

The Best of Both Modes: Separately Leveraging RGB and Depth for Unseen Object Instance Segmentation

Christopher Xie¹ Yu Xiang² Arsalan Mousavian² Dieter Fox^{2,1}

¹University of Washington ²NVIDIA

chrisxie@cs.washington.edu {yux, amousavian, dieterf}@nvidia.com

Abstract: In order to function in unstructured environments, robots need the ability to recognize unseen novel objects. We take a step in this direction by tackling the problem of segmenting unseen object instances in tabletop environments. However, the type of large-scale real-world dataset required for this task typically does not exist for most robotic settings, which motivates the use of synthetic data. We propose a novel method that separately leverages synthetic RGB and synthetic depth for unseen object instance segmentation. Our method is comprised of two stages where the first stage operates only on depth to produce rough initial masks, and the second stage refines these masks with RGB. Surprisingly, our framework is able to learn from synthetic RGB-D data where the RGB is non-photorealistic. To train our method, we introduce a large-scale synthetic dataset of random objects on tabletops. We show that our method, trained on this dataset, can produce sharp and accurate masks, outperforming state-of-the-art methods on unseen object instance segmentation. We also show that our method can segment unseen objects for robot grasping. Code, models and video can be found at the project website¹.

1 Introduction

For a robot to work in an unstructured environment, it must have the ability to recognize new objects that have not been seen before by the robot. Assuming every object in the environment has been modeled is infeasible and impractical. Recognizing unseen objects is a challenging perception task since the robot needs to learn the concept of “objects” and generalize it to unseen objects. Building such a robust object recognition module is valuable for robots interacting with objects, such as performing different manipulation tasks. A common environment in which manipulation tasks take place is on tabletops. Thus, in this paper, we approach this by focusing on the problem of unseen object instance segmentation (UOIS), where the goal is to separately segment every arbitrary (and potentially unseen) object instance, in tabletop environments.

In order to ensure the generalization capability of the module to recognize unseen objects, we need to learn from data that contains large amounts of various objects. However, in many robot environments, large-scale datasets with this property do not exist. Since collecting a large dataset with ground truth annotations is expensive and time-consuming, it is appealing to utilize synthetic data for training, such as using the ShapeNet repository which contains thousands of 3D shapes of different objects [1]. However, there exists a domain gap between synthetic data and real world data. Training directly on synthetic data only usually does not work well in the real world [2].

Consequently, recent efforts in robot perception have been devoted to the problem of Sim2Real, where the goal is to transfer capabilities learned in simulation to real world settings. For instance, some works have used domain adaptation techniques to bridge the domain gap when unlabeled real data is available [3, 4]. Domain randomization [5] was proposed to diversify the rendering of synthetic data for training. These methods mainly use RGB as input, and while that is desirable since there is evidence that models trained on (real world) RGB have been shown to produce sharp and accurate masks [6], this complicates the Sim2Real problem because state-of-the-art simulators typically cannot produce photo-realistic renderings. On the other hand, models trained with synthetic depth have been shown to generalize reasonably well (without fine-tuning) for simple settings such as bin-picking

¹<https://rse-lab.cs.washington.edu/projects/unseen-object-instance-segmentation/>

[7, 8]. However, in more complex settings, noisy depth sensors can limit the application of such methods. An ideal method would combine the generalization capability of training on synthetic depth and the ability to produce sharp masks by training on RGB.

In this work, we investigate how to utilize synthetic RGB-D images for UOIS in tabletop environments. We show that simply combining synthetic RGB images and synthetic depth images as inputs does not generalize well to the real world. To tackle this problem, we propose a simple two-stage framework that separately leverages the strengths of RGB and depth for UOIS. Our first stage is a Depth Seeding Network (DSN) that operates only on depth to produce rough initial segmentation masks. Training DSN with depth images allows for better generalization to the real world data. However, these initial masks from DSN may contain false alarms or inaccurate object boundaries due to depth sensor noise. In these cases, utilizing the textures in RGB images can significantly help.

