

Amodal Segmentation Based on Visible Region Segmentation and Shape Prior

Yuting Xiao,^{1,2} Yanyu Xu,⁴ Ziming Zhong,^{1,2} Weixin Luo,^{1,2} Jiawei Li,³ Shenghua Gao^{1,2}

¹ ShanghaiTech University

² Shanghai Engineering Research Center of Intelligent Vision and Imaging

³ Alibaba Group

⁴ Institute of High Performance Computing, A*STAR

{xiaoyt, zhongzm, luowx, gaoshh}@shanghaitech.edu.cn, xu_yanyu@ihpc.a-star.edu.sg, mingong.ljw@alibaba-inc.com

Abstract

Almost all existing amodal segmentation methods make the inferences of occluded regions by using features corresponding to the whole image. This is against the human's amodal perception, where human uses the visible part and the shape prior knowledge of the target to infer the occluded region. To mimic the behavior of the human and solve the ambiguity in the learning, we propose a framework, it firstly estimates a coarse visible mask and a coarse amodal mask. Then based on the coarse prediction, our model infers the amodal mask by concentrating on the visible region and utilizing the shape prior in the memory. In this way, features corresponding to background and occlusion can be suppressed for amodal mask estimation. Consequently, the amodal mask would not be affected by the occlusion when given the same visible regions. The leverage of shape prior makes the amodal mask estimation more robust and reasonable. Our proposed model is evaluated on three datasets. Experiments show that our proposed model outperforms existing state-of-the-art methods. The visualization of shape prior indicates that the category-specific feature in the codebook has certain interpretability. The code is available at <https://github.com/YutingXiao/Amodal-Segmentation-Based-on-Visible-Region-Segmentation-and-Shape-Prior>.

Introduction

Amodal segmentation aims to infer the amodal mask, including both the visible region and the possible invisible region of the target object. Different from semantic segmentation or traditional instance segmentation, amodal segmentation is designed to exploit the amodal perception capability (Zhu et al. 2017), where human could infer and perceive the whole semantic concept of the target object mainly according to the partially visible region of the target object.

Simulating the amodal perception is quite challenging. Lots of efforts have been done for the amodal segmentation, which could be roughly divided into two groups. Methods in the first group directly estimate both the visible and the amodal regions from the images (Qi et al. 2019; Follmann et al. 2019; Zhu et al. 2017), while methods in the second group use inferred depth order information to help the amodal mask prediction (Zhang et al. 2019). However,

all of them learn the mapping relationship from the feature corresponding to the whole view to the amodal mask. This processing brings explicitly ambiguity that the same image appearances of occlusion may require different predictions.

Motivated by the behavior of human leveraging only visible regions and memorizing the category-specific shape prior for amodal segmentation, we propose a solution that estimates the visible region of the target object and leverages the visible region and shape prior for amodal segmentation. Without the shape prior of an object, the amodal mask inferred by human might be in an arbitrary shape. Similarly, the estimated amodal mask might also suffer from the lack of shape prior, resulting in the arbitrary edges. Thus, our model leads to more robust results. As shown in Fig. 1, although the target (green vegetable) is occluded by different objects, our method could perceive the almost same occluded regions, like human ignoring the different occlusion contexts such as banana, apples, or cabbage. However, the existing baseline estimates the different occluded regions.

In particular, our proposed method consists of a coarse mask segmentation module, a visible mask segmentation module, and an amodal mask segmentation module. In the coarse mask segmentation module, we utilize the backbone of Mask R-CNN (He et al. 2017) with an amodal mask head and a visible mask head to predict the coarse amodal and visible mask respectively. In the visible mask segmentation module, we propose to leverage the amodal mask to refine the visible mask and a reclassification regularizer to alleviate the misleading effect of occlusion in classification. Specifically, we use the coarse amodal mask as the attention multiplying with the feature of region-of-interest for more accurate visible mask estimation. The coarse amodal mask alleviates the effect of the background and contains more information than the coarse visible mask, which provides a cue for the visible mask. Further, for the reclassification regularizer, we propose to apply the feature of the visible region instead of the whole view to for classification, which alleviates the influence of the occlusion and background. In the amodal mask segmentation module, we propose to use the feature of the visible region and the shape prior to refine the amodal mask. Unlike the diversity of the visible mask, the amodal mask has the inherent category-specific shape prior. To encode the shape prior, we design an auto-encoder for amodal mask encoding, and a codebook is used as the memory ob-




Figure 1: Performance comparison between ours and Mask R-CNN. We overlay different objects on top of the target green vegetables with Adobe Photoshop. Our method could estimate the almost same invisible regions, given different occlusions. For the Mask R-CNN using features corresponding to both occlusions and visible regions, the results are different.

tained from the K-Means of amodal ground-truth mask embeddings. We utilize the shape prior for amodal mask prediction in two aspects: Firstly, we utilize the shape prior to refine the amodal mask. Secondly, in the inference, after the final amodal mask prediction, we use the shape prior to post-process the score of proposed boxes to filter out the proposals with a low-quality amodal mask.

