

Mean Shift Mask Transformer for Unseen Object Instance Segmentation

Yangxiao Lu, Yuqiao Chen, Nicholas Ruoizzi, Yu Xiang

Abstract—Segmenting unseen objects from images is a critical perception skill that a robot needs to acquire. In robot manipulation, it can facilitate a robot to grasp and manipulate unseen objects. Mean shift clustering is a widely used method for image segmentation tasks. However, the traditional mean shift clustering algorithm is not differentiable, making it difficult to integrate it into an end-to-end neural network training framework. In this work, we propose the Mean Shift Mask Transformer (MSMFormer), a new transformer architecture that simulates the von Mises-Fisher (vMF) mean shift clustering algorithm, allowing for the joint training and inference of both the feature extractor and the clustering. Its central component is a hypersphere attention mechanism, which updates object queries on a hypersphere. To illustrate the effectiveness of our method, we apply MSMFormer to unseen object instance segmentation. Our experiments show that MSMFormer achieves competitive performance compared to state-of-the-art methods for unseen object instance segmentation¹.

I. INTRODUCTION

The ability to segment unseen objects in cluttered scenes, a problem known as Unseen Object Instance Segmentation (UOIS), is critical in a variety of robotic applications. For example, UOIS allows robots to separate objects from each other to facilitate grasping of these objects [1], [2], [3], [4].

Unlike category-based object instance segmentation [5], [6], [7], [8], [9], [10], [11], UOIS requires the ability to generalize to unseen objects, that is, unseen object instances and unseen object categories. To improve generalizability, the state-of-the-art UOIS method [12] employs a feature learning and clustering strategy, where pixel-wise feature representations are learned using deep neural networks, and then a traditional clustering algorithm, such as mean shift clustering, is utilized to cluster pixel features for object segmentation. Although the feature learning paradigm achieves competitive performance, its potential capabilities are limited by its dependence on traditional clustering algorithms: the cluster assignment steps in these algorithms are not differentiable, which prevents the clustering and feature learning from being combined into a single end-to-end framework. The parameters in the clustering algorithm cannot be learned together with the feature representations.

In this work, we introduce the Mean Shift decoder (MS decoder), a novel differentiable mean shift clustering module based on the transformer architecture [8]. Our design of the MS decoder converts iterations of the von Mises-Fisher

Yangxiao Lu, Yuqiao Chen, Nicholas Ruoizzi and Yu Xiang are with the Department of Computer Science, the University of Texas at Dallas, Richardson, TX 75080, USA {yangxiao.lu, yuqiao.chen, nicholas.ruozzi, yu.xiang}@utdallas.edu

¹Project page, appendix, video, and code are available at: <https://irvlutd.github.io/MSMFormer>

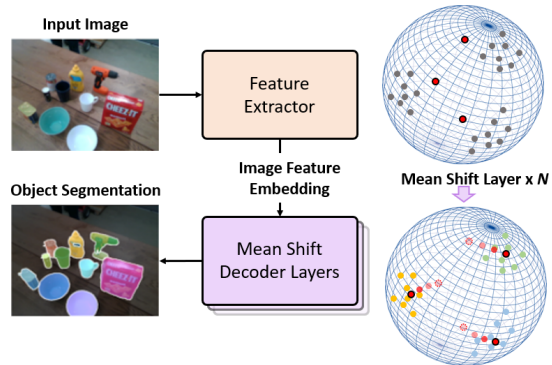


Fig. 1. Our framework consists of a feature extractor network to embed image pixels into a high dimensional hypersphere, and the mean shift decoder layers to compute cluster centers of the pixels for unseen object instance segmentation.

(vMF) mean shift clustering algorithm into multiple MS decoder layers. At each decoder layer, a set of object queries represent the cluster centers. Cross-attention is applied between the object queries and the pixel features to transform the object queries, which resembles the mean shift procedure to find the cluster centers. In addition, we propose to use L2 normalized object queries and pixel embeddings, and we introduce a hypersphere attention mechanism using the normalized queries, keys, and values.

In order to verify the effectiveness of our proposed MS decoder, we introduce a new network, Mean Shift Mask Transformer (MSMFormer), for UOIS by combining a feature learning backbone with the MS decoder for clustering. Since the MS decoder is differentiable, we can train this network end-to-end. Moreover, learning the weights in the MS decoder from data can further increase the capability of the model. Following previous work on UOIS, we train our network on a synthetic RGB-D dataset [13] and then test the trained network on the real-world Object Clutter Indoor Dataset (OCID) [14] and the Object Segmentation Database (OSD) [15]. Our method improves over the state-of-the-art methods on the OCID dataset and achieves comparable performance on the OSD dataset.

Our contributions are summarized as follows. 1) We introduce the mean shift decoder for differentiable mean shift clustering. 2) We introduce a hypersphere attention mechanism to improve training of transformers. 3) We propose MSMFormer for unseen object instance segmentation by combining feature learning and our mean shift decoder.