Thus, our second stage is a Region Refinement Network (RRN) that takes an initial mask of an object from DSN and an RGB image as input and outputs a refined mask. Our surprising result is that, conditioned on initial masks, our RRN can be trained on non-photorealistic synthetic RGB images without any of the domain randomization or domain adaptation approaches of Sim2Real. We posit that mask refinement is an easier problem than directly using RGB as input to produce instance masks. We empirically show robust generalization across many different objects in cluttered real world data. In fact, as we show in our experiments, our RRN works almost as well as if it were trained on real data. Our framework, including the refinement stage, can produce sharp and accurate masks even when the depth sensor reading is noisy. We show that it outperforms state-of-the-art methods trained using any combination of RGB and depth as input, including Mask-RCNN [6].

To train our method, we introduce a synthetic dataset of tabletop objects in house environments. The dataset consists of indoor scenes of random ShapeNet [1] objects on random tabletops. We use a physics simulator [9] to generate the scenes and render depth and non-photorealistic RGB. Despite this, training our proposed method on this dataset results in state-of-the-art results on the OCID dataset [10] and the OSD dataset [11] introduced for UOIS.

This paper is organized as follows. After reviewing related work, we discuss our proposed method. We then describe our generated synthetic dataset, followed by experimental results and a conclusion.

2 Related Works

Object Instance Segmentation. Object instance segmentation is the problem of segmenting every object instance in an image. Many approaches for this problem involve top-down solutions that combine segmentation with object proposals in the form of bounding boxes [6, 12]. However, when the bounding boxes contain multiple objects (e.g. heavy clutter in robot manipulation setups), the true segmentation mask is ambiguous and these methods struggle. More recently, a few methods have investigated bottom-up methods which assign pixels to object instances [13, 14, 15]. Most of these algorithms provide instance masks with category-level semantic labels, which do not generalize to unseen objects in novel categories.

One approach to adapting object instance segmentation techniques to unseen objects is to employ “class-agnostic” training, which treats all object classes as one foreground category. One family of methods exploits motion cues with class-agnostic training in order to segment arbitrary moving objects [16, 17]. Another family of methods are class-agnostic object proposal algorithms [18, 19, 20]. However, these methods will segment everything and require some post-processing method to select the masks of interest. We also train our proposed method in a class-agnostic fashion, but instead focus our notion of unseen objects in particular environments such as tabletop settings.

Sim2Real Perception. Training a model on synthetic RGB and directly applying it to real data typically fails [2]. Many methods employ some level of rendering randomization [21, 22, 23, 5, 24, 25], including lighting conditions and textures. However, they typically assume specific object instances and/or known object models. Another family of methods employ domain adaptation to bridge the gap between simulated and real images [3, 4]. Algorithms trained on depth have been shown to generalize reasonably well for simple settings [7, 8]. However, noisy depth sensors can limit the application of such methods. Our proposed method is trained purely on (non-photorealistic) synthetic RGB-D data and is accurate even when depth sensors are inaccurate, and can be trained without adapting or randomizing the synthetic RGB.

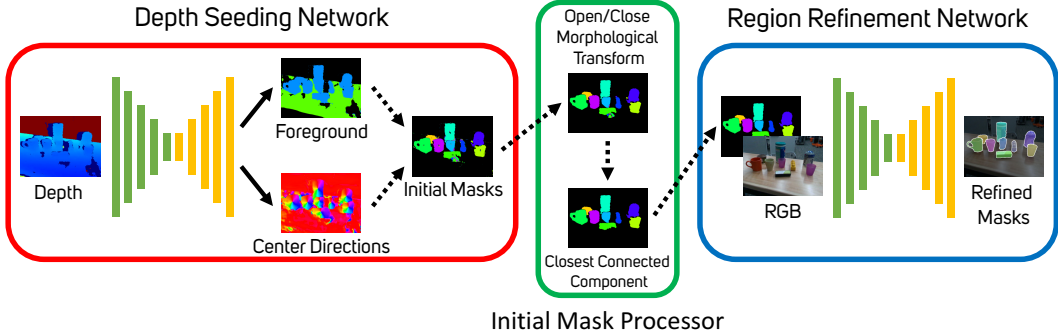


Figure 1: Overall architecture. The Depth Seeding Network (DSN) is shown in the red box, the Initial Mask Processor (IMP) in the green box, and the Region Refinement Network (RRN) in the blue box. The images come from a real example taken by an RGB-D camera in our lab. Despite the level of noise in the depth image (due to reflective table surface), our method is able to produce sharp and accurate instance masks.

