reconstruction. For each object  $o \in \mathcal{L}$ , we compute a voxel occupancy grid  $F_o \in \{0, 1\}^{d^3 \times 4}$  augmented with the surrounding objects' occupancies, as well as with voxel visibilities with respect to the camera. Let  $F_o^i$  denote the  $i^{\text{th}}$  channel of  $F_o$ .  $F_o^1$  is the voxel grid of object  $o$  alone, which is computed by voxelizing  $P_o$ , the point cloud segmented with the instance mask for  $o$ .  $F_o^2$  contains all other objects in  $L$  except for  $o$ .  $F_o^3$  consists of a mask of empty voxels, and  $F_o^4$  contains unobserved voxels. Note that  $F_o^3, F_o^4$  are computed using the camera extrinsics and intrinsics.  $F_o$  is centered at the center of mass of object  $o$  and has side length  $k\delta_o$ , where  $\delta_o$  is the maximum distance between points in  $P_o$ . In our implementation,  $k = 4$  to allow filling of occluded regions. Finally, we translate  $F_o$  so the table occupies the  $z = 0$  plane in our voxel grid.

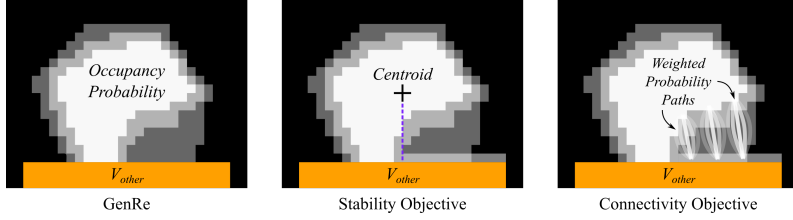


Figure 3: Impact of stability and connectivity objectives. Left: occupancy probabilities of an estimated shape, in greyscale. Adding the stability objective makes the object stable, and adding the connectivity objective fills in the gap between the shape and inferred base.

**3D Shape Estimation and Scene Reconstruction.** For each object  $o \in \mathcal{L}$ , we use  $F_o$  as input to a 3D reconstruction network  $\mathcal{C}$  which outputs the probability of  $o$ 's presence at each voxel as  $\mathcal{C}(F_o) = V_o \in [0, 1]^{d^3}$ . We use GenRe-Oracle [14] as our 3D reconstruction network. GenRe-Oracle is a modification of GenRe that uses depth, rather than RGB data. GenRe-Oracle projects observed pointclouds onto a sphere, inpaints them using a CNN, backprojects to a 3D voxel grid, and then refines the voxel grid using a 3D encoder-decoder architecture. Finally, we use marching cubes [34] after thresholding  $V_o$  to transform the output voxel probabilities into a mesh.

We post-process the meshes in order to make them suitable for physics simulation. First, we remove intersections between meshes to prevent inconsistent behavior in simulation, removing from the larger of the intersecting meshes. We then compute an approximate convex decomposition of each mesh using V-HACD [35].

**Manipulation Planning** We pose the task of manipulation planning in the form of an MDP consisting of an action space  $a \in A$ , a state space  $s \in S$ , a stochastic transition function  $G(s', s, a) = P(s'|s, a)$ , and a single RGB-D view of the corresponding environment. We solve this MDP by reconstructing every object in the view, instantiating a simulation of the environment from the robot viewpoint with the reconstructed objects, and then using a physics simulator [8] to approximate the transition function  $G$ . We use MPPI [11] to plan a sequence of actions in the simulator, and execute this plan in the real environment.

## 3.2 Loss Functions

GenRE [14] uses a weighted combination of cross entropy and a surface loss between reconstructed and ground truth voxels during training. However, in robotic settings, optimizing these losses alone are not sufficient to solve the downstream task of robotic manipulation, as we show in Section 4.4. This results in reconstructions with poor physical fidelity during the planning phase, often due to instability of poor reconstruction of occluded regions. We tackle this issue by designing auxiliary differentiable loss functions based on two physical priors: 1) objects are stable prior to manipulation, and 2) objects are a single connected component. Figure 3 gives an overview of these loss functions.

### 3.2.1 Stability Loss

Our stability loss provides a prior over object shape, even in occluded regions, by reasoning about hidden supports objects may have. An object is in static equilibrium if the net forces acting upon it are equal to zero [36]. This means that the center of mass is within the base of support of an object. Technically, the center of mass must be behind a pivot point (where the object rests on another object) along every direction  $s$  perpendicular to the force of gravity  $\vec{g}$ .