As shown in Fig. 1, although the target (green vegetable) is occluded by different objects, our method could perceive the almost same occluded regions with different occlusion contexts such as banana, apples, or cabbage. However, the existing baseline estimates the different occluded regions with features corresponding to different occlusions.

The contributions of this paper could be summarized as followings: (1) inspired by the behavior of human’s amodal perception, we propose a novel amodal segmentation model; (2) a cross-task attention based refinement strategy is proposed, where the amodal mask and the visible mask is used as attention to refine each other for improving the performance of amodal segmentation; (3) this is the first work that proposes to utilize the shape prior knowledge in the amodal segmentation, and two ways to use shape prior are discussed; (4) experimental results on three public datasets show our method outperforms the existing state-of-the-art methods.

Related Work

Visual Occlusion Learning

Occlusion is inevitable in a large amount of visual task regardless of detection (Huang et al. 2020; Chu et al. 2020), segmentation (Fu et al. 2019; Huang et al. 2019) or inpainting (Ren et al. 2019; Yu et al. 2019; Kar et al. 2015) tasks. Some researches about occlusion propose interesting and innovative methods or framework to solve the occlusion problem. In (Yang et al. 2019), Yang *et al.* propose a strategy to see the information behind the occlusion by moving the position of the camera. The BANet(Chen et al. 2020) computes the similarity of the pixels in the boundary region of instance to recognize whether these pixels belong to the occlusion or not. Some other works tend to explicitly remove the influence of occlusion, the conditional random fields (CRFs) is applied in (Winn and Shotton 2006) to represent the occlusion probability of object parts. The recent work (Huang

et al. 2020; Chu et al. 2020) deals with the occlusion problem in human crowd detection which is another task with heavy occlusion. The visible region is used to guide the full region of human in (Huang et al. 2020) which indicates that a pair of boxes is generated by a single proposal. Besides, in (Chu et al. 2020), Xuan *et al.* applied a similar measure which takes advantage of a proposal corresponding to multiple predictions to avoid the ambiguous regression. In 3D area, many methods (Zhang, Huang, and Wang 2020), (Ramamonjisoa, Du, and Lepetit 2020; Cheng et al. 2019) about solving the negative influence of occlusion are utilized for 3D keypoint, mesh and depth estimation.

Amodal Segmentation

A large number of researches has been accomplished on instance segmentation (Bai and Urtasun 2017; Ren and Zemel 2017; Xu et al. 2019; Neven et al. 2019; Chen et al. 2019), based on classic detection framework such as Faster-RCNN (Ren et al. 2015) or YOLO (Redmon et al. 2016). Mask R-CNN (He et al. 2017), one of the most representative two-stage approach, uses a mask head added in the Faster-RCNN to process the aggregated feature sampled by the ROIAlign module. In order to utilize the multiple scale information, the FPN (Lin et al. 2017) is proposed to detect instances on different scales. Further, the Path aggregation network (Liu et al. 2018) is proposed to boost information flow in feature hierarchy. Besides, some other works localize the position of instances by the center point of each instance (Xie et al. 2020) instead of the boundary box or iteratively deform an initial contour to the boundary box (Peng et al. 2020).

As a newly developing direction of instance segmentation, The earliest work about amodal segmentation is proposed by (Li and Malik 2016), which utilizing generated data by overlapping instances on other instances to train and test their method. They extend the boxes of each instance and refine the heatmap. And in (Zhu et al. 2017), the Sharp-Mask (Pinheiro, Collobert, and Dollár 2015) which predicts the object mask from coarse to fine is provided as the baseline model. Recently, several amodal segmentation datasets are released to help for the amodal segmentation research. The ORCNN (Follmann et al. 2019) uses a visible mask head and an amodal mask head to directly predict the visible mask and amodal mask respectively, obtaining the occlusion

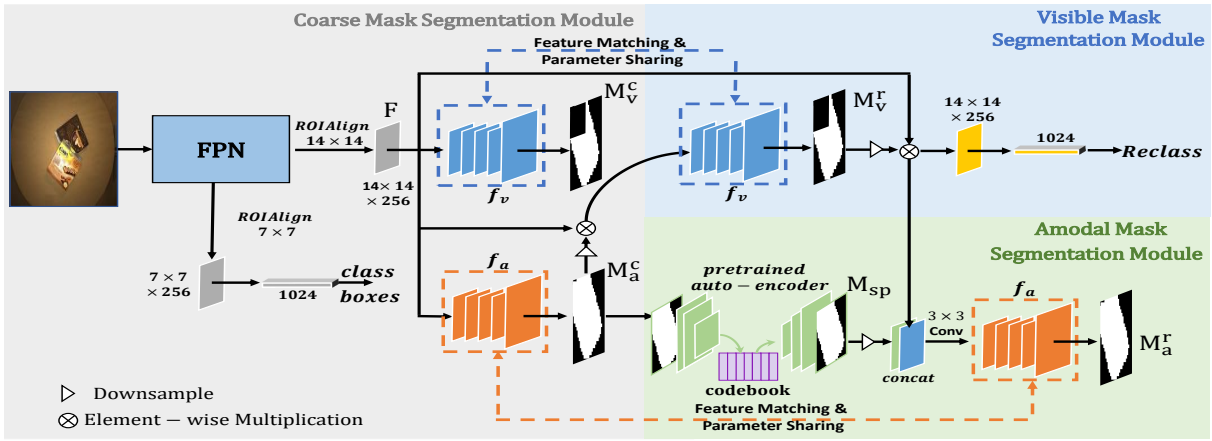


Figure 2: The overview of our approach.

mask by subtracts the visible mask from the amodal mask. The SLN (Zhang et al. 2019) modeling a depth order representation to infer the amodal mask. The method proposed in (Qi et al. 2019) utilizes an occlusion classifier to recognize whether an instance is occluded and ensemble the feature from box and class head by multi-level coding.