3 Method

Our framework consists of two separate networks that process Depth and RGB separately to produce instance segmentation masks. First, the Depth Seeding Network (DSN) takes a depth image and outputs a semantic segmentation and 2D directions to object centers. From these, we calculate initial instance segmentation masks with a Hough voting layer. These initial masks are expected to be quite noisy, so we use an Initial Mask Processor (IMP) to robustify the masks with standard image processing techniques. Lastly, we refine the processed initial masks using our Region Refinement Network (RRN). Because the Hough voting layer and IMP are non-differentiable, DSN and RRN are trained separately as opposed to end-to-end. The full architecture is shown in Figure 1.

3.1 Depth Seeding Network

3.1.1 Network Architecture

The DSN takes as input a 3-channel organized point cloud, $D \in \mathbb{R}^{H \times W \times 3}$, of XYZ coordinates. D is passed through an encoder-decoder architecture to produce two outputs: a semantic segmentation mask $F \in \mathbb{R}^{H \times W \times C}$, where C is the number of semantic classes, and 2D directions to object centers $V \in \mathbb{R}^{H \times W \times 2}$. We use $C = 3$ for our semantic classes: background, tabletop (table plane), and tabletop objects. Each pixel of V encodes a 2-dimensional unit vector pointing to the 2D center of the object. We define the center of the object to be the mean pixel location of the observable mask. Although we do not explicitly make use of the tabletop label in Section 5, it can be used in conjunction with RANSAC [26] in order to better estimate the table and get rid of false positive masks. For the encoder-decoder architecture, we use a U-Net [27] architecture where each 3×3 convolutional layer is followed by a GroupNorm layer [28] and ReLU. Sitting on top of this is two parallel branches of convolutional layers that produce the foreground mask and center directions.

In order to compute the initial segmentation masks from F and V , we design a Hough voting layer similar to [21]. First, we discretize the space of all 2D directions into M equally spaced bins. For every pixel in the image, we compute the percentage of discretized directions from all other pixels that point to it and use this as a score for how likely the pixel is an object center. We then threshold the scores to select object centers and apply non-maximum suppression. Given these object centers, each pixel in the image is assigned to the closest center it points to, which gives the initial masks as shown in the red box of Figure 1.

3.1.2 Loss Function

To train the DSN, we apply two different loss functions on the semantic segmentation F and the direction prediction V . For the semantic segmentation loss, we use a weighted cross entropy as this has been shown to work well in detecting object boundaries in imbalanced images [29]. The loss is $\ell_F = \sum_i w_i \ell_{CE}(\hat{F}_i, F_i)$ where i ranges over pixels, \hat{F}_i, F_i are the predicted and ground truth probabilities of pixel i , respectively, and ℓ_{CE} is the cross-entropy loss. The weight w_i is inversely proportional to the number of pixels with labels equal to F_i , normalized to sum to 1.

We apply a weighted cosine similarity loss to the direction prediction V . The cosine similarity is focused on the pixels belonging to the tabletop objects, but we also apply it to the background/tabletop pixels to have them point in a fixed direction to avoid potential false positives. The loss is given by

$$\ell_V = \frac{1}{2} \sum_{i \in O} \alpha_i \left(1 - \hat{V}_i^\top V_i\right) + \frac{\lambda_{bt}}{2|B \cup T|} \sum_{i \in B \cup T} \left(1 - \hat{V}_i^\top \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right)$$

where \hat{V}_i, V_i are the predicted and ground truth unit directions of pixel i , respectively. B, T, O are the sets of pixels belonging to background, table, and tabletop object classes, respectively. α_i is inversely proportional to the number of pixels with the same *instance* label as pixel i , giving equal weight to each instance regardless of size. We set $\lambda_{bt} = 0.1$. The total loss is given by $\ell_F + \ell_V$.

3.2 Initial Mask Processing Module

Computing the initial masks from F and V with the Hough voting layer often results in noisy masks (see an example in Figure 1). For example, these instance masks often exhibit salt/pepper noise and erroneous holes near the object center. As shown in Section 5, the RRN has trouble refining the masks when they are scattered as such. To robustify the algorithm, we propose to use two simple image processing techniques to clean the masks before refinement.