II. RELATED WORK

Semantic / Instance / Panoptic Segmentation. These three types of image segmentation tasks are based on object categories. Semantic segmentation labels each pixel in an

input image with a class label such as sky, person, table, etc. [16], [17]. Instance segmentation requires separating object instances among the same category such as multiple persons or cars. For example, Mask R-CNN [5] extends Faster R-CNN [18] by predicting an object mask for each detected object instance. Panoptic segmentation is proposed to unify both semantic and instance segmentation tasks, where every pixel needs to be labeled with a class label and an instance id [19], [20], [21]. It is tempting to solve panoptic segmentation tasks to obtain more information from the images. However, as discussed in Mask2Former [11], panoptic architectures do not guarantee good performance on instance segmentation tasks since they do not measure the abilities to rank predictions as instance segmentations. Mask2Former [11] is the first architecture that can achieve competitive performance on all of these segmentation tasks. Hence, we utilize their masked attention [11] in our method.

Unseen Object Instance Segmentation (UOIS). UOIS is object category agnostic. The goal is to segment unseen object instances from input images. The general idea for UOIS is to train with a large number of objects and hope the trained model can generalize to unseen objects. Since real images of many objects are difficult to collect, recent techniques utilize synthetic data for training [22], [13], [23]. In addition, depth images are widely used in UOIS to improve generalization. Recently, Zhang et al. [24] suggest exploiting test-time domain adaption to bridge the gap and boost the segmentation performance. Instead of depth input, Instance Stereo Transformer [25] generates object masks through the use of stereo image pairs. Unseen Clustering Network (UCN) [12] employs non-photorealistic RGB images and depth images to learn feature representations to segment unseen objects. It groups pixels via the von Mises-Fisher (vMF) mean shift clustering algorithm after learning the feature representation of pixels. To maximize image feature performance, our technique replaces the vMF mean shift clustering with a sequence of learnable transformer layers.

Differences between K-Means Mask Transformer [26] and Ours. K-means Mask Xformer [26] reformulates the cross-attention learning process as a k-means clustering procedure by proposing a new k-means cross-attention. However, the k-means cross-attention is not differentiable due to a hard assignment of pixel features to cluster centers, which makes training difficult. [26] needs to apply deep supervision after each decoder layer. In contrast, our proposed mean shift cross-attention using hypersphere attention is differentiable. In addition, the K-means Mask Xformer architecture is proposed for panoptic segmentation. Its performance on instance segmentation is not guaranteed. We were unable to successfully train a UOIS model using the K-means Mask Xformer architecture in our experiments.

vMF Mean Shift Clustering. The von Mises-Fisher (vMF) [27] distribution defines a distribution over unit vectors \mathbf{x} , $\|\mathbf{x}\| = 1$, on the d dimensional hypersphere. The vMF density function is defined as $P(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\kappa}) = C_d(\boldsymbol{\kappa}) \exp(\boldsymbol{\kappa} \boldsymbol{\mu}^T \mathbf{x})$, where $C_d(\boldsymbol{\kappa})$ is a normalization constant, $\boldsymbol{\mu}$ is a unit vector representing the mean direction of the distribution, and $\boldsymbol{\kappa}$ is

the concentration parameter that controls the concentration of the distribution around the mean direction.

The mean shift algorithm is an iterative procedure for finding the maxima of a distribution given a set of sampling points. In each mean shift iteration, the cluster centers will be updated with a shift towards the mean vector of its neighborhood, a set of sampling points within a certain window. The cluster centers will reach the positions with highest local densities once converged. Kobayashi and Otsu [28] proposed the use of the vMF distribution as a window function for sampling points on the hypersphere. In each mean shift iteration, each clustering center is updated by

$$\boldsymbol{\mu}_{t+1} = \frac{\sum_i^N \mathbf{x}_i \exp(\boldsymbol{\kappa} \boldsymbol{\mu}_t^T \mathbf{x}_i)}{\|\sum_i^N \mathbf{x}_i \exp(\boldsymbol{\kappa} \boldsymbol{\mu}_t^T \mathbf{x}_i)\|}, \quad (1)$$

where $\mathbf{x}_i, i = 1, \dots, N$ are sampling points from a dataset with size N . In this work, we adapt the concept of vMF mean shift clustering in our proposed mean shift mask transformer, which simulates the mean shift clustering procedure using a transformer decoder.

III. METHOD

In this section, we introduce the Mean Shift Mask Transformer (MSMFormer) as shown in Fig. 2. We utilize transformer decoders to convert object queries into object masks. By predicting N binary masks, a mask transformer architecture groups pixels into N segments. An inference pipeline in a mask transformer framework has three main steps. The first uses a backbone and an optional pixel decoder to generate feature maps F . The backbone extracts image features, which are then upsampled to higher-resolution features by the pixel decoder. Next, a transformer decoder that has been trained with a set of prediction losses produces mask embeddings C by updating object queries with the feature maps F . Lastly, binary mask predictions are generated from $\text{softmax}(C \times F)$. We use the pretrained Unseen Clustering Network (UCN) [12] to obtain robust RGB-D pixel embeddings since the backbone has been trained with a metric learning loss function [12], [29], [30].

A. Relationship between Cross-Attention and vMF Mean Shift Clustering

The key problem with the mask transformer-based segmentation frameworks is how to transform the object queries, i.e., some randomly initialized learnable embeddings, into informative mask embedding vectors. The design of our transformer decoder is inspired by the vMF mean shift clustering. To explain the development of our decoder layer, we begin by describing the similarity of the typical scaled dot-product attention [8] and one center update step in the vMF mean shift clustering, which motivates the hypersphere attention that we proposed. The cross-attention is subsequently modified into the attention version of one iteration in the clustering process.

