Active Coarse-to-Fine Segmentation of Moveable Parts from Real Images

Ruiqi Wang¹, Akshay Gadi Patil¹, Fenggen Yu¹, and Hao Zhang^{1,2}

¹ Simon Fraser University, Burnaby, Canada ² Amazon



Fig. 1: Our instance segmentation of moveable parts, with semantic labels, on real-world photos. Comparison is made with OPDFormer-C (OPD = *openable* part detection), the current state of the art, where small red \times s indicate erroneous or missed labels. With generalization to *non-openable* parts, e.g., on lamps and bottles (top right). As an application of accurate moveable part segmentation, we can manipulate 3D reconstructions of articulated objects (bottom right).

Abstract. We introduce the first *active learning* (AL) model for high-accuracy instance segmentation of moveable parts from RGB images of *real indoor scenes*. Specifically, our goal is to obtain *fully validated* segmentation results by humans while *minimizing manual effort*. To this end, we employ a transformer that utilizes a masked-attention mechanism to supervise the active segmentation. To enhance the network tailored to moveable parts, we introduce a *coarse-to-fine* AL approach which first uses an *object-aware* masked attention and then a *pose-aware* one, leveraging the hierarchical nature of the problem and a correlation between moveable parts and object poses and interaction directions. When applying our AL model to 2,000 real images, we obtain fully validated moveable part segmentations with semantic labels, by only needing to manually annotate 11.45% of the images. This translates to significant (60%) time saving over manual effort required by the best non-AL model to attain the same segmentation accuracy. At last, we contribute a dataset of 2,550 real images with annotated moveable parts, demonstrating its superior quality and diversity over the best alternatives.

1 Introduction

Most objects we interact with in our daily lives have dynamic movable parts, where the part movements reflect how the objects function. Perceptually, acquiring a visual and



Fig. 2: Overview of our pose-aware masked attention network for moveable part segmentation of articulated objects in real scene images. Utilizing a two-stage framework, we first derive a *coarse* segmentation by predicting the object mask, its 6 DoF pose, and the interaction direction, subsequently isolating the interaction surface of the objects. In the *fine* segmentation stage, we combine the object mask and interaction surface to form a refined mask, enabling the extraction of fine-grained instance segmentation of moveable parts.

actionable understanding of object functionality is a fundamental task. In recent years, motion perception and functional understanding of articulated objects have received increasing attention in vision, robotics, and VR/AR. Aside from per-pixel or per-point motion prediction, the detection and segmentation of moveable parts plays a vital role in embodied AI applications involving robot manipulation and action planning.

In this paper, we tackle the problem of *instance* segmentation of moveable parts in one or more articulated objects from RGB *images of real indoor scenes*, as shown in Figure 1. Note that we use the term articulated objects in a somewhat loose sense to refer to all objects whose parts can undergo motions; such motions can include opening a cabinet door, pulling a drawer³, and the movements of a lamp arm. Most prior works on motion-related segmentations [11, 17, 39] operate on point clouds, which are more expensive to capture than images while having lower resolution, noise, and outliers. Latest advances on large language models (LLMs) and vision-language models (VLMs) have led to the development of powerful generic models such as SAM [16], which can excel at generating quality object masks and exhibit robust zero-shot performance across diverse tasks owing to their extensive training data. However, these methods remain limited in comprehensive understanding of moveable object parts.

To our knowledge, OPD [14], for "openable part detection", and its follow-up, OPD-Multi [28], for "openable part detection for multiple objects", represent the state of the art in moveable part segmentation from images. However, despite the fact that both methods were trained on real object/scene images, there still remains a large gap between synthetic and real test performances: roughly 75% vs. 30% in segmentation accuracy [28]. The main reason is that manual instance segmentation on real images to form ground-truth training data is too costly. As a remedy, OPD and OPDMulti both opted to manually annotate 3D mesh or RGB-D *reconstructions* from real-world articulated object scans and project the obtained segmentation masks to 2D. Thus, for

³ Strictly speaking, articulations are realized by "two or more sections connected by a *flexible* joint," which would not include drawer sliding. However, as has been done in other works in vision and robotics, we use the term loosely to encompass more general part motions.

each reconstructed 3D scene, there is only a one-time annotation, in 3D, required, after which thousands of annotated images can be rendered. Clearly, such an indirect annotation still leaves a gap between rendered images of digitally reconstructed 3D models and real photographs, with both reconstruction errors and re-projection errors due to view discrepancies hindering the annotation quality on images.

To close the aforementioned gap by addressing the annotation challenge, we present an *active learning* (AL) [2, 24, 40] approach to obtain high-accuracy instance segmentation of moveable parts, with semantic labels, *directly on* real scene images containing one or more articulated objects. AL is a semi-supervised learning paradigm, relying on human feedback to continually improve the performance of a learned model.

Specifically, our goal in this work is to obtain *fully validated* segmentation results by humans while *minimizing manual segmentation efforts*. In other words, we would like the human to manually segment as few images as possible while ensuring that *all* the images in our dataset have been segmented accurately, either by a neural segmentation network that is trained by available ground-truth data or by human. To this end, we employ a transformer-based [7] segmentation network that utilizes a masked-attention mechanism [6]. To enhance the network for moveable part segmentation, we introduce a *coarse-to-fine* AL model which first uses an *object-aware* masked attention and then a *pose-aware* one, leveraging the hierarchical nature of the problem and a correlation between moveable parts and object poses and interaction directions.