We first define some notation here. Recall that  $V_o$  parameterizes a multivariate Bernoulli distribution over binarized voxel grids. For sample  $v \sim V_o$ , we define  $M(v)$  to be the center of mass of  $v$ . Furthermore, let  $i \in d^3$  index the voxel grid, and  $S = \{s : s \perp \vec{g}\}$  be the set of directions perpendicular to  $\vec{g}$ . Then, for each  $s \in S$ , let  $i^s$  and  $M^s(v)$  be the projections of  $i$  and  $M(v)$  onto the plane defined by  $s$  and  $\vec{g}$  that passes through the origin. We denote  $H_s(i)$  as the set of voxels belonging to other objects that support  $i$  in direction  $s$ , which can happen when  $i$  is directly above or leaning against such voxels. Finally,  $V_{\bar{o}}$  is the probabilities of other objects output by the 3D reconstruction network.

Given this notation, we can define our stability loss to be the probability that  $v$  is stable. Let  $E(v)$  be the event that  $v$  is stable. Then our stability objective is defined as

$$P(E(v)) = \prod_{s \in S} (1 - u_s) \quad (1)$$

$$u_s = \prod_{i \in d^3} \left[ 1 - V_o(i) P(i^s > M^s(v)) h_s(i) \right], \quad h_s(i) = 1 - \prod_{i' \in H_s(i)} (1 - V_o(i')) \quad (2)$$

$u_s$  is the probability that  $v$  is unstable in direction  $s$ . It is the probability that every voxel  $i$  is unstable; that is  $i$  either doesn't exist, doesn't support  $v$  along direction  $s$ , or isn't supported ( $h_s(i)$ ). Eq. (1) is intractable, so in order to take the gradient we introduce independence assumptions and the approximation that a voxel  $i$  exists only if  $V(i) \geq 0.5$  to derive an efficiently computable derivative of object stability with respect to each object voxel:

$$\frac{d \log P(E(v))}{dV_o(i)} = \sum_{s \in S} \frac{-u'_s}{1 - u_s} \quad (3)$$

$$u'_s = -P(i^s > M^s(v)) \hat{h}_s(i) \prod_{i_o \in d^3, i_o \neq i} \left[ 1 - P(i_o^s > M^s(v)) \mathbb{1}\{V(i_o) \geq 0.5\} \hat{h}_s(i_o) \right] \quad (4)$$

$$\hat{h}_s(i_o) = 1 - \prod_{i_b \in H_s(i_o)} \left[ 1 - \mathbb{1}\{V_o(i_b) \geq 0.5\} \right] \quad (5)$$

This gradient captures several intuitive properties of stability. If an object has even a single voxel supporting it in a particular direction then it is stable. For a direction  $s$ , if a single supported voxel  $i_o$  exists, then  $u'_s$  is close to zero, and the magnitude of the derivative in that direction will be small. Vice versa, when no supporting voxel is present,  $u'_s$  is close to 1 and the magnitude is nontrivial. This captures the idea that when supporting voxels are present, the effect on stability is small, but when no supporting voxels are present, the effect is large. Importantly, this prior is shape agnostic: it is not biased towards making an object stable by adding a base under existing voxels, for example, but rather only increases the probability of any voxel that would make the object stable, minimizing reconstruction deviation from the learned shape prior. A full derivation can be found in appendix A.

### 3.2.2 Connectivity Loss

Our connectivity loss imposes a prior on object shape even in occluded regions by allowing the network to infer connections between disjoint parts of observed objects. This complements the stability objective which frequently infers occluded bases of objects. We define  $v$  to be connected if for every pair of existent voxels  $a, b$ , there exists a path  $t = \{i_0, i_1, \dots\}$  between  $a$  and  $b$ . The probability that a path  $t$  exists in  $v$  is  $P(t) = \prod_{i \in t} V_o(i)$ . Let  $T(a, b)$  be the set of all possible paths between  $a$  and  $b$ ,  $C(v)$  be the event that  $v$  is connected, and  $C(a, b)$  be the event that there is a path between  $a$  and  $b$ . Then we define our connectivity objective as

$$P(C(v)) = \prod_{a, b \in d^3, a \neq b} \left[ V_o(a)V_o(b)P(C(a, b)) + 1 - V_o(a)V_o(b) \right] \quad (6)$$