Method

Problem Formulation

Given an image I , amodal segmentation aims to estimate the amodal mask M_a as well as the visible mask M_v for a region-of-interest (ROI). The visible mask M_v could be estimated directly from the image. The amodal mask M_a consists of both the visible part and the invisible part. The most challenging part of amodal segmentation is to estimate the invisible region based on the visible region and without being affected by the occlusions.

The human amodal perception means the ability to infer the global instance according to the partial observation. When inferring the invisible region, human utilizes the feature of the visible region and the shape prior of the object. Inspired by this, we formulate the amodal segmentation as learning a nonlinear mapping function that maps the ROI feature F of an image I to the visible mask M_v and the amodal mask M_a with the regularization of shape prior. As shown in Fig. 2, our proposed method consists of a coarse mask segmentation module, a visible mask segmentation module, and an amodal mask segmentation module.

The Coarse Mask Segmentation Module

The coarse mask segmentation module aims to extract the target visual features and predict the coarse amodal mask M_a^c and the coarse visible mask M_v^c . Following the existing works (Zhang et al. 2019; Follmann et al. 2019; Qi et al. 2019), we also employ a ResNet50 (He et al. 2016) based FPN (Lin et al. 2017) as the backbone to extract the visual feature of the region-of-interest F containing the target object. Specifically, in this module, both the visible mask head f_v and the amodal mask head f_a take the feature F as input.

The amodal mask head and the visible mask head have the same network structure that consists of 4 convolution layers and 1 deconvolution layer with different parameters.

There are four loss terms in this module, including a coarse amodal mask loss $\mathcal{L}_{BCE}(M_a^c, M_a^g)$, a coarse visible mask loss $\mathcal{L}_{BCE}(M_v^c, M_v^g)$, a classification loss \mathcal{L}_{cls} and an object bounding box regression loss \mathcal{L}_{reg} . The M_a^g and M_v^g are the amodal and visible ground-truth mask. Both of the \mathcal{L}_{cls} and \mathcal{L}_{reg} are the same as the loss function of the class and box head in Mask R-CNN (He et al. 2017). The $\mathcal{L}_{BCE}(\cdot, \cdot)$ is the binary cross-entropy loss.

The Visible Mask Segmentation Module

The visible segmentation module aims to further refine the visible mask, via the amodal mask and the reclassification regularizer. Because the amodal mask contains the visible region, we use it as the attention to multiple with the ROI feature F for the visible mask refinement. This operation enhances the capability to distinguish occlusion and target instance of the visible mask head. Further, it also alleviates the effect of background features for visible mask refinement. The loss term of visible mask refinement is

$$\mathcal{L}_v^r = \frac{1}{N} \sum_i^N \mathcal{L}_{BCE}(f_v(F_i \cdot M_{a,i}^c), M_{v,i}^g), \quad (1)$$

where N is the number of predicted instances. We denote the refined visible mask of i^{th} instance as $M_{v,i}^r$, which is the output of the visible mask head f_v in this term.

The reclassification regularizer aims to classify each instance by processing the feature of the visible region, which avoids the misleading effect of the feature corresponding to the occlusion and background. The input of the reclassification regularizer is the feature of refined visible region regarded as $F \cdot M_v^r$. The reclassification regularizer f_{rc} composes of two fully connected layers. In the inference, we obtain the class score of each instance by multiplying the score of the class head and reclassification regularizer. The loss term of the reclassification regularizer is

$$\mathcal{L}_{rc} = \frac{\lambda_{rc}}{N} \sum_i^N \mathcal{L}_{CE}(f_{rc}(F_i \cdot M_{v,i}^r), y_i), \quad (2)$$

where y_i is the class label of i^{th} instance. $\mathcal{L}_{CE}(\cdot, \cdot)$ is the cross-entropy loss. The hyper-parameter $\lambda_{rc} = 0.25$.

Feature matching Feature matching is commonly used in network model acceleration and compression (Ba and Caruana 2013; Li, Jin, and Yan 2017), where a compact student model mimics the feature maps extracted from a large teacher model to improve its accuracy. In this work, we also employ the feature matching strategy for the visible mask head to reduce the gaps between feature maps extracted in the coarse mask prediction and the refined mask prediction. Since the visible mask head has the same network structure and the same parameters in predicting the coarse and refined mask, and the only difference is the inputs. The feature matching loss between the coarse visible mask prediction and refined visible mask prediction helps the network concentrate more on visible region appearance for visible mask segmentation and alleviate the effect of the background.