As shown in Figure 2, in the *coarse* annotation stage, our AL model with objectaware attention predicts object masks, poses, and interaction directions, so as to help isolate interaction surfaces on the articulated objects. In the *fine* annotation stage, we combine the object masks and interaction surfaces to predict refined segmentation masks for moveable object parts, also with human-in-the-loop. Unlike prior works on active segmentation [29,37] which mainly focused on the efficiency of human annotations using point- or region-based supervision for fast labeling, we optimize the human-in-theloop pipeline to reduce AL iterations and samples required for manual annotation. Our network learns the regions-of-interests (ROIs) from the pose-aware masked-attention decoder for better segmentation sampling in AL iterations in the second stage, where we categorize samples in different branches for further training, testing, and annotation.

In summary, our main contributions include:

- We introduce the first AL framework for instance segmentation of moveable parts from RGB images of real indoor scenes. When applying our AL model to 2,000 real images, we obtain fully validated moveable part segmentations with semantic labels, by only needing to manually annotate 11.45% of the images. This translates to significant (60%) time saving over manual effort required by the best non-AL model, i.e., OPDFormer-C [28], to attain the same segmentation accuracy.
- Our coarse-to-fine AL model, with both object- and pose-aware masked-attention mechanisms, lead to reduced human effort and improved accuracy in moveable part segmentation over state-of-the-art (SOTA) methods: OPD [14] and OPDMulti [28].
- Our scalable AL model allows us to accurately annotate a dataset of 2,550 real photos of articulated objects in indoor scenes. We show the superior quality and diversity of our new dataset over current alternatives [14, 28], and the resulting improvements in segmentation accuracy.

2 Related Works

Articulated objects dataset. The last few years have seen the development of articulation datasets on 3D shapes. Of the many, ICON [10] build a dataset of 368 moving joints corresponding to various parts of 3D shapes from the ShapeNet dataset [5]. The Shape2Motion dataset [33] provides kinematic motions for 2,240 3D objects across 45 categories sourced from ShapeNet and 3D Warehouse [1]. The PartNet-Mobility dataset [36] consists of 2,374 3D objects across 47 categories from the PartNet dataset [21], providing motion annotations and part segmentation in 3D.

All these datasets are obtained via manual annotations and are *synthetic* in nature. Since sufficient training data is made available by these synthetic datasets, models trained on them can be used for fine-tuning on *real-world* 3D articulated object datasets with limited annotations. However, models trained exclusively on *synthetic* data cannot generalize well under real-world scenarios. Bridging the synthetic-real data gap remains a reoccurring challenge; see the supplementary material for details.

Recently, OPD [14] and its follow-up work OPDMulti [28], provide two 2D image datasets of real-world articulated objects: OPDReal and OPDMulti. In OPDReal, images are obtained from frames of RGB-D scans of indoor scenes containing a single object. OPDMulti, on the other hand, captures multiple objects. Both datasets come with 2D segmentation labels on all *openable* parts along with their motion parameters. However, due to the nature of annotation process, the 2D part segmentation masks obtained via 3D-to-2D projection do not fully cover all openable parts in the image. Also, in OPDReal, objects are scanned from within a limited distance range. Practical scenarios and use cases are likely going to have large camera pose and distance variations. Although OPDMulti incorporates such viewpoint variations, a large portion of it contains frames without any articulated objects, which directly affects trained models.

To overcome these limitations, we contribute a 2D image dataset of moveable objects present in the real world (furniture stores, offices, homes), captured using iPhone 12, 12 Pro and 14. We then use our *coarse-to-fine* AL framework (Figure 3 and Section 4) to learn generalized 2D segmentations for moveable object parts.

Part segmentation in images. Early approaches [31, 32, 35] to 2D semantic part segmentation developed probabilistic models on human and animal images. While not addressing the 2D semantic part segmentation problem as such, [3, 12, 15, 20, 22] tackled the problem of estimating 3D articulations from human images, which requires an understanding of articulated regions in the input image.

Recently, the development of large visual models, such as SAM [16], has addressed classical 2D vision tasks, such as object segmentation, surpassing all existing models. Such large pre-trained models can be directly employed for *zero-shot* segmentation on new datasets. Follow-up works [8, 43] to SAM aim at multi-modal learning by generalizing to natural language prompts. For the task of moveable part segmentation in real scene images, we observe an unsatisfactory performance using such models. This is expected since they were never trained on any moveable parts datasets, and therefore, lack an understanding of articulated objects. To our knowledge, OPDMulti [28] is the SOTA model that can segment moveable parts in an input image, and is built on



Fig. 3: Our coarse-to-fine Active Learning (AL) training pipeline. The *coarse* AL applys on interaction directions and retains high-quality predictions while manually rectifying the rest. These rectified predictions form a constructive prior for refined mask prediction. Subsequently, the *fine* AL stage utilizes these refined masks, employing an iterative training method with continuous human intervention for accurate part mask annotation.

the Mask2Former architecture [6]. In our work, we use a transformer architecture in a *coarse-to-fine* manner to obtain moveable part segmentation (Section 4).