Scaled Dot-Product Attention [8]. The scaled dot-product attention is defined as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right) \mathbf{V} = \frac{\exp\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)}{C_n} \mathbf{V}, \quad (2)$$

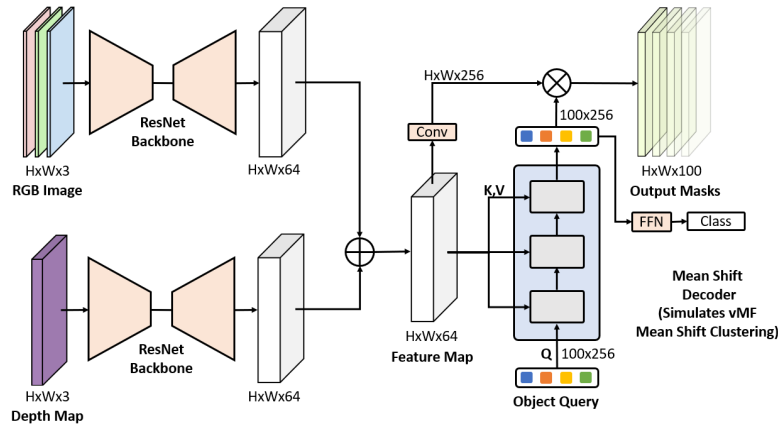


Fig. 2. The Mean Shift Mask Transformer consists of backbones and a mean shift decoder. The backbones can be any network that produces pixel embeddings. The depth map input and its backbone are optional depending on the task. A series of mean shift decoder layers are applied to transform $N = 100$ object queries. Then these transformed object queries group related pixels to generate masks and predict the classes of each mask. “FFN” indicates a Feed-Forward Network.

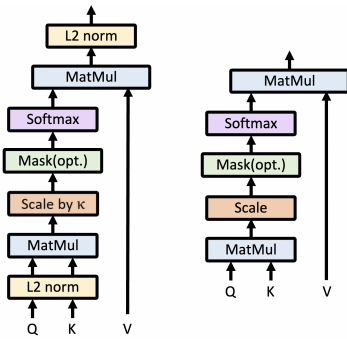


Fig. 3. (left) Our proposed hyper- sphere attention. (right) The standard scaled dot-product attention. Mask(Opt.) stands for optional attention masks.

where C_n stands for a normalization term. $\mathbf{Q} \in \mathbb{R}^{N \times D_k}$ denotes the N D_k -dim query vectors. $\mathbf{K} \in \mathbb{R}^{M \times D_k}$, $\mathbf{V} \in \mathbb{R}^{M \times D_v}$ are M D_k -dim key vectors and M D_v -dim value vectors, respectively. This function calculates the affinity between queries \mathbf{Q} and keys \mathbf{K} as the weights assigned to values \mathbf{V} . Next, it outputs the weighted sum of the values. The *softmax* function decomposes into an exponential function in the numerator and a normalization term in the denominator.

Comparing Eq. (1) and Eq. (2), both exponential functions produce outputs that represent the similarity of two items. These outputs are scaled by a factor (κ or $\frac{1}{\sqrt{d_k}}$) and then normalized. Ignoring the subtle differences between the above two equations, the weighted sums of values become the outputs. Since the attention process is similar to the update step of the vMF mean shift clustering, we propose the following attention variant.

Our **Hypersphere Attention** (HSAtten) is defined as:

$$\text{HSAtten}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = g(\text{softmax}(\kappa g(\mathbf{Q}) g(\mathbf{K})^T) \mathbf{V}), \quad (3)$$

where κ is a positive scalar hyperparameter. $g(\mathbf{x}) = \frac{\mathbf{x}}{\|\mathbf{x}\|}$ is used to ℓ_2 normalize these vectors into unit vectors. As shown in Fig. 3, we first compute the dot products of the normalized query vectors and key vectors as the cosine similarities between \mathbf{Q} and \mathbf{K} , which are then scaled by κ . Then the weights of values are obtained by softmax. Finally, we normalize the weighted sum of value vectors as the output. We call this attention variant *hypersphere attention* because these ℓ_2 normalized vectors are on a hypersphere.

Fig. 3 demonstrates the difference between our hypersphere attention and the scaled dot-product attention. We replace the scaled dot-product attention with our hypersphere attention to build our mean shift self-attention and cross-attention mechanisms in the next section.

B. Masked Mean Shift Cross-Attention

Simply using hypersphere attention is not enough to obtain good mask embeddings. Therefore, we encapsulate our attention in transformer-decoder layers. Further, inspired by masked attention [11], we add attention masks into the hypersphere attention in order to only allow a window of image feature embeddings affecting queries.

Cross-attention. The cross-attention modules are used to assemble pixel features that belong to the same object and update the corresponding object queries, defined as follows:

$$\hat{\mathbf{X}} = \mathbf{X} + \text{softmax}(\mathbf{Q} \times \mathbf{K}^T) \times \mathbf{V}, \quad (4)$$

where $\mathbf{X} \in \mathbb{R}^{N \times D}$ denotes N object queries with D channels, and $\hat{\mathbf{X}}$ is the updated queries using image features. The key-value pairs $\mathbf{K} \in \mathbb{R}^{HW \times D}$, $\mathbf{V} \in \mathbb{R}^{HW \times D}$ are produced via a linear transformation of the image features, where H and W are the height and width of the image, respectively.