In our implementation, we use feature in the last two convolution layers of the visible mask head to measure the feature matching loss. We denote the loss function of the feature matching of the visible mask head as L_{vfm} , and use a subscript (j) on the visible mask head f_v to denote the feature maps in the i^{th} layer. Here $j = 4, 5$, which means we only use features in the higher level convolution layers. Then the feature matching loss of the visible mask head is

$$\mathcal{L}_{vfm} = \frac{1}{N \cdot S} \sum_{i,j} \lambda_j \mathcal{L}_S(f_v^{(j)}(\mathbf{F}_i), f_v^{(j)}(\mathbf{F}_i \cdot \mathbf{M}_{a,i}^c)), \quad (3)$$

where N, S is the number of instances and the number of convolution layers of the visible mask head respectively. And we apply cosine similarity \mathcal{L}_S (Zhu et al. 2019) in feature matching. We set the hyper-parameters $\lambda_4 = 0.01$, $\lambda_5 = 0.05$ and $\lambda_j = 0$ ($j \in \{1, 2, 3\}$).

The Amodal Mask Segmentation Module

The amodal mask segmentation module is designed to refine the coarse amodal mask by using the feature of the visible region and the shape prior. Inferring the amodal mask from the visible region appearance helps our model alleviate the misleading effect of the occlusion feature. Besides, different from the visible mask, for each category, the shapes of amodal masks are explicitly more stable, without the effect of arbitrary occlusions. The feature of the visible region can be obtained by using the refined visible mask \mathbf{M}_v^r from the visible mask segmentation module as visible attention.

In the pre-training phase, to model the shape prior, we collect the amodal ground-truth masks in the training set and use an auto-encoder to obtain the embedding of each amodal mask. Then, for each category, we partition the embeddings into K clusters via K-Means clustering and use the centers of these clusters as a codebook. Thus the codebook could memorize the embeddings of category-specific shape prior. This codebook is used to further refine the amodal mask.

In the training phase, for a predicted coarse amodal mask \mathbf{M}_a^c , we feed it into the encoder of the pre-trained auto-encoder to obtain the embedding. We use an L2 distance to find k nearest embeddings in the category-specific codebook according to the predicted category. This is denoted




Figure 3: The illumination of the shape prior post-process. The predicted amodal mask of box B gets a higher IoU value than the predicted amodal mask of box A. But the class score of box B (0.97) is lower than box A (0.99), which results in the suppression of box B in NMS. The shape prior post-process multiplies the shape prior similarity with the class score to refine the score of object. The score of A is $0.99 \times 0.69 \approx 0.68$ while the score of B is $0.97 \times 0.91 \approx 0.88$. The box B is remained while the box A is suppressed by NMS.

as shape prior search to obtain the shape prior embeddings. Then we feed these k nearest embeddings into the decoder to get decoded shape prior masks $\mathbf{M}_{sp}^k = f_{sp}(\mathbf{M}_a^c)$ ($\mathbf{M}_a^c \in \mathbb{R}^{H \times W}$, $\mathbf{M}_{sp}^k \in \mathbb{R}^{k \times H \times W}$). The f_{sp} denotes the operation using an auto-encoder with category-specific codebook for shape prior search. From this operation, we obtain the category-specific shape prior masks \mathbf{M}_{sp}^k which are the most similar to the coarse amodal mask in the shape prior. Then, we concatenate the feature of the visible region $\mathbf{F} \cdot \mathbf{M}_v^r$ and the k nearest shape prior amodal masks \mathbf{M}_{sp}^k as the input of the amodal mask head for amodal mask refinement. This process is to imitate the perception of human that infers the amodal mask objects by focusing on the appearance at the visible region and using the shape prior knowledge. The total loss term of this processing can be denoted by

$$\mathcal{L}_a^r = \frac{1}{N} \sum_i \mathcal{L}_{CE}(f_a(\text{cat}(\mathbf{F}_i \cdot \mathbf{M}_{v,i}^r, \mathbf{M}_{sp,i}^k)), \mathbf{M}_{a,i}^g), \quad (4)$$

where the f_a is the amodal mask head whose output is \mathbf{M}_a^r . The $\text{cat}(\cdot, \cdot)$ is the matrix concatenate operation. In our implementation, we set $k = 16$. Besides, similar to the equation (3) of the visible segmentation module, we also apply feature matching to the amodal mask head to enhance the capacity of focusing on the visible region. The loss function of feature matching in the amodal mask head is

$$\mathcal{L}_{afm} = \frac{1}{N \cdot S} \sum_{i,j} \lambda_j \mathcal{L}_S(f_a^{(j)}(\mathbf{F}_i), f_a^{(j)}(\mathbf{F}_i \cdot \mathbf{M}_{v,i}^r)). \quad (5)$$

In the inference phase, the shape prior can also be used to help the network further improve amodal segmentation. As shown in Fig. 3, we use the difference between the refined amodal mask and its nearest counterpart in the codebook $\|\mathbf{M}_a^r - \mathbf{M}_{sp}^k\|_2$ as a measurement to rank the scores of bounding box proposals, and filter out the proposals with larger differences. Here the difference is measured by an L1 distance. Here we set $k = 1$. If the predicted amodal mask shape of an instance is explicitly different from the nearest