Active learning for image segmentation. Active learning (AL) is a well-known technique for improving model performance with limited labeled data. This, in turn, allows the expansion of labeled datasets for downstream tasks. Prior works [4,25–27,38] have presented different AL frameworks to acquire labels with minimum cost for 2D segmentation tasks. There exist AL algorithms for such tasks [23,34] that are specifically designed to reduce the domain gap by aligning two data distributions. We cannot borrow such methods to reduce the domain gap between synthetic and real scene images of moveable objects because of the large feature differences: our synthetic images contain a single object without background or texture, and most objects have an empty interior.

More recently, [29,37] employed AL to refine initial 2D segmentation masks through key point or region selection, requiring little human guidance. These works focus on minimize labeling efforts over the prediction results, whose supervisions effectively correct prediction errors in object masks. However, due to potentially multiple moveable parts, such point/region selection is ambiguous for articulated objects. As such, we design an AL framework based on our two-stage network that reduces manual effort by focusing on: (a) using an improved part segmentation model for generating better samples (Section 4.1), and (b) employing a *coarse-to-fine* strategy to optimize the AL working flow (Section 4.2).

3 Problem Statement

Given a set of images D captured from the real-world scene, our input is a single RGB image $I \in D$ containing one or more articulated objects o_i from one or more categories $c_i \in \{\text{cabinet}, \text{dishwasher}, \text{fridge}, \text{microwave}, \text{oven}, \text{washer}\}$. We assume that each object o_i has more than one moveable parts $P = \{p_1, \ldots, p_k\}$ according to its functionality. Our first goal is to predict the 2D bounding box b_i , the segmentation mask m_i represented by a 2D polygon and the semantic label $l_i \in \{\text{door}, \text{drawer}\}$ for each moveable part. Extending the above goal, we also aim to build a labeled image dataset that provides accurate 2D segmentation masks and labels for all p_i 's, for all $I \in D$.

4 Method

To address the above problem, we propose an active learning setup that consists of a transformer-based learning framework coupled with a human-in-the-loop feedback process. To this end, we present an end-to-end pose-aware masked-attention network (Fig 2) that works in a *coarse-to-fine* manner for part segmentation and label prediction. By making use of *coarse* and *fine* features from the network, segmentation masks are further refined by humans in the AL setup (Fig 3), resulting in precise moveable part masks, while minimizing human efforts spent on manual segmentation.

4.1 Pose-aware masked-attention network

Fig 2 provides a comprehensive depiction of our network architecture, encompassing two distinct stages. In the *coarse* stage, the network processes a single RGB image and computes a refined mask based on outputs from multiple heads, which accurately pinpoints the region containing moveable parts. This stage filters out noise predictions on background and extraneous portions of the object. Subsequently, *fine* stage takes the refined mask and image features to generate part masks, bounding boxes, and semantic labels for all moveable parts of all articulated objects in the images.

Coarse stage. There are three steps in *coarse* stage. First, the input image is passed through a backbone object detector network based on MaskRCNN [9], producing multiscale feature maps f and 2D object bounding boxes b^o . A pixel decoder [42] upsamples f for subsequent processing in the *fine* stage. Second, we use a modified version of the multi-head attention-based encoder and decoder [42] to process f. Inspired by [13], we replace the original object query embedding module in [42] by our new object query embedding with normalized centre coordinates (c_x, c_y) , width and height (w, h) from the detected 2D bounding box, enabling the decoder to generate new object query embeddings containing both local global information and estimate 6DoF pose from the 2D bounding box. Third, the decoded queries are passed into multiple task-specific MLP heads for (a) object class prediction, (b) 6DoF object pose estimation, (c) object interaction direction prediction and (d) object mask prediction.

We obtain the object class with a fully connected network with 3 layers followed by a softmax activation. For 6DoF pose estimation, we use two identical MLP heads with 3 linear layers with ReLU activation and different output dimensions – one for estimating camera translation $\tilde{t} = (\tilde{t}_x, \tilde{t}_y, \tilde{t}_z)$, and the other for estimating the camera rotation matrix $\tilde{R} \in SO(3)$ as described in [41].

The MLP head for interaction direction prediction outputs a set of 6 possible interaction directions $d \in \{\pm x, \pm y, \pm z\}$ corresponding to the 6DoF coordinates. Using b^o and the estimated 6DoF object pose, we can obtain the corresponding 3D *oriented* bounding box B^o , which tightly fits the b^o . From among the eight vertices in B^o , we select vertices of the face along the interaction direction as the representative 2D box for the interaction surface, and use it to crop the input image. This cropped image is further multiplied with the object 2D binary mask to filter out background pixels, obtaining the refined binary object mask m_r , which guides the subsequent *fine* stage to focus exclusively on the relevant features of the articulated object. **Fine stage.** There is just one component to this stage, which is the masked-attention decoder from Mask2Former [6] (see Figure 2). It is made up of a cascade of three identical layers, L_i 's. L_1 takes as input image features f_{pd} and the refined mask, m_r , and outputs a binary mask which is fed to the next layer. Eventually, the binary mask at the output of L_3 is multiplied with f_{pd} resulting in moveable part segmentation in the RGB space. We call this our *pose-aware masked-attention decoder*.