The derivative of this equation is intractable because it requires considering every path  $t$  between every vertex pair  $(a, b)$ . To resolve this, we note that relative to the most likely path  $t^*$  between  $a$  and  $b$ , most paths have small probability. Thus, for any other voxel  $c$ , we may ignore low probability paths passing through  $c$  when calculating their contribution to the connectivity of  $a$  and  $b$  and only consider the most likely path from  $a$  to  $b$  passing through  $c$ . With this approximation, our per-voxel derivative of Eq. (6) is

$$\frac{d \log P(C(v))}{dV_o(c)} = \sum_{a, b \in d^3, a \neq b \neq c} \frac{V_o(a)V_o(b) \frac{d}{dV_o(c)} P(C(a, b))}{V_o(a)V_o(b)P(C(a, b)) + 1 - V_o(a)V_o(b)} \quad (7)$$

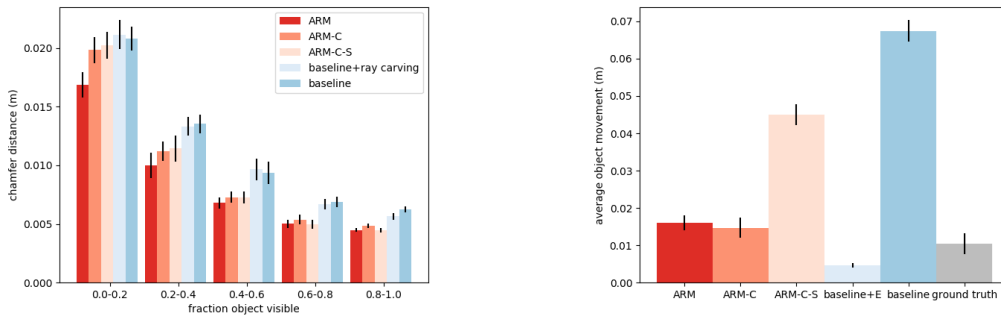


Figure 4: (left) Chamfer distances on held-out objects, broken down by observation occlusion. (right) Average stability of reconstructed objects. Error bars are a 90% confidence interval.

where  $P(C(a,b)) = P\left(\bigcup_{t \in T(a,b)} t\right) \approx P(t^* \cup t^c)$ ,  $t^*$  is the path from  $a$  to  $b$  with the highest probability of existing, and  $t^c$  is the path from  $a$  to  $b$  that includes  $c$  with the highest probability of existing. This approximation preserves several desirable properties of the exact gradient. First, it only encourages connecting the object by reinforcing the most likely paths, rather than the physically shortest paths. By considering each most likely path from  $a$  to  $b$  that passes through  $c$ , it also produces dense connections, rather than only amplifying the shortest path between  $a$  to  $b$ , which would often result in shapes connected by single voxel width paths.

## 4 Experiments

### 4.1 Implementation Details

We implement ARM using UOIS-Net [33] for instance segmentation, and the GenRE depth backbone [14] for 3D reconstruction. We use MuJoCo [8] as a physics simulator for our reconstructed environment. To train ARM, we create a large dataset of cluttered tabletop scenes in MuJoCo using ShapeNet [37] tables and objects. We divide the ShapeNet objects into training and test sets, containing 4803 and 3368 unique objects respectively. For each scene, we drop between 5 and 20 randomly selected objects onto a table to ensure cluttered scenes and stacked objects with complex stability relationships. We render several views with randomized camera positions, using a custom OpenGL renderer to produce realistic images. Each network is trained with ADAM for approximately 100,000 iterations with a batch size of 16. Stability and connectivity loss gradients are only applied on occluded voxels, as all other voxels are observed to be either empty or occupied. Additional implementation and training details are in Appendix C.

### 4.2 Baselines

Our main baseline that we compare against is GenRE-Oracle [14], which we denote as baseline. In order to test the most direct way of using the information about observed occupied and unoccupied voxels encoded in the four channel representation, we introduce a simple modification to GenRE to give baseline+ray carving where we remove all observed empty voxels after reconstruction. Additionally, in order to test their significance, we train two ablations of our method, ARM-C, ARM without the connectivity prior and ARM-C-S, ARM without the connectivity or stability prior.