Methods	D2SA					Visible		KINS					Visible		COCOA cls						
	Amodal							Amodal							Amodal					Visible	
	AP	AP50	AP75	AR	AP (Occ)	AP	AR	AP	AP50	AP75	AR	AP (Occ)	AP	AR	AP	AP50	AP75	AR	AP (Occ)	AP	AR
MRCNN	63.6	83.9	68.0	65.2	NA	68.9	70.1	30.0	54.5	30.1	19.4	NA	28.0	19.2	33.6	56.5	35.8	34.1	NA	30.1	31.5
MRCNN(C8)	64.8	84.0	70.7	65.6	NA	69.8	70.5	30.7	54.3	31.4	19.7	NA	28.7	19.3	34.7	57.5	36.9	35.4	NA	31.8	32.8
ORCNN	64.2	83.5	69.1	65.2	45.2	69.6	70.4	30.6	54.2	31.2	19.6	34.2	28.7	20.0	28.0	53.6	25.3	29.8	17.4	30.8	32.2
SLN	25.1	30.8	29.4	19.2	NA	NA	NA	6.6	10.7	6.9	6.1	NA	NA	NA	14.4	23.6	15.8	17.1	NA	NA	NA
Our method	70.2	85.1	75.8	69.1	51.1	72.2	71.8	32.0	55.3	33.3	20.9	37.4	29.8	19.8	35.4	56.0	38.6	37.1	22.1	34.5	36.4

Table 1: The comparison on the D2SA dataset, the KINS dataset, and the COCOA cls dataset. Because some methods only output the amodal mask prediction, the AP (Occluded) and visible mask prediction performance of them are unavailable (NA).

shape prior in the memory, it should be treated as a low-quality prediction even with a high class score.

The Implementation Details

Our model could predict the amodal mask and the visible mask. In the visible mask segmentation module and amodal mask segmentation module, we use the coarse amodal mask and refined visible mask as attention. However, if directly using the predicted masks in the warm-up phase, these inaccurate predictions might destroy the following parts. Thus, we design a weighting operation, where each instance is assigned an amodal weight and a visible weight to measure the weight of an instance in optimization. This operation can be found in supplementary. The final loss function is

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{reg} + \mathcal{L}_a^c + \mathcal{L}_v^c + \mathcal{L}_a^r + \mathcal{L}_v^r + \mathcal{L}_{rc} + \mathcal{L}_{afm} + \mathcal{L}_{vfm}. \quad (6)$$

Stochastic Gradient Descent (SGD) (Zinkevich et al. 2010) with weight decay is used for optimization in the training.

Experiments

Experimental Setting

We implement our proposed model based on Detectron2 (Wu et al. 2019) on the PyTorch framework. The main parameter setting is: For the D2SA dataset, batch size(2), learning rate (0.005), and the number of iteration (70000). For the KINS dataset, batch size(1), learning rate (0.0025), and the number of iteration (48000). For the COCOA cls dataset, batch size (2), learning rate (0.0005), and the number of iteration (10000).

Datasets. We evaluate the model performance for amodal segmentation on three datasets: the D2SA (D2S amodal) (Follmann et al. 2019), the KINS dataset (Qi et al. 2019), the COCOA cls dataset (Zhu et al. 2017).

The D2SA dataset is built based on the D2S(Densely Segmented Supermarket) dataset with 60 categories of instances. It contains 2000 images in the training set and 3600 images in the validation set, where the annotations of the amodal mask are generated by overlapping one to another.

The KINS dataset is built based on the KITTI dataset (Geiger, Lenz, and Urtasun 2012). It consists of 7474 images in the training set and 7517 images in the validation set. Different from the D2SA dataset, its amodal ground-truth is manually annotated. There are 7 categories about the autonomous driving task in the KINS dataset.

The COCOA cls dataset (Zhu et al. 2017) is built based on the COCO dataset (Lin et al. 2014). It consists of 2476 images in the training set and 1223 images in the validation set. There are 80 categories in this dataset.

Metrics. Following the (Zhu et al. 2017; Zhang et al. 2019), we use the mean average precision (AP) and mean average recall (AR) to evaluate performances. The AP (Occluded) is computed on the instances whose occlusion rate is larger than 15%. We use the evaluation api of the COCO dataset (Lin et al. 2014) for fair comparisons.

Baselines. We use following related methods for comparison.

(1) MRCNN indicates Mask RCNN (He et al. 2017) which predicts the amodal masks by a mask head consisting of 4 convolution layers and 1 deconvolution layer. We train two Mask R-CNN baselines predicting amodal mask and visible mask respectively since the Mask R-CNN cannot predict both of them simultaneously. The Mask RCNN (C8) means that the mask head has 8 convolution layers.

(2) ORCNN (Follmann et al. 2019) uses an amodal mask head and a visible mask head to infer the amodal masks and visible masks. The invisible mask is obtained by abstracting the visible mask from the amodal mask.

(3) SLN (Zhang et al. 2019) claims the importance of depth information in amodal segmentation. It uses a semantics-aware distance map to predict the amodal mask by utilizing depth order information.