Loss functions. We formulate the training loss as below

$$L = L_{class} + L_{dir} + L_{om} + L_{pos} + L_{fine} \tag{1}$$

where L_{class} is the binary-cross entropy for object class prediction, L_{dir} is the crossentropy loss for interaction direction prediction, L_{om} is the binary mask loss for object mask prediction. We define the loss for pose estimation as $L_{pos} = \lambda_t L_t + \lambda_{rot} L_{rot}$, where L_t is the L2-loss of the translation head and L_{rot} is the geodesic loss [19] of the rotation head. We set λ_t and λ_{rot} to 2 and 1 respectively. We use a pixel-wise crossentropy loss for the *fine* stage.

When pre-training, we jointly train our two-stage network in an end-to-end fashion (see Section 6). During fine-tuning on real images with part annotations, we fix MLP weights since ground truth poses and object masks are not available.

4.2 Coarse-to-fine active learning strategy

Our active learning setup, consisting of human-in-the-loop feedback, unfolds in a coarseto-fine manner (see Figure 3). We independently run AL workflow on outputs of both *coarse* and *fine* stages from Section 4.1.

In *coarse* AL part, *Coarse* stage generates predictions for the *test set*. In our experiment setup, the *test set* is the *enhancement set*. During this phase, users validate interaction direction predictions and rectify inaccuracies. With ground-truth interaction directions established, refined masks m_r are computed and input into the *fine* stage.

In *fine* AL part, part segmentation mask and label outcome from the *Fine* stage are subject to user evaluation, categorized as perfect, missed, or fair. Specifically: i) A perfect prediction implies coverage of all moveable parts in the final segmentation masks, without any gaps, as well as accurate class labels for each segmented part; ii) A missed prediction effectively refers to a null segmentation mask, and/or erroneous class labels; iii) A fair prediction denotes an output segmentation mask that may exhibit imperfections such as gaps or rough edges, and/or may have inaccuracies in some part class labels. We provide extensive examples of these scenarios in our supplements. During the AL process, perfect predictions, we employ the labelme [30] annotation interface to manually annotate the part mask polygons, and include such images in the next-iteration training set. Fair predictions, on the other hand, remain in the *test set* for re-evaluation.

The AL workflow on the *fine* stage continues iteratively until all images within the *test set* transition to the training set, becoming well-labeled and eventually leaving the test set vacant. Benefiting from the verified ground-truth interaction direction established in the *coarse* AL part, the *Fine* stage hones in on features of the target surface,

Table 1: Dataset statistics across six articulated object categories for OPDReal, OPDMulti and our datasets. Microwave and Oven categories are merged due to their co-occurrence in real scenes. Compared to OPDReal, our dataset is relatively more balanced in terms of sample distribution of different categories. OPDMulti does not provide category-wise information, and only 19K out of 64K total images are valid with target and annotation. Parts/img shows the average parts annotated for each image. Our dataset exhibits the most object and part diversity.

				Category	/		Total	Parts/img
		Storage	Fridge	Dishwasher	Mic.&Oven	Washer		
OPDReal [14]	Objects	231	12	3	12	3	284	
	Images	27,394	1,321	186	823	159	30K	2.22
	image %	91.67%	3.93%	0.62%	2.75%	0.53%	100%	
	Parts	787	27	3	13	3	875	
OPDMulti [28]	Objects	-	-	-	-	-	217	
	Images	-	-	-	-	-	19K/64K	1.71
	Parts	-	-	-	-	-	688	
Ours	Objects	176	51	31	62	13	333	
	Images	925	370	315	775	175	2550	4.22
	image %	36.27%	14.51%	12.35%	30.39%	6.8%	100%	4.55
	Parts	896	159	31	62	13	1161	

omitting noisy object parts. This streamlined focus notably expedites the annotation process. Further insights into the human verification and annotation procedures will be provided in our supplementary materials.

5 Datasets and Metrics

Datasets. We use three real image datasets in our experiments: (1) OPDReal [14], (2) OPDMulti [28], and (3) our dataset. Our dataset images are obtained from the real world by taking photographs of articulated objects in indoor scenes from furniture stores, offices, and homes, captured using iPhone 12, 12Pro and 14. Images are captured with varying camera poses and distances from the objects, and an image can contain more than one object, with multiple moveable parts per object. Differences to OPD and OPD-Multi datasets are explained in Section 2.

We consider six object categories: Storage, Fridge, Dishwasher, Microwave, Washer, and Oven. A comparison of dataset statistics is presented in Table 1. OPDReal comprises of \sim 30K images, with each image depicting a single articulated object. OPD-Multi contains \sim 64K images. Among these, only 19K images are considered "valid", containing at least one articulated object. Our dataset has a total of 2,550 images, with each image showcasing objects from several categories. We organize our dataset according to the primary object depicted in each image. Our dataset stands out by offering the highest diversity of objects and parts among the compared datasets, including 333 different articulated objects and 1,161 distinct parts.