Mean Shift Cross-attention. In our case, object queries can be viewed as cluster centers on a hypersphere. By substituting the scaled dot-product attention with our hypersphere attention, we obtain:

$$\hat{\mathbf{X}} = \mathbf{X} + g(\text{softmax}(\kappa \times g(\mathbf{Q}) \times g(\mathbf{K})^T) \times \mathbf{V}), \quad (5)$$

where $g(\mathbf{x}) = \frac{\mathbf{x}}{\|\mathbf{x}\|}$ is used to ℓ_2 normalize embeddings. The above equation employs the full feature maps to update the queries. However, global context in the cross-attention layers leads to the slow convergence of models with transformer decoders [31], [32]. Cheng et al. [11] hypothesize that local features are sufficient to update queries and self-attention can collect global context. Similarly, in mean shift clustering, the points contained within a window are sufficient to update a cluster center. Attention masks enable the cross-attention layers to focus on the local regions near each cluster center. Multiple mean shift self-attention layers can make queries

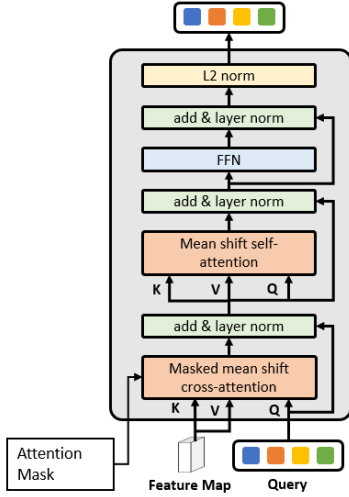


Fig. 4. Our mean shift decoder layer. We have L_2 norm after FFN.

unique, which saves us from merging cluster centers in mean shift clustering.

After introducing attention masks, our masked mean shift cross-attention works as follows:

$$\mathbf{X}_l = \mathbf{X}_{l-1} + g(\text{softmax}(\mathcal{M}_{l-1} + \kappa g(\mathbf{Q}_l)g(\mathbf{K}_l)^T) \mathbf{V}_l), \quad (6)$$

where \mathbf{Q}_l , \mathbf{K}_l and \mathbf{T}_l stands for the queries, keys and values in l -th transformer decoder layer. \mathbf{X}_{l-1} denotes the input of the l -th decoder layer. The attention mask \mathcal{M}_{l-1} [11] at feature location (x, y) is computed with the mask prediction \mathbf{M}_{l-1} via

$$\mathcal{M}_{l-1}(x, y) = \begin{cases} 0 & \text{if } \mathbf{M}_{l-1}(x, y) = 1 \\ -\infty & \text{otherwise} \end{cases}. \quad (7)$$

The mask prediction of the $(l-1)$ -th transformer decoder layer is first resized to the resolution of \mathbf{K}_l . It is then binarized to be $\mathbf{M}_{l-1} \in \{0, 1\}^{N \times HW}$ using a threshold of 0.5. \mathbf{M}_0 is the binary mask from the prediction of the raw queries which has not been processed by the transformer decoder.

C. Mean Shift Decoder Layer

As each iteration of vMF mean shift clustering outputs unit vectors. To obtain unit vectors in our framework, we impose ℓ_2 normalization on the output of the Feed-Forward Network (FFN) in each decoder layer. As shown in Fig. 4, our mean shift decoder layer (MS decoder layer) consists of three components: mean shift cross-attention, mean shift self-attention, and FFN with ℓ_2 normalization. In this way, we simulate one iteration of the clustering with one mean shift decoder layer. As shown in Fig. 2, the mean shift decoder consists of multiple MS decoder layers to replace the mean shift clustering process. Built on the backbones and MS decoder, we propose the Mean Shift Mask Transformer (MSMFormer) for segmentation tasks (Fig. 2). The backbones generate pixel embeddings to guide object queries in the MS decoder. After the queries find their destinations, the queries can be viewed as the final mask embeddings. They are then multiplied with the pixel embeddings to generate pixel similarities. The pixels with positive similarities form object masks [10], [11].

A. Implementation Details

Backbones. For RGB-D input, we choose two pretrained 34-layer, stride-8 ResNet (ResNet34-8s) from [12] as the backbone for RGB images and depth images, which are pretrained on the Tabletop Object Dataset [30]. With only RGB input, we choose a pretrained ResNet-50 backbone from Detectron2 [35] and the pixel decoder from Mask2Former [11] to compute the feature maps, which achieves better performance than the ResNet-34 backbone in [12].

Loss Functions. We adopt the same loss function as Mask2Former [11]. The mask loss consists of the binary cross-entropy loss and the dice loss [36]: $\mathcal{L}_{\text{mask}} = \lambda_{\text{ce}} \mathcal{L}_{\text{ce}} + \lambda_{\text{dice}} \mathcal{L}_{\text{dice}}$, where λ_{ce} and λ_{dice} are both set as 5.0. The final loss is the weighted sum of mask loss and classification loss: $\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{mask}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}}$, where λ_{cls} is 2.0 for predictions that can match with a ground truth object and 0.1 for predictions that cannot match any ground truth object. We use the AdamW optimizer [37] with batch size 4 and learning rate $1e-4$. We train the model for one epoch on the Tabletop Object Dataset from [13] that contains 280,000 images.

Post-processing. We adopt the post-processing from Mask2Former [11] for instance segmentation. Each mask prediction has a class confidence score. Since the network always outputs 100 mask predictions, we only consider those with a score of 0.7 or higher as reliable masks.