Performance Comparison

We compare our model with all comparative methods on the datasets mentioned, and the performance comparisons are shown on Table 1. We can see that our model always outperforms other methods. Compared with Mask R-CNN and Mask R-CNN (C8), the improvement resulting from directly adding the depth of the mask head is explicitly lower than the improvement achieved by our model. Our method gets better performance mainly due to our reasonable design of the network rather than the expansion of the network.

We also show some qualitative results estimated by Mask R-CNN, ORCNN, and our method in Fig. 4. We can see that our method can segment more accurately than other methods, owing to the help of the attention on the visible region and the shape prior. For the 1st and 2nd row, the predictions of our method are not misled by the feature of occlusions such as the glass bottle and cabbage. For the 3rd and 4th row,




Figure 4: Examples show the images, the ground-truth amodal masks, estimations of Mask RCNN, ORCNN and ours.

Ablation Study					D2SA				KINS					
Visible Attention	Reclass	Shape Prior Refinement	Shape Prior Post-process	Feature Matching	Amodal		Visible		Amodal		Visible			
					AP	AR	AP (Occ)	Visible AP	Visible AR	AP	AR	AP (Occ)	AP	AR
1	✓	✓	✓		66.67	65.68	46.40	70.50	70.24	31.61	20.16	36.84	29.44	19.61
2	✓		✓	✓	69.02	68.09	50.31	71.66	70.83	31.65	20.55	36.93	29.25	19.81
3	✓	✓		✓	68.84	67.42	49.08	71.53	70.70	31.70	20.33	37.25	29.45	19.63
4	✓	✓	✓		68.15	68.30	50.03	70.71	70.43	31.87	20.46	37.48	29.53	19.70
5	✓	✓	✓	✓	69.98	68.87	51.01	71.92	71.15	31.94	20.60	37.55	29.61	19.88
6	✓	✓	✓	✓	70.27	69.17	51.17	72.28	71.85	32.08	20.90	37.57	29.88	19.88

Table 2: The ablation studies results on the D2SA dataset and the KINS dataset.

our method keeps robust even the occlusion rate is large.

Ablation Studies

We conduct the ablation studies on both the D2SA dataset and the KINS dataset. All the results are shown on Table 2.

The Effect of Visible Attention. To evaluate the effect of visible attention in refining the amodal mask, we design the baseline refining the amodal mask without using the visible mask as attention. The input of the amodal mask head for predicting refined amodal mask is the concatenation of ROI feature F and shape prior M_{sp}^k . Experimental results are shown at the 1st and 5th rows on Table 2.

The Effect of Reclassification Regularizer. To validate the effectiveness of the reclassification regularizer. We conduct the experiments at the 2nd and 5th rows on Table 2. The experimental results show the importance of the reclassification regularizer.

The Effect of the Shape Prior Refinement. To investigate the effect of shape prior in refinement, we plan to train the baseline at the 3rd. We only utilize the feature of the refined visible region to refine the amodal mask. The gap between the results at the 3rd row the 5th row shows the im-

portance of shape prior knowledge, which agrees with the usage of shape prior in the human’s amodal perception.

The Effect of the Shape Prior Post-process. To investigate the effect of shape prior post-process, we plan to remove the shape prior post-process based on the baseline at the 5th row. The result is shown at the 4th row. Compared with the results at the 5th row, the gap shows that the utilizing of shape prior post-process can achieve improvement.

The Effect of Feature Matching. To validate the effect of feature matching in our method, we train a baseline via removing the feature matching on both the amodal and visible mask head. The result is shown in the 5th row on Table 2. Compared with our model at 6th row, the feature matching could achieve further better performance, which indicates that the feature matching could further help the amodal mask head learn the feature extracted from the visible regions.

Attention Analysis

To evaluate the effect of amodal and visible attention in our model, we propose 4 different types to utilize amodal and visible attention. The details are shown in supplementary. We conduct the experiments on Table 3 to compare different

Types of Utilizing Attention	D2SA						KINS							
	Amodal					Visible		Amodal					Visible	
	AP	AP50	AP75	AR	AP (Occ)	AP	AR	AP	AP50	AP75	AR	AP (Occ)	AP	AR
(a) Both Self Attention	64.55	82.53	68.96	65.62	44.31	70.50	71.02	31.18	54.24	32.08	20.15	36.89	29.27	19.76
(b) Only Visible Attention	66.98	83.99	72.87	67.78	47.53	70.92	71.50	31.32	54.79	32.23	20.25	36.98	29.40	19.81
(c) Cross Attention	67.11	83.90	72.81	67.70	47.75	71.79	71.77	31.57	54.58	32.54	20.61	37.13	29.65	20.09
(d) Ours	67.33	84.10	72.97	68.06	47.91	71.88	72.43	31.69	54.52	32.96	20.75	37.30	29.68	20.01

Table 3: The experimental results of the attention analysis on the D2SA dataset and the KINS dataset.