In terms of the moveable part annotation, both OPDReal and OPDMulti generate annotations on a 3D mesh reconstructed from the RGB-D scans, and project these 3D annotations back to the 2D image space to get 2D part masks. This process is prone to reconstruction and projection errors. We, on the other hand, create annotations on the captured images directly using our *coarse-to-fine* active learning framework. See the supplementary material for annotation quality comparisons.

Table 1 shows that the majority (91.67%) of data samples in OPDReal belong to the Storage category, with the rest distributed among the remaining categories. In contrast, our dataset offers a more uniform data distribution across all six categories.

Metrics. To evaluate model performance and AL efficiency, we use the following:

- Mean Average Precision (mAP): We report mAP@IoU=0.5 for correctly predicting the part label and 2D mask segmentation with IoU ≥ 0.5. This metric, which is applied to 2D mask segmentation, is more precise for evaluating segmentation quality than BBox mAP used by OPD [14], which only assesses boundary accuracy and overlooks finer details such as mask edges and internal holes.
 - The *ground-truth* (GT) segmentations over an image dataset to measure mAP are obtained by applying AL over the dataset with full validation by humans.
- AL iterations: We report the number of iterations required during active learning. This metric represents the efficiency of the overall AL pipeline.
- Annotated images: We report the numbers of images and corresponding parts required for manual annotation for each iteration during AL. This metric helps us evaluate the efficiency of the AL sampling process.
- Total lab time: We report the total lab time required for labeling a dataset. For methods which employ AL, it includes time spent on compulsory sampling after each iteration and manual annotation in each iteration. For methods without AL, it calculates the time spent on manual annotation of all failed predictions. This metric provides an overview of human effort required for all methods. See Section 3.4 in the Supplementary Materials for details of human efforts in our AL process.

6 Experiments

We start our experiments by rendering synthetic images from the PartNet-Mobility dataset [36] with diverse articulation states, enabling us to obtain sufficient annotations for training 2D segmentation networks and support transfer learning applications. The synthetic dataset contains \sim 32K images, evenly distributed across categories, and randomly partitioned into training (90%) and test sets (10%).

We implement our network in PyTorch on two NVIDIA Titan RTX GPUs. All images are resized to 256×256 for training. For pre-training on PartNet-Mobility, we use the Adam optimizer with an initial learning rate (*lr*) of 2.5*e*-4, reducing it by $\gamma = 0.1$ at 1K and 1.5K epochs separately over a total of 2K epochs. When fine-tuning on real images, we use the same *lr* and γ at 3.5K and 4K epochs over a total of 4.5K epochs.

6.1 Competing methods

We compare our active *coarse-to-fine* part segmentation model with three 2D segmentation methods and also analyze two variants of our proposed approach.

Table 2: Comparing segmentation accuracy against competing methods and variants of our method on the *unseen test set* of 2,000 real images. In the table, "AL" indicates whether the method uses active learning. All methods take the *train* set as the training data, and are evaluated on the *test set*. Methods in the last two columns perform AL on the *enhancement set*. The "Time" row represents the total lab time metric described in Section 5, only for methods using AL.

	Method								
	Grounded-SAM [8]] OPD-C [14]	OPDFormer-C [28]	$Ours_{w/oAL}$	$Ours_{f-AL}$	Ours			
AL	-	-	-	-	\checkmark	\checkmark			
segm mAP ↑	23.1	45.2	68.4	77.3	91.2	91.3			
Time (hr) \downarrow	-	-	-	-	1.675	1.6			

- Grounded-SAM [8], which combines Grounding-DINO [18] and Segment Anything [16], is a foundational vision-language model that can be used for zero-shot 2D object detection and segmentation, and supports text prompts. In our experiments, we set the text prompt as [door, drawer] for segmentation results.
- OPD-C [14], which is the first work for detecting openable parts in images based on MaskRCNN [9]. This is the base variant without camera pose for training.
- OPDFormer-C [28], a follow-up of OPD-C based on Mask2Former [6], is the SOTA for openable part detection of multiple articulated objects in images.
- **Ours**_{w/oAL} is a variant that does not use human feedback it infers part segmentation results based only on the transformer-based model.
- **Ours**_{f-AL} is variant of our approach that uses only the*fine*stage of the AL framework. That is, verification and annotation of just the part masks is done.</sub>

6.2 Evaluation on Our Dataset

We perform two key evaluations on our dataset, one for segmentation accuracy and one for annotation efficiency, while comparing to SOTA alternatives, with or without AL.

We work on 2,550 images with a split of 50/500/2000 into *train/enhancement/test* sets. The train set has been fully annotated manually and it is used by all methods, except for Grounded-SAM, for fine-tuning. The *enhancement set*, initially unlabeled, is employed by AL models to progressively improve the learning. The *test set* of 2,000 images is unseen by all methods, including AL, when evaluating segmentation accuracy. When assessing annotation efforts, we apply the methods on both the 500-image set and the 2,000-image set to examine how the efficiency achieved by our AL model scales.