For a single instance mask, we first apply an opening operation, which consists of mask erosion followed by mask dilation [30], removing the salt/pepper noise issues. Next we apply a closing operation, which is dilation followed by erosion, which closes up small holes in the mask. Finally, we select the closest connected component to the object center and discard all other components. Note that these operations are applied to each instance mask separately.

3.3 Region Refinement Network

3.3.1 Network Architecture

This network takes as input a 4-channel image, which is a cropped RGB image concatenated with an initial instance mask. The full RGB image is cropped around the instance mask with some padding for context, concatenated with the (cropped) mask, then resized to 224×224 . This gives an input image $I \in \mathbb{R}^{224 \times 224 \times 4}$. The output of the RRN is the refined mask probabilities $R \in \mathbb{R}^{224 \times 224}$, which we threshold to get the final output. We use the same U-Net architecture as in the DSN. To train the RRN, we apply the loss ℓ_F with two classes (foreground vs. background) instead of three.

3.3.2 Mask Augmentation

In order to train the RRN, we need examples of perturbed instance masks. This problem can be seen as a data augmentation task where we augment the ground truth mask into something that resembles an initial mask (after the IMP). We detail the different augmentation techniques used below:

- Translation/rotation: We translate the mask by sampling a displacement vector proportionally to the mask size. Rotation angles are sampled uniformly in $[-10^\circ, 10^\circ]$.
- Adding/cutting: For this augmentation, we choose a random part of the mask near the edge, and either remove it (cut) or copy it outside of the mask (add). This reflects the setting when the initial mask egregiously overflows from the object, or is only covering part of it.
- Morphological operations: We randomly choose multiple iterations of either erosion or dilation of the mask. The erosion/dilation kernel size is set to be a percentage of the mask size, where the percentage is sampled from a beta distribution. This reflects inaccurate boundaries in the initial mask, e.g. due to noisy depth sensors.
- Random ellipses: We sample the number of ellipses to add or remove in the mask from a Poisson distribution. For each ellipse, we sample both radii from a gamma distribution, and sample a rotation angle as well. This augmentation requires the RRN to learn to remove irrelevant blots outside of the object and close up small holes within it.

4 Tabletop Object Dataset

Many desired robot environment settings (e.g. kitchen setups, cabinets) lack large scale training sets to train deep networks. To our knowledge, there is also no large scale dataset for tabletop objects.




Figure 2: Examples from our Tabletop Object Dataset. RGB, depth, and instance masks are shown.

To remedy this, we generate our own synthetic dataset which we name the Tabletop Object Dataset (TOD). This dataset is comprised of 40k synthetic scenes of cluttered objects on a tabletop in home environments. We use the SUNCG house dataset [31] for home environments and ShapeNet [1] for tables and arbitrary objects. We only use ShapeNet tables that have convex tabletops, and filter the ShapeNet object classes to roughly 25 classes of objects that could potentially be on a table.

Each scene in the dataset is of a random table from ShapeNet inside a random room from a SUNCG house. We randomly sample anywhere between 5 and 25 objects to put on the table. The objects are either randomly placed on the table, on top of another object (stacked), or generated at a random height and orientation above the table. We use PyBullet [9] to simulate physics until the objects come to rest and remove any objects that fell off the table. Next, we generate seven images (RGB, depth, and ground truth instance masks) using PyBullet’s rendering capabilities. One view is of only background, another is of just the table in the room, and the rest are taken with random camera viewpoints with the tabletop objects in view. The viewpoints are sampled at a height of between .5m and 1.2m above the table and rotated randomly with an angle in $[-12^\circ, 12^\circ]$. The images are generated at a resolution of 640×480 with vertical field-of-view of 45 degrees. The segmentation has a tabletop (table plane, not including table legs) label and instance labels for each object.

We show some example images of our dataset in Figure 2. The rightmost two examples show that some of our scenes are heavily cluttered. Note that the RGB looks non-photorealistic. In particular, PyBullet is unable to load textures of some ShapeNet objects (see gray objects in leftmost two images). PyBullet was built for reinforcement learning, not computer vision, thus its rendering capabilities are insufficient for photorealistic tasks [9]. Despite this, our RRN is able to learn to snap masks to object boundaries from this synthetic dataset.