### 4.3 Reconstruction Quality

We quantitatively compare the visual reconstruction quality of ARM to our baselines on reconstruction of cluttered scenes generated with held-out test objects and ground truth segmentations in Figure 4 (left). ARM outperforms the baseline at all occlusion levels, improving Chamfer loss by 28% overall. This improvement is especially pronounced on highly occluded objects, where the stability and connectivity objectives combined allow ARM to estimate occluded bases and fill gaps between those bases and the observed parts of objects. However, low Chamfer distance alone is not sufficient

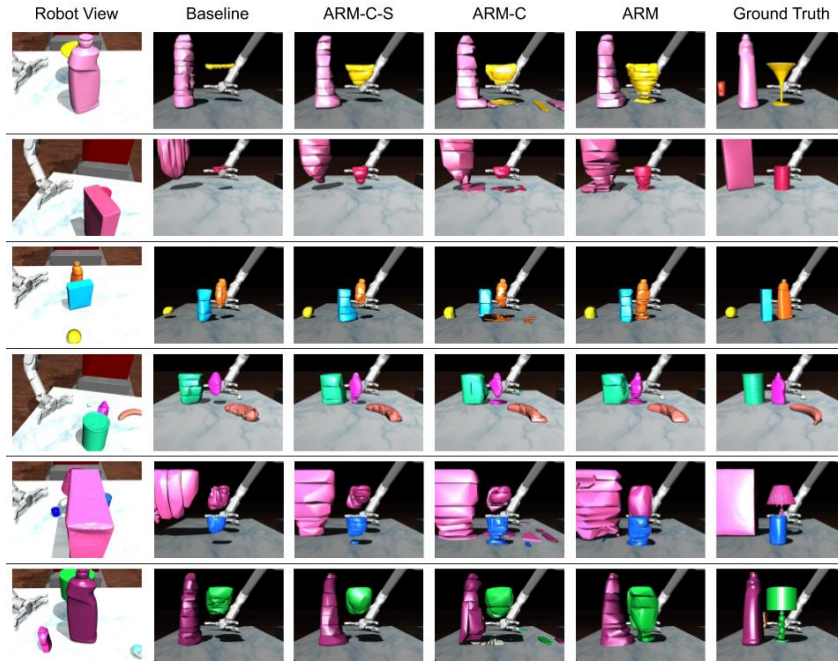


Figure 5: Qualitative reconstruction results. ARM is able to infer both bases and occluded object regions.

for accurate simulation; reconstructed objects must also exhibit similar physics to the ground truth. In Figure 4 (right), we measure scene stability by placing each reconstructed scene into a physics engine, simulating forward for five seconds with gravity as the only force, and measuring the  $L_2$  displacement of the reconstructed object centers. As a simple yet effective baseline for stability, we consider baseline+E, where we extrude the reconstructed mesh (from baseline) down to the table to ensure mesh stability. Both the baseline and ARM-C-S frequently reconstruct unstable objects that fall or tumble. Adding the stability prior improves object stability to near that of the ground truth meshes. Note that ground truth meshes move a small amount because Mujoco considers only one point of contact between each pair of mesh geometries, which can cause meshes to slowly move.

In Figure 5, we provide qualitative reconstruction results. The baseline and ARM-C-S only reconstruct the visible portions of occluded meshes at the top, producing reconstructions that will tumble over as soon as simulation begins. ARM-C reconstructs the bases to produce stable meshes, but still leaves large voids in the middle of reconstructions which frequently cause manipulation to fail. On the other hand, ARM is able to both reconstruct bases and effectively reason about occluded regions, producing a tapered reconstruction for the yellow wine glass in the top row, but filling in a cylinder for the blue tin can in the fifth row.