Segmentation accuracy on test set. Table 2 compares four non-AL and two AL methods. Among four non-AL methods (columns 1-4), Grounded-SAM is without finetuning and has the lowest performance. This demonstrates that current generic large foundational models are still limited in understanding object parts without adequate training on well-labeled data. Despite the small (50-image) *train set*, models fine-tuned on it produce significant improvements. Specifically, $Ours_{w/oAL}$ model surpasses all competing methods with over 75% segmentation mAP, while OPDFormer-C falls short of 70%, and OPD-C scores below 50%. This discrepancy stems from the architectural

Table 3: Comparison of manual segmentation efforts required for different methods in annotating segmentation masks for two sets of images of different sizes. In the table, "AL" indicates whether the method uses active learning for labeling. All methods are trained on the original 50-image *train* set, when annotating the 500-image set, When annotating the 2,000-image set, we add the 500 images with ground-truth segmentations to the train set (50+500=550 images).

Dataset size Row ID		Method	AL	Iterations	#Images / Parts.	Time (hr) \downarrow
	1	Grounded-SAM	-	-	479 / 1,711	7.54
	2	OPD-C	-	-	483 / 1,720	7.58
	3	OPDFormer-C	-	-	324 / 1,248	5.62
500 imagaa	4	$Ours_{w/oAL}$	-	-	210 / 762	3.58
500 images	5	OPD-C	\checkmark	7	261 / 887	5.7
	6	OPDFormer-C	\checkmark	7	158 / 560	3.9
	7	$Ours_{f-AL}$	\checkmark	3	83 / 239	1.675
	8	Ours	\checkmark	3	64 / 184	1.6
	1	Grounded-SAM	-	-	1,888 / 8,119	35.5
	2	OPD-C	-	-	1,420 / 6,391	28.3
2,000 images	3	OPDFormer-C	-	-	792 / 3,511	16.3
	4	$Ours_{w/oAL}$	-	-	625 / 2,815	13.4
	5	$Ours_{f-AL}$	\checkmark	4	289 / 983	7.5
	6	Ours	\checkmark	4	229 / 158	6.5

designs of OPD-C and OPDFormer-C, which were built on vanilla MaskRCNN and Mask2Former for general segmentation tasks but fail to capture the nuances of articulated objects, where movable parts are closely tied to object pose and interaction directions. In contrast, our network effectively leverages the hierarchical structure of the scene, objects, and parts therein, resulting in a much better performance.

As seen in the last two columns of Table 2, by performing AL on the *enhancement set*, the performance is significantly boosted over non-AL methods, reaching over 90% accuracy, with less than 1.7 hours spent on manual segmentation. Figure 6 shows qualitative results of different methods on our *test set*.

The segmentation accuracy of our two AL alternatives is close since they share the identical network architecture. But they differ in AL training strategies, which impacts labeling efficiency. On the 500-image *enhancement set*, our *coarse-to-fine* AL strategy leads to a slight improvement (only 4.5%) on human annotation effort. We show next that on a larger set to perform AL, the improvement becomes more significant.

Annotation efficiency comparison. Table 2 shows that with 1.6 hours of manual segmentation to process images with missed predictions, our AL model is able to fully validate the moveable part segmentations and semantic labels for the 500-image *enhancement set*. To obtain the same GT annotations on this set, with respect to an non-AL method such as Grounded-SAM, one must manually correct all images with erroneous or imperfect segmentations. Specifically, Grounded-SAM could only yield less than 5% perfectly annotated images, with the rest (479) needing manual processing.

In Table 3 (top), we compare manual efforts, including number of images, parts, and lab times, required across different methods to obtain GT for the 500-image set.

Table 4: Ablation study on our key components.

Row ID Mask Pose Interaction direction AL segm mAP ↑ Time (hr) ↓

Table 5: Quantitative comparison against competing segmentation methods and our model variant on OPDReal, OPDMulti test set.

			-	68.4	-							
2 √	-	-	-	74.9	-							
3 √	\checkmark	√	-	77.3	-		segm mAP (↑)					
4 -	-	-	\checkmark	87.3	3.9	G	ounded-SA	M OPD-C C	PDFormer-C	¹ Ours		
5 √	-	-	√	89.1	3.1	01	ounded bit	morbee	or bronner c	, Ours _{w/oAL}		
6 -	\checkmark	\checkmark	√	90.8	2.4	OPDReal	16.5	44.5	46.3	51.6		
7 🗸	\checkmark	\checkmark	~	91.3	1.6	OPDMulti	8.0	25.6	27.6	31.5		

Ours_{*w*/*oAL*} shows the best efficiency among non-AL methods, but it still takes 3.58 hours to annotate 210 images with 762 parts. Rows 5-8 underscore the benefits of AL for annotation efficiency. By employing AL, in rows 5 & 6, OPD-C and OPDFormer-C demonstrate marked improvements over their non-AL versions. However, they still require 5.7 and 3.9 hours, respectively. Due to their tendency to generate noisy predictions on irrelevant parts of the object or background, most predictions are categorized as *fair* as described in Section 4.2, leading to more iterations in AL and additional time spent on sampling. In contrast, as shown in rows 7 & 8, both variants of our AL model complete in 3 iterations, with our *coarse-to-fine* AL methods requiring the least images for labeling and minimum time efforts.