Pixel Confidence Map. We preserve positive pixel similarities that belong to object masks via the ReLU function. Then we normalize the similarities by dividing them by their maximum value, scaling them from 0 to 1. This process yields pixel confidence heatmaps, which serves as a tool for identifying contour uncertainty [38].

Second Stage as Zoom-in Refinement. The segmentation accuracy can be improved by using a two-stage clustering process [12]. Each segmentation mask from the first stage is fed into another MSMFormer for object segmentation again. In this way, nearby objects or stacked objects can be separated using the second-stage segmentation network.

Inference Time. Given a 480×640 RGB-D image, the running time of MSMFormer is on average 0.278s for the 1st stage and 0.072s per ROI in the 2nd stage on an NVIDIA A100 GPU. In contrast, the reported running time of UCN [12] is 0.2s for the 1st stage and 0.05s per object in the 2nd stage on a TITAN XP GPU.

B. Evaluation Results

Datasets. We evaluate the performance on the Object Cluster Indoor Dataset (OCID) [14] and the Object Segmentation Database (OSD) [15] for UOIS in tabletop scenes. The OCID dataset contains 2,390 RGB-D images, with up to 20 objects and an average of 7.5 objects per image. The OSD dataset contains 111 RGB-D images, with up to 15 objects and an average of 3.3 objects per image.

Evaluation Metrics. We evaluate the object segmentation performance using precision, recall and F-measure according to [13], [12]. For these three metrics, we first compute the values between all pairs of predicted objects and ground truth objects. Then we use the Hungarian method with pairwise

TABLE I

UNSEEN OBJECT INSTANCE SEGMENTATION RESULTS ON OCID AND OSD. +: THE RESULT USING A SECOND-STAGE NETWORK, I.E., AFTER ZOOM-IN REFINEMENT. *: MODELS TRAINED WITH THE UOAI-SIM DATASET IN [33]. #: MODELS WITH AMODAL PERCEPTION.

Method	Input	OCID (2390 images)							OSD (111 images)						
		Overlap			Boundary				Overlap			Boundary			
		P	R	F	P	R	F	%75	P	R	F	P	R	F	%75
MRCNN [5]	RGB	77.6	67.0	67.2	65.5	53.9	54.6	55.8	64.2	61.3	62.5	50.2	40.2	44.0	31.9
UCN [12]	RGB	54.8	76.0	59.4	34.5	45.0	36.5	48.0	57.2	73.8	63.3	34.7	50.0	39.1	52.5
UCN+ [12]	RGB	59.1	74.0	61.1	40.8	55.0	43.8	58.2	59.1	71.7	63.8	34.3	53.3	39.5	52.6
Mask2Former [11]	RGB	67.2	73.1	67.1	55.9	58.1	54.5	54.3	60.6	60.2	59.5	48.2	41.7	43.3	32.4
UOAI-Sim-Net [33]*	RGB #	66.5	83.1	67.9	62.1	70.2	62.3	73.1	84.2	83.7	83.8	72.2	72.8	72.1	76.7
MSMFormer* (Ours)	RGB	70.2	84.4	70.5	64.5	74.9	64.9	75.3	59.3	82.0	67.9	42.9	72.0	52.4	72.4
MSMFormer (Ours)	RGB	72.9	68.3	67.7	60.5	56.3	55.8	52.9	63.4	64.7	63.6	48.6	47.4	47.0	40.2
MSMFormer+ (Ours)	RGB	73.9	67.1	66.3	64.6	52.9	54.8	52.8	63.9	63.7	62.7	51.6	45.3	47.0	41.1
MRCNN [5]	Depth	85.3	85.6	84.7	83.2	76.6	78.8	72.7	77.8	85.1	80.6	52.5	57.9	54.6	77.6
UOIS-Net-2D [13]	Depth	88.3	78.9	81.7	82.0	65.9	71.4	69.1	80.7	80.5	79.9	66.0	67.1	65.6	71.9
UOIS-Net-3D [34]	Depth	86.5	86.6	86.4	80.0	73.4	76.2	77.2	85.7	82.5	83.3	75.7	68.9	71.2	73.8
UOAI-Sim-Net [33]*	Depth #	89.9	90.9	89.8	86.7	84.1	84.7	87.1	84.9	86.4	85.5	68.2	66.2	66.9	80.8
UOAI-Sim-Net [33]*	RGBD #	70.7	86.7	71.9	68.2	78.5	68.8	78.7	85.3	85.4	85.2	72.7	74.3	73.1	79.1
UCN [12]	RGBD	86.0	92.3	88.5	80.4	78.3	78.8	82.2	84.3	88.3	86.2	67.5	67.5	67.1	79.3
UCN+ [12]	RGBD	91.6	92.5	91.6	86.5	87.1	86.1	89.3	87.4	87.4	87.4	69.1	70.8	69.4	83.2
Mask2Former [11]	RGBD	78.6	82.8	79.5	69.3	76.2	71.1	69.3	75.6	79.2	77.3	54.1	64.0	58.0	65.2
MSMFormer* (Ours)	RGBD	78.2	83.8	76.6	69.1	78.4	69.6	75.3	80.6	86.9	83.1	65.9	72.9	68.7	80.2
MSMFormer (Ours)	RGBD	88.4	90.2	88.5	84.7	83.1	83.0	80.3	79.5	86.4	82.8	53.5	71.0	60.6	79.4
MSMFormer+ (Ours)	RGBD	92.5	91.0	91.5	89.4	85.9	87.3	86.0	87.1	86.1	86.4	69.0	68.6	68.4	80.4