#### 4.4 Robot Manipulation

**Robot Manipulation Tasks** To evaluate the efficacy of our method on robot manipulation tasks, we create a suite of robotics manipulation tasks across a range of challenging objects in cluttered scenes. We consider three important robot tasks: grasping, pushing, and rearrangement, which entails grasping and pushing/pulling. In each task, the robot first creates a 3D reconstruction of the environment from an RGBD observation. It then plans a trajectory with the 3D reconstruction, and finally executes the trajectory in the ground truth environment. We execute each task on 14 different target objects from the YCB dataset [3] and 12 from a set of challenging, highly non-convex objects downloaded from online 3D repositories, all previously unseen during training (see Figure 5 for examples). For each task and target object we consider 10 occlusion intervals, or fraction of the target manipulation object visible to the robot, from  $[0, 0.1)$  to  $[0.9, 1)$ , and no occlusion. For each task, target object, and occlusion interval, we generate three scenes by randomly placing unseen YCB objects until the occlusion of target manipulation object is within the desired range, for total

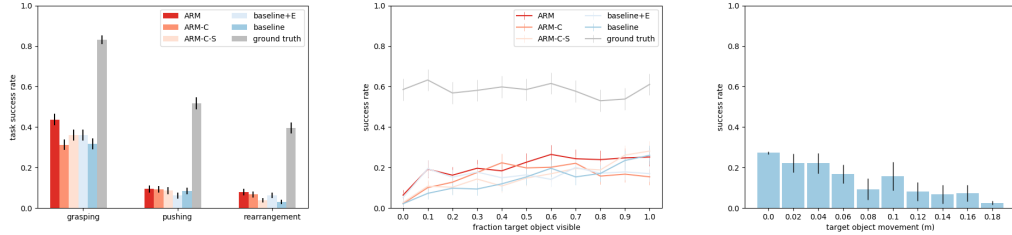


Figure 6: (left) Success rates on manipulation tasks using models generated by different reconstruction algorithms. (middle) Manipulation success rate vs. target object occlusion. Visibilities are binned in increments of 0.1, so 0.0 includes all visibilities in  $[0,0.1)$ . (right) Task success vs. target object stability. Error bars are a 90% confidence interval.

of 2574 tasks. To isolate the effects of different 3D reconstruction algorithms, we use ground truth instance segmentations. More details on our cluttered robot manipulation benchmark are available in Appendix D.

**Robot Task Performance** Figure 6 (left) shows average task success rates on the manipulation tasks for each of the methods. ARM achieves the best performance across all tasks, improving the baseline and the extrusion baseline by 42% and 25% respectively. In Figure 6 (middle), we break down performance by target object occlusion, and find that while ARM and the baseline perform similarly at very low levels of occlusion, at almost all levels of visibility below 80% ARM performs the best. Notably, while the baseline success rate on 10% visible objects is only about 25% the success rate on completely unoccluded objects, ARM’s success rate on 10% visible objects is 75% its success rate on unoccluded objects, showing that our stability and connectivity priors are less prone to performance degradation in the face of occlusion, which is due to our stability and connectivity priors. While ARM gives significant improvements over the baseline, this task suite is still quite challenging: most tasks involve high levels of target object occlusion, and even poor reconstructions of self-occluded object regions can cause manipulation failure. Lastly, we analyze the impact of object stability on performance in Figure 6 (right), where we plot task success vs. target object stability, showing a correlation between task success and object stability.

## 5 Conclusion

We have shown that directly applying 3D object reconstruction methods in cluttered robotic environments can produce reconstructions with low physical fidelity, which often leads to unsuccessful task execution. We proposed a modular framework, ARM, that includes a stability prior, a connectivity prior, and a multi-channel input representation to deliver more physically faithful reconstructions. ARM generates 3D reconstructions that are not only better by standard visual loss metrics, but more importantly they allow for significantly better robot task performance in challenging cluttered scenes. We hope our reconstruction system will enable further model-based learning and control applications. While ARM enables significant improvements over the baseline on reconstruction and manipulation of occluded objects, performance on highly occluded objects is still far from that of the ground truth. To enable further research on this challenging task, we will publicly release our large, high-quality cluttered dataset and robot evaluation benchmarks.

## 6 Acknowledgements

We would like to thank the authors of UOIS [33], GenRE [14], Dex-Net 2.0 [38] and POLO [39] for releasing code used in this paper. We would like to thank Jonathan Tremblay for help on early versions of this project. We would like to thank Aravind Rajeswaran, Kendall Lowrey, and Colin Summers for advice and help with MuJoCo. This work was supported by an NDSEG Fellowship and ONR grant N00014-18-1-2826. The GPU machine used for this research was donated by Nvidia. This work was partially funded by the National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), National Science Foundation NRI (#1637748), the Office of Naval Research, the RCTA, Amazon, and Honda Research Institute USA.