In Table 3 (bottom), we report all the numbers to obtain GT annotations for a larger set of (2,000) images. What is most notable is that the efficiency gain in manual annotation time by our *coarse-to-fine* AL strategy has improved from less than 5%, for the smaller image set of 500, to more than 13% (6.5 hours vs. 7.5 hours). This demonstrates that our *coarse-to-fine* AL approach is particularly beneficial for large-scale annotation tasks, where the time saved on annotation significantly outweighs the extra time spent on sampling. Please check our supplement for detailed AL iterations.

Ablation study. Table 4 highlights the need and contributions of key components of our method on improving prediction accuracy and minimizing human efforts. Columns 2-5 respectively indicate the presence of: **Mask** object mask head; **Pose** object pose estimation head; **Interaction direction** interaction direction prediction head; **AL** active learning. Row 4 and 5 uses the *fine* AL stage on part mask alone due to the absence of pose and interaction direction prediction module. Row 6 and 7 use our *coarse-to-fine* AL strategy. Results in Table 4 clearly justify the *coarse-to-fine* design in our method, which gives the best performance (see row 7).

6.3 Evaluation on OPDReal and OPDMulti

In addition, we assess the performance of different models on the OPDReal and OPD-Multi dataset, using their respective train and test splits.

As shown in Table 5, $Ours_{w/oAL}$ method outperforms the rest. However, with more than 70% training data, all methods still fail to achieve >55% accuracy on OPDReal. This limitation primarily stems from data skewness towards the Storage category in OPDReal, which constitutes more than 90% of total samples, and results in poor generalization across other object categories. See category-wise results in the supplement.



Fig. 4: Qualitative results on OPDReal and OPDMulti test set. See supp. for more results.



Fig. 5: Part-level reconstruction and manipulation of the bottle and dishwasher

Performance on OPDMulti is further compromised by an abundance of noisy data in its test set [28]. From the qualitative results in Figure 4, we observe that some openable parts are cluttered or missed in the GT annotation, while our method accurately segments these parts. This discrepancy also contributes to the low accuracy.

7 Application

Our work demonstrate practical applications in part based reconstruction and manipulation of articulated objects from images. Given a set of multi-view RGB images of an articulated object, our model predicts precise segmentation masks of moveable parts in each image. This enables part based 3D reconstruction using masked images for both moveable parts and the main body of the object. The resulting 3D models of parts allow for easy manipulation of moveable parts to unseen states in 3D as shown in Figure 5.

8 Conclusion

We present the first active segmentation framework for high-accuracy instance segmentation of moveable parts in real-world RGB images. Our active learning framework, integrating human feedback, iteratively refines predictions in a *coarse-to-fine* manner, and achieves close-to-error-free performance on the test set. By leveraging correlations between the scene, objects, and parts, we demonstrate that our method can achieve state-of-the-art performance on challenging scenes with multiple cross-categories objects, and significantly reduce human efforts for dataset preparation.

Additionally, we contribute a high-quality and diverse dataset of articulated objects in real-world scene, complete with precise moveable part annotations. We will expand it further to support the vision community for understanding scene from images. We also hope our work catalyzes future motion- or functionality-aware vision tasks.



Fig. 6: Qualitative results on *test set* from our dataset. We visualize predictions results on different object categories using 3 competing methods and our final model. Our method outputs better segmentation masks over moveable parts across multiple objects in the image with clear separation of parts and small parts segmentation (Row 1, 4, 5). Our results also show that the *coarse-to-fine* segmentation framework can effectively reduce segmentation errors from unwanted objects (Row 2) and object side surfaces (Row 2, 3, 6, 8). More results in the supplementary materials.

Acknowledgements

This work was supported in part by NSERC. We thank all anonymous reviewers and area chairs for their valuable comments, Hanxiao Jiang, Hang Zhou for insightful discussion, and Mingrui Zhao for help with data annotation.

References

- Trimble Inc. 3D Warehouse. https://3dwarehouse.sketchup.com/ (2023), accessed: 2023-3-4
- Aggarwal, C.C., Kong, X., Gu, Q., Han, J., Yu, P.S.: Active learning: A survey. In: Data Classification: Algorithms and Applications, pp. 571–597 (2014)
- Ballan, L., Taneja, A., Gall, J., Van Gool, L., Pollefeys, M.: Motion capture of hands in action using discriminative salient points. In: Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VI 12. pp. 640–653. Springer (2012)
- 4. Casanova, A., Pinheiro, P.O., Rostamzadeh, N., Pal, C.J.: Reinforced active learning for image segmentation. In: International Conference on Learning Representations (2020)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1290–1299 (2022)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
- Grounded-SAM Contributors: Grounded-Segment-Anything (Apr 2023), https:// github.com/IDEA-Research/Grounded-Segment-Anything
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
- Hu, R., Li, W., Van Kaick, O., Shamir, A., Zhang, H., Huang, H.: Learning to predict part mobility from a single static snapshot. ACM Transactions on Graphics (TOG) 36(6), 1–13 (2017)
- Huang, J., Wang, H., Birdal, T., Sung, M., Arrigoni, F., Hu, S.M., Guibas, L.J.: Multibodysync: Multi-body segmentation and motion estimation via 3d scan synchronization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7108–7118 (2021)
- Huang, Z., Xu, Y., Lassner, C., Li, H., Tung, T.: Arch: Animatable reconstruction of clothed humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3093–3102 (2020)
- Jantos, T., Hamdad, M., Granig, W., Weiss, S., Steinbrener, J.: PoET: Pose Estimation Transformer for Single-View, Multi-Object 6D Pose Estimation. In: 6th Annual Conference on Robot Learning (CoRL 2022) (2022)
- Jiang, H., Mao, Y., Savva, M., Chang, A.X.: Opd: Single-view 3d openable part detection. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX. pp. 410–426. Springer (2022)
- Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7122–7131 (2018)