5 Experiments

5.1 Implementation Details

In this section, all models are trained from scratch for 100k iterations of SGD with a batch size of 8. We use a fixed learning rate of $1e-2$ for our models and follow the Detectron schedule for Mask RCNN. All images have a resolution $H = 480, W = 640$. During DSN training, we augment depth with multiplicative gamma noise and additive Gaussian Process noise, similar to [7]. Inputs to the RRN are augmented as described in Section 3.3.2. All experiments run on a NVIDIA RTX 2080ti.

Method	Overlap			Boundary		
	P	R	F	P	R	F
GCUT [32]	21.5	51.5	25.7	10.2	46.8	15.7
SCUT [33]	45.7	72.5	43.7	43.1	65.1	42.6
LCCP [34]	58.4	89.1	63.8	53.6	82.6	60.2
V4R [35]	65.3	81.4	69.5	62.5	81.4	66.6
Ours	88.8	81.7	84.1	83.0	67.2	73.3

Figure 3: Comparison with baselines on the ARID20 and YCB10 subsets of OCID [10].

5.2 Datasets

We use TOD to train all of our models. We do not fine tune on any real data. We evaluate quantitatively and qualitatively on the OCID dataset [10] and the OSD dataset [11], which have 2346 images of semi-automatically constructed labels and 111 manually labeled images, respectively. We note that OCID contains images with objects on a tabletop, and images with objects on a floor. Despite our method being trained in tabletop settings, it generalizes to floor settings as well.

Method	Input	OCID [10]						OSD [11]					
		Overlap			Boundary			Overlap			Boundary		
		P	R	F	P	R	F	P	R	F	P	R	F
Mask RCNN	RGB	70.7	25.6	30.2	45.0	14.8	16.7	67.3	27.8	34.9	24.9	10.0	12.0
Mask RCNN	Depth	83.5	75.2	78.1	73.4	59.3	64.4	74.4	72.1	72.1	49.8	41.5	44.0
Mask RCNN	RGB-D	80.8	73.9	76.1	68.2	58.4	61.8	74.4	72.7	73.4	53.1	48.1	49.8
Ours	DSN: RGB-D	84.2	57.6	62.2	72.9	44.5	49.9	72.2	63.7	66.1	58.5	43.4	48.4
Ours	DSN: Depth	88.3	78.9	81.7	82.0	65.9	71.4	80.7	80.5	79.9	66.0	67.1	65.6

Table 1: Evaluation of our method against state-of-the-art instance segmentation algorithm Mask RCNN trained on different input modes.

RRN training data	OCID [10]						OSD [11]					
	Overlap			Boundary			Overlap			Boundary		
	P	R	F	P	R	F	P	R	F	P	R	F
TOD	88.3	78.9	81.7	82.0	65.9	71.4	80.7	80.5	79.9	66.0	67.1	65.6
OID [37]	87.9	79.6	81.7	84.0	69.1	74.1	81.2	83.3	81.7	69.8	73.7	70.8

Table 2: Comparison of RRN when training on TOD and real images from Google OID [37].

5.3 Metrics

We use the precision/recall/F-measure (P/R/F) metrics as defined in [17]. This metric promotes methods that segment the desired objects and penalize methods that provide false positives. Specifically, the precision, recall, and F-measure are computed between all pairs of predicted mask and ground truth masks. The Hungarian method is used to compute a matching between predictions and ground truth. Given this matching, the final P/R/F is computed by

$$P = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_i |c_i|}, \quad R = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_j |g_j|}, \quad F = \frac{2PR}{P + R}$$

where c_i denotes the set of pixels belonging to predicted object i , $g(c_i)$ is the set of pixels of the matched ground truth object of c_i , and g_j is the set of pixels for ground truth object j . We denote this as Overlap P/R/F. See [17] for more details.