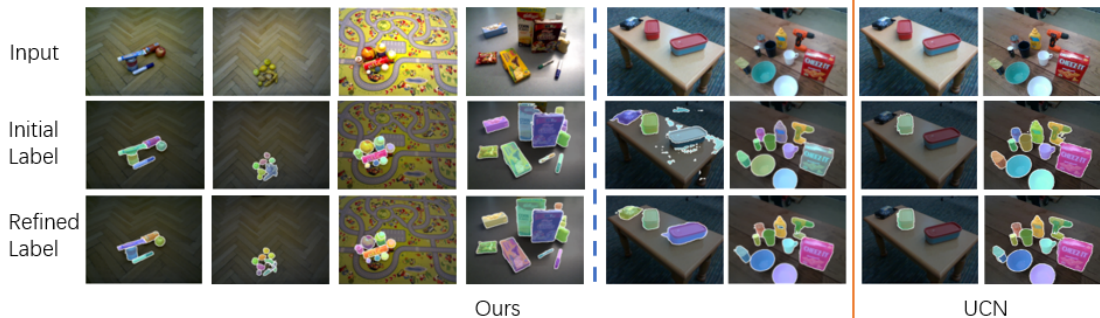


Fig. 5. Examples of the two-stage predictions. The first six columns show our method’s predictions, while the last two columns display UCN’s predictions [12]. Given the same inputs, in the fifth column, our method segments all the objects whereas UCN only finds one of them. In the last image, UCN over-segments the bottom-left can during zoom-in refinement while ours does not.

F-measure to match predictions with ground truth. Based on this matching, the final precision, recall and F-measure are obtained by $P = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_i |c_i|}$, $R = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_j |g_j|}$, $F = \frac{2PR}{P+R}$, where c_i denotes the segmentation of predicted object i , $g(c_i)$ denotes the segmentation of the matched ground truth object of c_i , and g_j is the segmentation for ground truth object j . Overlap P/R/F denote the aforementioned three metrics since the true positives can be viewed as the overlap segmentation of the whole object. In addition, Boundary P/R/F are used to evaluate how sharp the predicted boundary matches against the ground truth boundary, where the true positives are counted by the pixel overlap of the two boundaries. Moreover, Overlap F-measure 75% is the percentage of segmented objects with Overlap F-measure $\geq 75\%$ [39].

Comparison to other Methods. In Table I, we compare our model with existing methods on OCID and OSD. For OCID, our RGB-D model achieves a new state-of-the-art of 87.3 Boundary F-measure. The first-stage model already reaches 83.0 Boundary F-measure which demonstrates the advantage of MSMFormer on accurately segmenting objects. With only RGB images, our model also outperforms other methods on the OCID dataset.

For OSD, the UOAI-Sim-Net [33] achieves the best performance by using a new synthetic dataset UOAI-Sim for

training, which employs Hierarchical Occlusion Modeling trained with the amodal and visible masks. We found that training MSMFormer only on the visible masks of this new synthetic dataset improves our results on OCID but not on OSD. The main reason is that OSD contains images with heavy occlusions. The UOAI-Sim dataset provides more training images in these scenarios.

Synthetic training datasets impose limitations on our model’s performance. The UOAI-Sim dataset with superior quality RGB images in comparison to the TableTop dataset, contributes to the enhancement of MSMFormer’s performance. Our model demonstrates robust scalability when exposed to improved datasets.

Effect of Second-stage Model. For RGB-D input, the results of the model with or without the second-stage refinement are presented in the final two rows of Table I. The second-stage model improves most of the metrics in both data sets. For RGB input, the first-stage model makes incorrect predictions on backgrounds such as walls and colorful blankets. The second stage model does not remove these false alarms. Therefore, it cannot enhance the overall segmentation performance. Fig. 5 shows some segmentation results before and after refinement using RGB-D input. In these examples, the boundaries are refined and more consistent, and some merged objects are separated correctly.

TABLE II

TWO KINDS OF DECODER LAYERS RESULTS. *Masked* DENOTES THE TRANSFORMER DECODER LAYER WITH MASKED ATTENTION [11], WHEREAS *MS* IS OUR MS DECODER LAYER. +: THE TWO-STAGE PERFORMANCE.

Method	OCID (2390 images)							OSD (111 images)							
	Overlap			Boundary				%75	Overlap			Boundary			%75
	P	R	F	P	R	F	P		R	F	P	R	F		
Masked	88.4	90.0	88.2	85.4	82.4	83.0	78.8	72.4	80.5	76.2	45.6	63.4	52.5	65.3	
Masked+	90.3	89.2	89.4	87.7	89.2	85.0	81.5	80.1	79.6	79.7	64.1	59.8	61.4	66.5	
MS (Ours)	88.4	90.2	88.5	84.7	83.1	83.0	80.3	79.5	86.4	82.8	53.5	71.0	60.6	79.4	
MS+ (Ours)	92.5	91.0	91.5	89.4	85.9	87.3	86.0	87.1	86.1	86.4	69.0	68.6	68.4	80.4	

TABLE III

ABLATION STUDY ON DIFFERENT NUMBERS OF MS DECODER LAYERS.