- 16 R. Wang et al.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
- Li, X., Wang, H., Yi, L., Guibas, L.J., Abbott, A.L., Song, S.: Category-level articulated object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3706–3715 (2020)
- Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
- Mahendran, S., Ali, H., Vidal, R.: 3d pose regression using convolutional neural networks. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 2174–2182 (2017)
- Mehta, D., Sridhar, S., Sotnychenko, O., Rhodin, H., Shafiei, M., Seidel, H.P., Xu, W., Casas, D., Theobalt, C.: Vnect: Real-time 3d human pose estimation with a single rgb camera. Acm transactions on graphics (tog) 36(4), 1–14 (2017)
- Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., Su, H.: PartNet: A largescale benchmark for fine-grained and hierarchical part-level 3D object understanding. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- 22. Mueller, F., Bernard, F., Sotnychenko, O., Mehta, D., Sridhar, S., Casas, D., Theobalt, C.: GANerated hands for real-time 3d hand tracking from monocular rgb. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 49–59 (2018)
- Ning, M., Lu, D., Wei, D., Bian, C., Yuan, C., Yu, S., Ma, K., Zheng, Y.: Multi-anchor active domain adaptation for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9112–9122 (2021)
- 24. Ren, P., Xiao, Y., Chang, X., Huang, P.Y., Li, Z., Gupta, B.B., Chen, X., Wang, X.: A survey of deep active learning (2020). https://doi.org/10.48550/ARXIV.2009.00236, https://arxiv.org/abs/2009.00236
- Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. In: International Conference on Learning Representations (2018)
- Shin, G., Xie, W., Albanie, S.: All you need are a few pixels: Semantic segmentation with pixelpick. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops. pp. 1687–1697 (October 2021)
- 27. Sinha, S., Ebrahimi, S., Darrell, T.: Variational adversarial active learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
- Sun, X., Jiang, H., Savva, M., Chang, A.X.: OPDMulti: Openable part detection for multiple objects. In: Proc. of 3D Vision (2024)
- Tang, C., Xie, L., Zhang, G., Zhang, X., Tian, Q., Hu, X.: Active pointly-supervised instance segmentation. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII. pp. 606–623. Springer (2022)
- Wada, K.: labelme: Image Polygonal Annotation with Python. https://github.com/ wkentaro/labelme (2016)
- Wang, J., Yuille, A.L.: Semantic part segmentation using compositional model combining shape and appearance. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1788–1797 (2015)
- Wang, P., Shen, X., Lin, Z., Cohen, S., Price, B., Yuille, A.L.: Joint object and part segmentation using deep learned potentials. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1573–1581 (2015)
- Wang, X., Zhou, B., Shi, Y., Chen, X., Zhao, Q., Xu, K.: Shape2motion: Joint analysis of motion parts and attributes from 3d shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8876–8884 (2019)

- Wu, T.H., Liou, Y.S., Yuan, S.J., Lee, H.Y., Chen, T.I., Huang, K.C., Hsu, W.H.: D 2 ada: Dynamic density-aware active domain adaptation for semantic segmentation. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIX. pp. 449–467. Springer (2022)
- Xia, F., Wang, P., Chen, X., Yuille, A.L.: Joint multi-person pose estimation and semantic part segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6769–6778 (2017)
- Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A.X., Guibas, L.J., Su, H.: SAPIEN: A simulated part-based interactive environment. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Xie, B., Yuan, L., Li, S., Liu, C.H., Cheng, X.: Towards fewer annotations: Active learning via region impurity and prediction uncertainty for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8068–8078 (2022)
- Xie, S., Feng, Z., Chen, Y., Sun, S., Ma, C., Song, M.: Deal: Difficulty-aware active learning for semantic segmentation. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (November 2020)
- Yan, Z., Hu, R., Yan, X., Chen, L., Van Kaick, O., Zhang, H., Huang, H.: Rpm-net: recurrent prediction of motion and parts from point cloud. ACM Transactions on Graphics (TOG) 38(6) (2019)
- 40. Zhan, X., Wang, Q., Huang, K.h., Xiong, H., Dou, D., Chan, A.B.: A comparative survey of deep active learning (2022). https://doi.org/10.48550/ARXIV.2203.13450, https://arxiv.org/abs/2203.13450
- Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5745–5753 (2019)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: International Conference on Learning Representations (2020)
- 43. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. arXiv preprint arXiv:2304.06718 (2023)