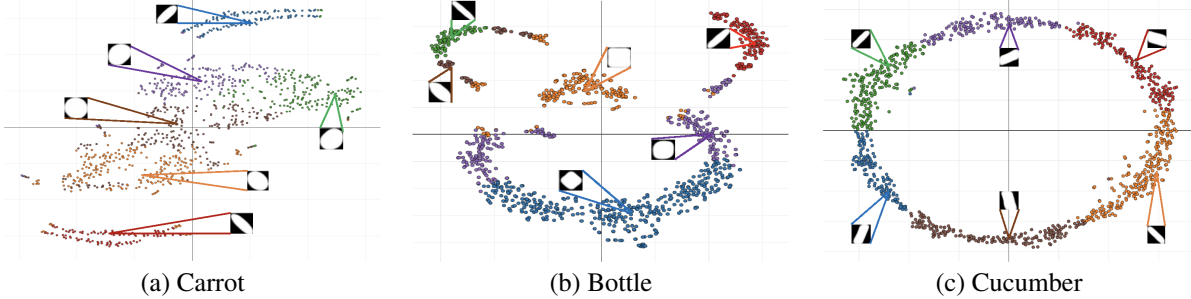


Figure 5: Visualization of shape prior embeddings in shape codebook. We use t-SNE (Van der Maaten and Hinton 2008) to visualize the latent space of each category in the codebook. This proves that the learned shape prior has certain interpretability.

ways to use amodal and visible attention. The shape prior and reclassification regularizer are removed in this section.

(a) Both Self-Attention uses the coarse amodal mask and coarse visible mask as attention to refine amodal mask and visible mask respectively. The input of the amodal mask head and the visible mask head in attention-based refinement are the feature of the amodal region $\mathbf{F} \cdot \mathbf{M}_a^c$ and the feature of the visible region $\mathbf{F} \cdot \mathbf{M}_v^c$ respectively.

(b) Only Visible Attention uses the coarse visible mask as attention to refine the visible mask prediction, which indicates $\mathbf{M}_v^r = f_v(\mathbf{F} \cdot \mathbf{M}_v^c)$. Then, the refined visible mask \mathbf{M}_v^r is utilized as attention to refine amodal mask prediction. The refined amodal mask is obtained by $\mathbf{M}_a^r = f_a(\mathbf{F} \cdot \mathbf{M}_v^r)$.

(c) Cross Attention uses the coarse amodal mask and coarse visible mask as attention to refine the visible and amodal mask prediction respectively. The formula is $\mathbf{M}_a^r = f_a(\mathbf{F} \cdot \mathbf{M}_v^c)$ and $\mathbf{M}_v^r = f_v(\mathbf{F} \cdot \mathbf{M}_a^c)$.

(d) Ours uses the coarse amodal mask as attention to refine the visible mask, $\mathbf{M}_v^r = f_v(\mathbf{F} \cdot \mathbf{M}_a^c)$. Then, using the refined visible mask as attention to refine the amodal mask, $\mathbf{M}_a^r = f_a(\mathbf{F} \cdot \mathbf{M}_v^r)$.

We can observe that utilizing the visible mask as attention to refine amodal mask (b,c,d) achieves explicitly better performance on amodal mask prediction than using the coarse amodal mask (a). This indicates that applying visible attention for amodal mask prediction is more reasonable than applying amodal attention. Using the amodal mask as attention (c,d) to refine visible mask can get better performance on visible mask prediction than using the coarse visible mask (a,b). This indicates the effect of applying amodal attention in visible mask prediction. Using the refined visible mask as attention to refine amodal mask (d) has slight

improvement than using the coarse visible mask (c). This result shows that using more accurate visible attention can obtain improvement in amodal mask prediction.

Visualization of Shape Prior in the Codebook

We also show some category-specific shape prior such as carrot, bottle and cucumber, via t-SNE (Van der Maaten and Hinton 2008) in Fig. 5. For each category, we partition the latent feature into 1024 clusters via K-Means, since we need to use many redundant shape prior items to store the various changes of rotation. Thus, there exist a huge number of shape prior in each category in Fig. 5 (a)-(c). In particular, Fig. 5 (c) represents the learned codebook of cucumber, just like a circle. As we know, cucumber is a formable object with little shape changes, while the 6 amodal masks reflect its rotation changing process. This visualization shows that the shape prior has certain interpretability.

Conclusion

In this work, we propose a novel model to mimic the human amodal perception using the shape prior to imagine the invisible regions mainly based on the feature of visible regions. However, almost existing methods use the appearance of the whole region-of-interest to infer the amodal masks, which is against the human amodal perception. And this strategy brings the ambiguity that the same appearance of occlusion may require different predictions. To simulate the imagination from visible region and shape prior, we use the visible mask as attention to focus on the visible regions and build a codebook to store the collected amodal shape prior embeddings. The experimental results indicate our method outperforms other state-of-the-art methods.

Acknowledgments

The work was supported by National Key R&D Program of China (2018AAA0100704), NSFC #61932020, Science and Technology Commission of Shanghai Municipality (Grant No. 20ZR1436000) and ShanghaiTech-Megvii Joint Lab.