While the above metric is quite informative, it does not take object boundaries into account. To remedy this, we introduce a Boundary P/R/F measure to complement the Overlap P/R/F. Using the same Hungarian matching used to compute Overlap P/R/F, we compute Boundary P/R/F by

$$P = \frac{\sum_i |c_i \cap D[g(c_i)]|}{\sum_i |c_i|}, \quad R = \frac{\sum_i |D[c_i] \cap g(c_i)|}{\sum_j |g_j|}, \quad F = \frac{2PR}{P + R}$$

where we overload notation and denote c_i, g_j to be the set of pixels belonging to the boundaries of predicted object i and ground truth object j , respectively. $D[\cdot]$ denotes the dilation operation, which allows for some slack in the prediction. Roughly, these metrics are a combination of the \mathcal{F} -measure in [36] along with the Overlap P/R/F as defined in [17].

We report all P/R/F measures in the range $[0, 100]$ ($P/R/F \times 100$).

5.4 Quantitative Results

Comparison to baselines. We compare to baselines shown in [10], which include GCUT [32], SCUT [33], LCCP [34], and V4R [35]. In [10], these methods were only evaluated on the ARID20 and YCB10 subsets of OCID, so we compare our results these subsets as well. These baselines are designed to provide over-segmentations (i.e., they segment the whole scene instead of just the objects of interest). To allow a more fair comparison, we set all predicted masks smaller than 500 pixels to background, and set the largest mask to table label (which is not considered in our metrics). Results are shown in Figure 3. Because the baselines aim to over-segment the scene, the precision is in general low while the recall is high. LCCP is designed to segment convex objects (most objects in OCID are convex), but its predicted boundaries are noisy due to operating on depth only. Both SCUT and V4R utilize models trained on real data. V4R was trained on OSD [11] which has an extremely similar data distribution to OCID, giving V4R a substantial advantage. Our method, despite never having seen any real data, significantly outperforms these baselines on F-measure.

DSN	IMP		RRN	Boundary			Method	Input	RRN	Boundary		
	O/C	CCC		P	R	F				P	R	F
✓				35.0	58.5	43.4	Mask RCNN	RGB		24.9	10.0	12.0
✓			✓	36.0	48.1	39.6	Mask RCNN	RGB	✓	52.9	26.1	31.6
✓	✓			49.2	55.3	51.7	Mask RCNN	Depth		49.8	41.5	44.0
✓	✓		✓	59.0	64.1	60.7	Mask RCNN	Depth	✓	64.7	52.4	56.4
✓	✓	✓		53.8	54.7	53.6	Mask RCNN	RGB-D		53.1	48.1	49.8
✓	✓	✓	✓	66.0	67.1	65.6	Mask RCNN	RGB-D	✓	60.6	54.7	56.5

Table 3: (left) Ablation experiments on OSD [11]. O/C denotes the Open/Close morphological transform, while CCC denotes Closest Connected Component of the IMP module. (right) Refining Mask RCNN results with our RRN on OSD.

Effect of input mode. Next, we evaluate how different input modes affect the results by training Mask RCNN [6] on different combinations of RGB and depth from TOD. In Table 1, we compare Mask RCNN trained on RGB, depth, and RGB-D and compare it to our proposed model on the full OCID and OSD datasets. Training Mask RCNN on synthetic RGB only does not generalize at all. Training on depth drastically boosts performance. When training on RGB-D, we posit that Mask RCNN relies heavily on depth since adding RGB to depth results in little change. However, our method (line 5, Table 1) exploits RGB and depth separately, leading to better results while being trained on the exact same synthetic dataset. Furthermore, we show that when our DSN is trained on RGB concatenated with depth (line 4, Table 1), we see a drop in performance, suggesting that training directly on (non-photorealistic) RGB is not the best way of utilizing the synthetic data.

Degradation of training on non-photorealistic simulated RGB. To quantify how much non-photorealistic RGB degrades performance, we train an RRN on real data. This serves as an approximate upper bound on how well the synthetically-trained RRN can perform. We use the instance masks from the Google Open Images dataset (OID) [37] and filter them to relevant object classes, resulting in roughly 220k instance masks on real RGB images. Results are shown in Table 2. Both models share the same DSN and IMP. The Overlap measures are roughly the same, while the RRN trained on OID has slightly better performance on the Boundary measures. This suggests that while there is still a gap, our method is surprisingly not too far off, considering that we train with non-photorealistic synthetic RGB. We conclude that mask refinement with RGB is an easier task to transfer from Sim2Real than directly segmenting from RGB.