## References

- [1] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [2] 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*, 37(4-5):421–436, 2018.
- [3] 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 *International Conference on Advanced Robotics (ICAR)*, 2015.
- [4] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *Robotics: Science and Systems (RSS)*, 2018.
- [5] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox. Deepim: Deep iterative matching for 6d pose estimation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [6] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning (CoRL)*, 2018.
- [7] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. *GitHub repository*, 2016.
- [8] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [9] M. Dogar and S. Srinivasa. A framework for push-grasping in clutter. *Robotics: Science and Systems (RSS)*, 2011.
- [10] M. Stilman, J.-U. Schamburek, J. Kuffner, and T. Asfour. Manipulation planning among movable obstacles. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2007.
- [11] G. Williams, A. Aldrich, and E. Theodorou. Model predictive path integral control using covariance variable importance sampling. *arXiv preprint arXiv:1509.01149*, 2015.
- [12] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.
- [13] G. Gkioxari, J. Malik, and J. Johnson. Mesh r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [14] X. Zhang, Z. Zhang, C. Zhang, J. B. Tenenbaum, W. T. Freeman, and J. Wu. Learning to Reconstruct Shapes from Unseen Classes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [15] E. Smith, S. Fujimoto, and D. Meger. Multi-view silhouette and depth decomposition for high resolution 3d object representation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [16] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [17] S. Tulsiani, S. Gupta, D. F. Fouhey, A. A. Efros, and J. Malik. Factoring shape, pose, and layout from the 2d image of a 3d scene. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] Y. Nie, X. Han, S. Guo, Y. Zheng, J. Chang, and J. J. Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

- [19] N. Kulkarni, I. Misra, S. Tulsiani, and A. Gupta. 3d-relnet: Joint object and relational network for 3d prediction. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [20] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [21] D. Shin, C. C. Fowlkes, and D. Hoiem. Pixels, voxels, and views: A study of shape representations for single view 3d object shape prediction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [22] X. Han, H. Laga, and M. Bennamoun. Image-based 3d object reconstruction: State-of-the-art and trends in the deep learning era. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [23] A. Kundu, Y. Li, and J. M. Rehg. 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [24] E. Sucar, K. Wada, and A. Davison. Neural object descriptors for multi-view shape reconstruction. *arXiv preprint arXiv:2004.04485*, 2020.
- [25] Z. Jia, A. Gallagher, A. Saxena, and T. Chen. 3d-based reasoning with blocks, support, and stability. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [26] Y. Du, Z. Liu, H. Basevi, A. Leonardis, B. Freeman, J. Tenenbaum, and J. Wu. Learning to exploit stability for 3d scene parsing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [27] Y. Chen, S. Huang, T. Yuan, S. Qi, Y. Zhu, and S.-C. Zhu. Holistic++ scene understanding: Single-view 3d holistic scene parsing and human pose estimation with human-object interaction and physical commonsense. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [28] J. Varley, C. DeChant, A. Richardson, J. Ruales, and P. Allen. Shape completion enabled robotic grasping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [29] M. Van der Merwe, Q. Lu, B. Sundaralingam, M. Matak, and T. Hermans. Learning continuous 3d reconstructions for geometrically aware grasping. In *IEEE Conference on Robotics and Automation (ICRA)*, 2020.
- [30] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee. Learning 6-dof grasping interaction via deep geometry-aware 3d representations. In *IEEE Conference on Robotics and Automation (ICRA)*, 2018.
- [31] J. Lundell, F. Verdoja, and V. Kyrki. Robust grasp planning over uncertain shape completions. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [32] J. Lundell, F. Verdoja, and V. Kyrki. Beyond top-grasps through scene completion. In *IEEE Conference on Robotics and Automation (ICRA)*, 2020.
- [33] 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 (CoRL)*, 2019.
- [34] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.
- [35] K. Mamou, E. Lengyel, and A. Peters. Volumetric hierarchical approximate convex decomposition. *Game Engine Gems 3*, pages 141–158, 2016.
- [36] P. P. Urone, K. Dirks, and M. Sharma. *Statics and Torque*, page 289316. OpenStax.

- [37] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [38] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg. Dexnet 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. *arXiv preprint arXiv:1703.09312*, 2017.
- [39] K. Lowrey, A. Rajeswaran, S. Kakade, E. Todorov, and I. Mordatch. Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control. In *ICLR*, 2019.