Stage	#layers	OCID(2390 images)						
		Overlap			Boundary			
		P	R	F	P	R	F	%75
1st	4	74.4	83.2	78.4	45.8	67.0	54.0	72.9
	6	88.4	90.2	88.5	84.7	83.1	83.0	80.3
	8	88.4	89.2	88.1	84.8	82.8	82.9	80.0
2nd	6	91.8	89.6	90.5	88.3	83.8	85.6	83.0
	8	92.5	91.0	91.5	89.4	85.9	87.3	86.0
	10	91.6	89.7	90.3	88.3	83.7	85.5	83.2

C. Ablation Studies

Our Mean Shift Decoder vs. Standard Scaled Dot-product Attention Decoder. We train Mask2Former [11] using RGB-D images with ResNet50 backbones on the Tabletop Object Dataset. As shown in Table I, our MSMFormer significantly outperforms Mask2Former on both OCID and OSD. The main reason is that Mask2Former overfits the synthetic training set and does not generalize well to real images. It cannot benefit from the pretrained backbones trained with the metric learning loss from UCN [12]. Therefore, we use the same pretrained network from UCN to compare our MS decoder layer and the Masked-attention Transformer decoder [11]. They are trained for one epoch on the Tabletop Object dataset. In Table II, our MS decoder layer outperforms the Masked-attention decoder in most metrics on both datasets.

Number of Mean Shift Decoder Layers. The first-stage model processes RGB-D images with size 480×640 , whereas the second stage processes ROI images with size 224×224 . The input image size has an impact on the number of MS decoder layers required, which can be viewed as the number of iterations in vMF mean shift clustering. Table III shows that with more than 6 decoder layers, the network performance is stable. The second-stage segmentation is based on the output of the best first-stage model with 6 MS decoder layers. In Table III, even though the results of second-stage networks with various numbers of layers are close, the 8-layer model has the best performance.

Failure Cases and Limitations. As shown in Fig. 6, most failures are under-segmentation and over-segmentation. In RGB-D images, as some objects are adjacent or stacked together, their embeddings are still close in feature space. These objects are hard to separate even with the second-stage refinement. In Fig. 6, for example, the first image has three boxes next to each other, which makes the network struggle to separate them. However, the pixel confidence heatmap can still provide some uncertainty measurement of the under-segmented masks. When an object can be covered by other objects, over-segmentation occurs since the model is misled by the view. For instance, in the last image, the keyboard under other objects is over-segmented. MSMFormer is cur-

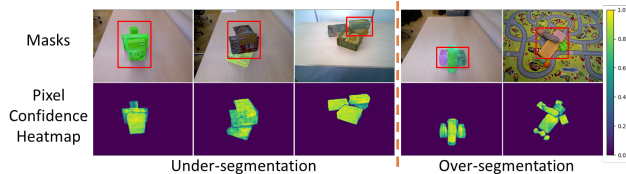


Fig. 6. Examples of failure cases and their pixel confidence heatmaps. The contour uncertainty can be found in some under-segmentation cases.

rently limited by these challenging scenarios.

These issues may be resolved in future work by generating more synthetic training images of these difficult scenarios.

Testing with a Real Robot. We have tested the performance of MSMFormer on a Fetch mobile manipulator. RGB-D images from the Fetch head camera are fed into the network to segment objects in tabletop scenes. Once objects are segmented, we have enabled the robot to pick up objects by applying the Contact-GraspNet [2] for grasp planning and MoveIt [40] for motion planning of grasping. Please see the appendix for additional results and the video for the real-world robot grasping demonstration.

V. CONCLUSION AND FUTURE WORK

In this work, we have proposed a novel transformer attention variant, called *hypersphere attention*, motivated by the von Mises-Fisher mean shift clustering algorithm. With this attention, we designed the masked mean shift cross-attention and mean shift self-attention. Next, we introduced the mean shift decoder for differentiable mean shift clustering to build the Mean Shift Mask Transformer (MSMFormer) for object segmentation tasks. In our model, the feature representation and the clustering for object segmentation can be trained jointly in an end-to-end fashion. In this way, our MSMFormer obtains very competitive results on the Object Clutter In-door Dataset (OCID) and the Object Segmentation Database (OSD) for Unseen Object Instance Segmentation (UOIS). We also demonstrated that our method is effective for UOIS on real-world images.

For future work, we will explore several ways to further improve the unseen object segmentation performance. One idea is to utilize photo-realistic synthetic images for training. Another idea is to leverage a small set of real-world images to fine-tune the segmentation network for domain adaptation. In addition, we plan to study UOIS beyond tabletop scenes such as bin-picking settings and objects inside cabinets, which will facilitate robot manipulation in these scenarios.

ACKNOWLEDGMENT

This work was supported in part by the DARPA Perceptually-enabled Task Guidance (PTG) Program under contract number HR00112220005.