Ablation studies. We report ablation studies on OSD to evaluate each component of our proposed method in Table 3 (left). We omit the Overlap P/R/F results since they follow similar trends to Boundary P/R/F. Running the RRN on the raw masks output by DSN without the IMP module actually hurts performance as the RRN cannot refine such noisy masks. Adding the open/close morphological transform and/or the closest connect component results in much stronger results, showing that the IMP is key in robustifying our proposed method. In these settings, applying the RRN significantly boosts Boundary P/R/F showing that it effectively sharpens the masks. In fact, Table 3 (right) shows that applying the RRN to the Mask RCNN results effectively boosts the Boundary P/R/F, showing the efficacy of the RRN despite being directly trained on non-photorealistic RGB. Note that even with this refinement, the Mask RCNN results are outperformed by our method.

5.5 Qualitative Results

We show qualitative results on OCID of baseline methods, Mask RCNN (trained on RGB-D), and our proposed method in Figure 4 (top). It is clear that the baseline methods suffer from over-segmentation issues; they segment the table and background into multiple pieces. For the methods that utilize RGB as an input (GCUT and SCUT), the objects are often over-segmented as well. Methods that operate on depth alone (LCCP and V4R) result in noisy object boundaries due to noise in the depth sensors. The main failure mode for Mask RCNN is that it tends to undersegment objects; a close inspection of Figure 4 shows that Mask RCNN erroneously segments multiple objects as one. This is the typical failure mode of top-down instance segmentation algorithms in clutter. On the other hand, our bottom-up method utilizes depth and RGB separately to provide sharp and accurate masks.

In Figure 4 (middle), we qualitatively show the effect of the RRN. The first row shows the initial masks after the IMP module and the second row shows the refined masks. These images were taken around our lab with an Intel RealSense D435 RGB-D camera to demonstrate the robustness of our method to camera viewpoint variations and distracting backgrounds (OCID/OSD have relatively simple backgrounds). Due to noise in the depth sensor, it is impossible to get sharp and accurate



Figure 4: Qualitative results. (Top) Comparison on OCID [10]. (Middle) Mask refinements on images taken from our lab. (Bottom) Failure modes on OSD [11].

predictions from depth alone without using RGB. Our RRN can provide sharp masks even when the boundaries of objects are occluding other objects (images 2 and 5).

We show some failure modes of our method in Figure 4 (bottom) on the OSD dataset. In images 1, 5, and 6, false positives contributed by the DSN cannot be undone by the RRN. In images 2 and 3, we see examples of missed objects due to the center of the mask being occluded by an object, which is a limitation of the 2D center voting procedure. Lastly, when an object is split into two (images 3, 4, and 5), our method predicts two separate objects.

5.6 Application in Grasping Unknown Objects

We use our model to demonstrate manipulation of unknown objects in a cluttered environment using a Franka robot with panda gripper and wrist-mounted RGB-D camera. The task is to collect objects from a table and put them in a bin. Objects are segmented using our method and the point cloud of the closest object to the camera is fed to 6-DOF GraspNet [38] to generate grasps, with other objects considered obstacles. Video of the experiments can be found at the project website. While our method segments the objects correctly most of the time, it does exhibit some failures, including the over-segmentation of the drill. We conducted the experiment 3 times and our method was able to successfully complete the task with 1-2 extra grasping attempts in each trial. The failures stem from either imperfections in segmentation or inaccurate generated grasps.

6 Conclusion

We proposed a framework that separately leverages RGB and depth to provide sharp and accurate masks for unseen object instance segmentation. Our two-stage framework produces rough initial masks using only depth, then refines those masks with RGB. Surprisingly, our RRN can be trained on non-photorealistic RGB and generalize quite well to real world images. We demonstrated the efficacy of our approach on multiple datasets for UOIS in tabletop environments.

Acknowledgements

This work was funded in part by an NDSEG fellowship and ONR grant 63-6094. We thank NVIDIA for generously providing a DGX used for this research via the NVIDIA Robotics Lab and the UW NVIDIA AI Lab (NVAAIL). We also thank Tim Patten for providing baseline results on OCID.