References

- Ba, L. J.; and Caruana, R. 2013. Do Deep Nets Really Need to be Deep? *Advances in Neural Information Processing Systems* 2654–2662.
- Bai, M.; and Urtasun, R. 2017. Deep watershed transform for instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5221–5229.
- Chen, K.; Pang, J.; Wang, J.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Shi, J.; Ouyang, W.; et al. 2019. Hybrid task cascade for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4974–4983.
- Chen, Y.; Lin, G.; Li, S.; Omar, B.; Wu, Y.; Wang, F.; Feng, J.; Xu, M.; and Li, X. 2020. BANet: Bidirectional Aggregation Network with Occlusion Handling for Panoptic Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Cheng, Y.; Yang, B.; Wang, B.; Yan, W.; and Tan, R. T. 2019. Occlusion-aware networks for 3d human pose estimation in video. In *Proceedings of the IEEE International Conference on Computer Vision*, 723–732.
- Chu, X.; Zheng, A.; Zhang, X.; and Sun, J. 2020. Detection in Crowded Scenes: One Proposal, Multiple Predictions. *Proceedings of the IEEE International Conference on Computer Vision*.
- Follmann, P.; Nig, R. K.; Rtinger, P. H.; Klostermann, M.; and Ttger, T. B. 2019. Learning to see the invisible: End-to-end trainable amodal instance segmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1328–1336. IEEE.
- Fu, J.; Liu, J.; Wang, Y.; Zhou, J.; Wang, C.; and Lu, H. 2019. Stacked deconvolutional network for semantic segmentation. *IEEE Transactions on Image Processing*.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 3354–3361. IEEE.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, X.; Ge, Z.; Jie, Z.; and Yoshie, O. 2020. NMS by Representative Region: Towards Crowded Pedestrian Detection by Proposal Pairing. *Proceedings of the IEEE International Conference on Computer Vision*.
- Huang, Z.; Wang, X.; Huang, L.; Huang, C.; Wei, Y.; and Liu, W. 2019. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 603–612.
- Kar, A.; Tulsiani, S.; Carreira, J.; and Malik, J. 2015. Amodal completion and size constancy in natural scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, 127–135.
- Li, K.; and Malik, J. 2016. Amodal instance segmentation. In *European Conference on Computer Vision*, 677–693. Springer.
- Li, Q.; Jin, S.; and Yan, J. 2017. Mimicking Very Efficient Network for Object Detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, S.; Qi, L.; Qin, H.; Shi, J.; and Jia, J. 2018. Path aggregation network for instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8759–8768.
- Neven, D.; Brabandere, B. D.; Proesmans, M.; and Gool, L. V. 2019. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8837–8845.
- Peng, S.; Jiang, W.; Pi, H.; Bao, H.; and Zhou, X. 2020. Deep Snake for Real-Time Instance Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Pinheiro, P. O.; Collobert, R.; and Dollár, P. 2015. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, 1990–1998.
- Qi, L.; Jiang, L.; Liu, S.; Shen, X.; and Jia, J. 2019. Amodal instance segmentation with kins dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3014–3023.
- Ramamonjisoa, M.; Du, Y.; and Lepetit, V. 2020. Predicting Sharp and Accurate Occlusion Boundaries in Monocular Depth Estimation Using Displacement Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14648–14657.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Ren, M.; and Zemel, R. S. 2017. End-to-end instance segmentation with recurrent attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6656–6664.

- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Ren, Y.; Yu, X.; Zhang, R.; Li, T. H.; Liu, S.; and Li, G. 2019. StructureFlow: Image Inpainting via Structure-aware Appearance Flow. In *Proceedings of the IEEE International Conference on Computer Vision*, 181–190.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9(11).
- Winn, J.; and Shotton, J. 2006. The layout consistent random field for recognizing and segmenting partially occluded objects. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Xie, E.; Sun, P.; Song, X.; Wang, W.; Liu, X.; Liang, D.; Shen, C.; and Luo, P. 2020. Polarmask: Single shot instance segmentation with polar representation .
- Xu, W.; Wang, H.; Qi, F.; and Lu, C. 2019. Explicit shape encoding for real-time instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5168–5177.
- Yang, J.; Ren, Z.; Xu, M.; Chen, X.; Crandall, D. J.; Parikh, D.; and Batra, D. 2019. Embodied Amodal Recognition: Learning to Move to Perceive Objects. In *Proceedings of the IEEE International Conference on Computer Vision*, 2040–2050.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE International Conference on Computer Vision*, 4471–4480.
- Zhang, T.; Huang, B.; and Wang, Y. 2020. Object-Occluded Human Shape and Pose Estimation From a Single Color Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7376–7385.
- Zhang, Z.; Chen, A.; Xie, L.; Yu, J.; and Gao, S. 2019. Learning semantics-aware distance map with semantics layering network for amodal instance segmentation. In *Proceedings of the 27th ACM International Conference on Multimedia*, 2124–2132.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9308–9316.
- Zhu, Y.; Tian, Y.; Metaxas, D.; and Dollár, P. 2017. Semantic amodal segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1464–1472.
- Zinkevich, M.; Weimer, M.; Li, L.; and Smola, A. J. 2010. Parallelized stochastic gradient descent. In *Advances in neural information processing systems*, 2595–2603.