REFERENCES

- [1] A. Mousavian, C. Eppner, and D. Fox, “6-dof graspnet: Variational grasp generation for object manipulation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2901–2910.
- [2] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 438–13 444.
- [3] L. Wang, Y. Xiang, W. Yang, A. Mousavian, and D. Fox, “Goal-auxiliary actor-critic for 6d robotic grasping with point clouds,” in *Conference on Robot Learning*. PMLR, 2022, pp. 70–80.
- [4] A. Goyal, A. Mousavian, C. Paxton, Y.-W. Chao, B. Okorn, J. Deng, and D. Fox, “Ifor: Iterative flow minimization for robotic object rearrangement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14 787–14 797.
- [5] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [6] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, “End-to-end object detection with transformers,” *CoRR*, vol. abs/2005.12872, 2020. [Online]. Available: <https://arxiv.org/abs/2005.12872>
- [7] Z. Tian, C. Shen, and H. Chen, “Conditional convolutions for instance segmentation,” in *European conference on computer vision*. Springer, 2020, pp. 282–298.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [9] X. Jia, B. De Brabandere, T. Tuytelaars, and L. V. Gool, “Dynamic filter networks,” *Advances in neural information processing systems*, vol. 29, 2016.
- [10] B. Cheng, A. Schwing, and A. Kirillov, “Per-pixel classification is not all you need for semantic segmentation,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 17 864–17 875, 2021.
- [11] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, “Masked-attention mask transformer for universal image segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1290–1299.
- [12] Y. Xiang, C. Xie, A. Mousavian, and D. Fox, “Learning rgb-d feature embeddings for unseen object instance segmentation,” in *Conference on Robot Learning (CoRL)*, 2020.
- [13] 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*. PMLR, 2020, pp. 1369–1378.
- [14] M. Suchi, T. Patten, D. Fischinger, and M. Vincze, “Easylab: A semi-automatic pixel-wise object annotation tool for creating robotic rgb-d datasets,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6678–6684.
- [15] A. Richtsfeld, T. Mörwald, J. Prankl, M. Zillich, and M. Vincze, “Segmentation of unknown objects in indoor environments,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 4791–4796.
- [16] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [17] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
- [18] S. Ren, K. He, R. B. Girshick, and J. Sun, “Faster R-CNN: towards real-time object detection with region proposal networks,” *CoRR*, vol. abs/1506.01497, 2015. [Online]. Available: <http://arxiv.org/abs/1506.01497>
- [19] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, “Panoptic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9404–9413.
- [20] B. Cheng, M. D. Collins, Y. Zhu, T. Liu, T. S. Huang, H. Adam, and L.-C. Chen, “Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 12 475–12 485.
- [21] Y. Li, H. Zhao, X. Qi, Y. Chen, L. Qi, L. Wang, Z. Li, J. Sun, and J. Jia, “Fully convolutional networks for panoptic segmentation with point-based supervision,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [22] L. Shao, Y. Tian, and J. Bohg, “Clusternet: 3d instance segmentation in rgb-d images,” *arXiv preprint arXiv:1807.08894*, 2018.
- [23] 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 *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 7283–7290.
- [24] L. Zhang, S. Zhang, X. Yang, and Z. Liu, “Unseen object instance segmentation with fully test-time rgb-d embeddings adaptation,” *arXiv preprint arXiv:2204.09847*, 2022.
- [25] M. Durner, W. Boerdijk, M. Sundermeyer, W. Friedl, Z.-C. Márton, and R. Triebel, “Unknown object segmentation from stereo images,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4823–4830.
- [26] Q. Yu, H. Wang, S. Qiao, M. Collins, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, “k-means mask transformer,” in *European Conference on Computer Vision*. Springer, 2022, pp. 288–307.
- [27] K. V. Mardia, P. E. Jupp, and K. Mardia, *Directional statistics*. Wiley Online Library, 2000, vol. 2.
- [28] T. Kobayashi and N. Otsu, “Von mises-fisher mean shift for clustering on a hypersphere,” in *2010 20th International Conference on Pattern Recognition*, 2010, pp. 2130–2133.
- [29] B. De Brabandere, D. Neven, and L. Van Gool, “Semantic instance segmentation with a discriminative loss function,” *arXiv preprint arXiv:1708.02551*, 2017.
- [30] C. Xie, Y. Xiang, Z. Harchaoui, and D. Fox, “Object discovery in videos as foreground motion clustering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9994–10 003.
- [31] P. Gao, M. Zheng, X. Wang, J. Dai, and H. Li, “Fast convergence of detr with spatially modulated co-attention,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3621–3630.
- [32] Z. Sun, S. Cao, Y. Yang, and K. M. Kitani, “Rethinking transformer-based set prediction for object detection,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 3611–3620.
- [33] S. Back, J. Lee, T. Kim, S. Noh, R. Kang, S. Bak, and K. Lee, “Unseen object amodal instance segmentation via hierarchical occlusion modeling,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 5085–5092.
- [34] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, “Unseen object instance segmentation for robotic environments,” *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1343–1359, 2021.
- [35] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, “Detectron2,” <https://github.com/facebookresearch/detectron2>, 2019.
- [36] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in *2016 fourth international conference on 3D vision (3DV)*. IEEE, 2016, pp. 565–571.
- [37] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” *arXiv preprint arXiv:1711.05101*, 2017.
- [38] C. Xie, A. Mousavian, Y. Xiang, and D. Fox, “Rice: Refining instance masks in cluttered environments with graph neural networks,” in *Conference on Robot Learning*. PMLR, 2022, pp. 1655–1665.
- [39] P. Ochs, J. Malik, and T. Brox, “Segmentation of moving objects by long term video analysis,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 6, pp. 1187–1200, 2013.
- [40] S. Chitta, I. Sukan, and S. Cousins, “Moveit![ros topics],” *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 18–19, 2012.