References

- [1] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012, 2015.
- [2] F. Zhang, J. Leitner, M. Milford, B. Upcroft, and P. Corke. Towards vision-based deep reinforcement learning for robotic motion control. In *Australasian Conference on Robotics and Automation (ACRA)*, 2015.
- [3] E. Tzeng, C. Devin, J. Hoffman, C. Finn, P. Abbeel, S. Levine, K. Saenko, and T. Darrell. Adapting deep visuomotor representations with weak pairwise constraints. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [4] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, et al. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [5] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [6] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [7] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. In *Robotics: Science and Systems (RSS)*, 2017.
- [8] M. Danielczuk, M. Matl, S. Gupta, A. Li, A. Lee, J. Mahler, and K. Goldberg. Segmenting unknown 3d objects from real depth images using mask r-cnn trained on synthetic data. In *IEEE Conference on Robotics and Automation (ICRA)*, 2019.
- [9] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019.
- [10] M. Suchi, T. Patten, and M. Vincze. Easylabel: A semi-automatic pixel-wise object annotation tool for creating robotic rgb-d datasets. In *IEEE Conference on Robotics and Automation (ICRA)*, 2019.
- [11] A. Richtsfeld, T. Mörwald, J. Prankl, M. Zillich, and M. Vincze. Segmentation of unknown objects in indoor environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [12] L.-C. Chen, A. Hermans, G. Papandreou, F. Schroff, P. Wang, and H. Adam. Masklab: Instance segmentation by refining object detection with semantic and direction features. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] B. De Brabandere, D. Neven, and L. Van Gool. Semantic instance segmentation with a discriminative loss function. *arXiv preprint arXiv:1708.02551*, 2017.
- [14] D. Neven, B. De Brabandere, M. Proesmans, and L. Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] D. Novotny, S. Albanie, D. Larlus, and A. Vedaldi. Semi-convolutional operators for instance segmentation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [16] C. Xie, Y. Xiang, Z. Harchaoui, and D. Fox. Object discovery in videos as foreground motion clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [17] A. Dave, P. Tokmakov, and D. Ramanan. Towards segmenting everything that moves. *arXiv preprint arXiv:1902.03715*, 2019.

- [18] P. O. Pinheiro, R. Collobert, and P. Dollár. Learning to segment object candidates. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [19] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár. Learning to refine object segments. In *European Conference on Computer Vision (ECCV)*, 2016.
- [20] W. Kuo, A. Angelova, J. Malik, and T.-Y. Lin. Shapemask: Learning to segment novel objects by refining shape priors. *arXiv preprint arXiv:1904.03239*, 2019.
- [21] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems (RSS)*, 2018.
- [22] 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.
- [23] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox. Deepim: Deep iterative matching for 6d pose estimation. In *European Conference Computer Vision (ECCV)*, 2018.
- [24] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel. Asymmetric actor critic for image-based robot learning. In *Robotics: Science and Systems (RSS)*, 2018.
- [25] F. Sadeghi and S. Levine. CAD2RL: Real single-image flight without a single real image. In *Robotics: Science and Systems(RSS)*, 2017.
- [26] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981.
- [27] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.
- [28] Y. Wu and K. He. Group normalization. In *European Conference on Computer Vision (ECCV)*, 2018.
- [29] S. Xie and Z. Tu. Holistically-nested edge detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [30] J. Serra. *Image analysis and mathematical morphology*. Academic Press, Inc., 1983.
- [31] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser. Semantic scene completion from a single depth image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [32] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision (IJCV)*, 2004.
- [33] T. T. Pham, T.-T. Do, N. Sünderhauf, and I. Reid. Scenecut: Joint geometric and object segmentation for indoor scenes. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [34] S. Christoph Stein, M. Schoeler, J. Papon, and F. Worgotter. Object partitioning using local convexity. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [35] E. Potapova, A. Richtsfeld, M. Zillich, and M. Vincze. Incremental attention-driven object segmentation. In *IEEE-RAS International Conference on Humanoid Robots*, 2014.
- [36] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [37] R. Benenson, S. Popov, and V. Ferrari. Large-scale interactive object segmentation with human annotators. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [38] A. Mousavian, C. Eppner, and D. Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.