参赛队员姓名：杨景程

中学：深圳国际交流学院

省份：广东省

国家/地区：中国

指导教师姓名：吴邵汐

指导教师单位：深圳国际交流学院

论文题目：DAISnet: Depth-aware Amodal Instance Segmentation using 3D Feature Extraction and Occlusion Refinement

DAISnet: Depth-aware Amodal Instance Segmentation using 3D Feature Extraction and Occlusion Refinement

Jingcheng Yang

Shenzhen College of International Education, Shenzhen, China

Abstract. In recent years, classical computer vision problem, e.g. instance segmentation, has made great progress. However, amodal perception, has yet not been paid much attention. Amodal instance segmentation, as the representative of amodal perception, requires additional segmentation of occluded regions of an object. Amodal instance segmentation is difficult, as it requires a high-level understanding of visual inputs and features. Amodal instance segmentation is also important to down-stream tasks like robotics vision and autonomous driving as it offers vehicles additional spacial information to traverse in complex situations, while enhancing machines' ability to navigate path to target and object retrieval in complex environments.

On the principle that depth is an inherent part of occlusion, we propose a novel depth-aware amodal instance segmentation network - DAISnet, to fully utilize depth information in amodal instance segmentation. We introduced methods to implicitly and explicitly use depth information for amodal instance segmentation. DAISnet's backbone is entirely trained using depth features, making it a fully depth-aware network. We propose processing methods like regional depth normalization and tangent normal computation for implicit depth and 3D feature extraction. We propose an Occlusion Edge refinement layer that utilizes localized depth to infer potential regions of occlusion. We improve upon works on shape prior memorization by creating a 3D encoder that is capable of creating embeddings based on depth and 2D shape features. It is expensive to obtain pixel accurate depth and amodal segmentation ground truths in depth and amodal segmentation problems. To address this, we propose a new Real-time Generated Synthetic Dataset, also named as DAIS-RGSD to overcome this problem. DAIS-RGSD includes pixel accurate depth and amodal masks, as well as photo-realistic RGB images. We evaluated our method using the KINS benchmark, and received remarkable results proving the importance of depth information and the effectiveness of DAISnet.

Keywords: Instance Segmentation, Amodal Perception, Deep Learning, Synthetic Data Collection, Depth Prediction

Table of Contents

1 Introduction . . . 1
2 Related Work . . . 2
2.1 Instance Segmentation . . . 2
2.2 Monocular Depth Estimation . . . 2
2.3 Amodal Instance Segmentation . . . 3
3 Method . . . 3
3.1 Overview . . . 3
3.2 Synthetic Data Collection . . . 4
3.3 Monocular Depth Prediction . . . 5
3.4 Normalized Tangent Map . . . 7
3.5 Depth-Aware Feature Backbone . . . 10
3.6 ROI Pooling . . . 10
3.7 Occlusion Edge Refinement . . . 11
3.8 3D Shape Prior Refinement . . . 12
4 Experiment . . . 13
4.1 Datasets . . . 13
4.2 Model Set-up . . . 14
4.3 Metrics . . . 15
5 Evaluation and Ablation Study . . . 17
5.1 Backbone . . . 17
5.2 Depth . . . 17
5.3 Normal Tangent . . . 18
5.4 Occlusion Edge Refinement . . . 19
5.5 3D Shape Prior Refinement . . . 20
5.6 Combinations . . . 20
5.7 Limitation and Future Works . . . 20
6 Conclusion . . . 21
7 References . . . 22
8 Acknowledgement . . . 24

1 Introduction

Amodal instance segmentation is the segmentation of the visible and occluded region of all objects from an image, as shown in Fig 1. Visible objects can be segmented with a high degree of accuracy thanks to progress made in the last few years. Segmenting occluded regions of an object, is much more difficult to achieve. There is no single probable solution for some occluded regions as Fig 2 demonstrates, and it requires a high level understanding of the image semantics to achieve reasonable predictions. The annotation of amodal instance segmentation is also expensive, as it requires annotators to thoroughly examine all possibilities. Although the task appears ill-posed, past works have demonstrated that high quality amodal instance segmentation is achievable and helpful. Amodal instance segmentation is also a necessary stepping stone for tasks that require greater amodal awareness, such as 3D amodal segmentation and amodal reconstruction. Amodal instance segmentation is also useful in numerous down-stream tasks. Self-driving vehicles can make better decisions on road using occlusion and spacial information. Occlusion information offers robust information on the order of vehicles for better navigation, while spacial information can help deduce viable parking regions. Amodal instance segmentation is also important in robotic vision. Navigation towards an occluded object, efficient retrieval of objects in complex environments, and preemptive measures all benefit from accurate amodal instance segmentation predictions.

Fig. 1: Amodal instance segmentation segments the entire object while (modal) instance segmentation only segments visible regions

Past works regarding amodal instance segmentation involves the use of Mask R-CNN architecture which have achieved state of the art performance in visible instance segmentation. [1–5] Direct implementation of such architecture have yielded reasonable results. Most amodal instance segmentation networks improve upon this architecture. Notably, Xiao et al.[6] introduced a code-book mechanism that conducts feature mapping based on pretrained 2D shape awareness. Amodal mask are refined through this process, improving the accuracy of such predictions. However, there are limitations in representing 3D objects in a 2D context, which neglects actual spacial relationship between objects, as well as it's 3D features. We believe depth is inherent to occlusion, as it provides information on possible occluded regions and additional features of visible regions. More information about the visible regions of an object can help make better predictions on amodal predictions.

We improve upon previous solutions by creating a depth-aware amodal instance segmentation network (DAISnet) through a monocular depth prediction layer pretrained on synthetic and real-life data. Depth information is incorporated into the backbone feature extraction pipeline to make our network fully depth-aware. Depth information is regionally normalized and computed into a normal tangent map through a normal tangent encoder. Normal maps help describe the change in 3D surface while being distance invariant. To best utilize tangent and depth information, we introduce the Edge Occlusion and Tangent Occlusion modules which relies on normalized depth and tangent information of objects to predict possible regions of occlusion, which is essentially a rough prediction of amodal masks. We base this from the principle that areas directly touching visible objects with smaller depth (closer to the observer) have a very high possibility of being the occluder. We refine amodal mask predictions using depth information for 3D shape prior mapping, which provides rich features even under circumstances with obscure shapes. We used a 3D real-time graphics engine to generate synthetic data of objects with occlusion, photo-realistic colors, depth and amodal segmentation masks. This workflow allows us to generate large quantities of data with accurate pixel-level ground truth based on various scenarios for training.

Our contribution can be summarized as:

Fig. 2: There can be multiple amodal predictions given the nature of uncertainty. P3 is the closest to ground truth, but both P1 and P2 are possible predictions. This shows that it's difficult to obtain accurate results in amodal instance segmentation

- A novel depth-aware network for amodal instance segmentation.
- A synthetic dataset that can be used to train depth prediction, 3D feature extraction and amodal instance segmentation tasks.
- A novel approach at extracting depth information using normal tangent maps and regional depth normalization.
- A post refinement RoI mask head for amodal completion using depth and tangent information.
- A 3D codebook mechanism for amodal mask refinement based on 3D shape prior.

2 Related Work

2.1 Instance Segmentation

Instance Segmentation segment and classify instances' visible part separately from the image. Amodal instance segmentation extends directly from research in instance segmentation, which itself extended from research regarding object detection. Ross at el. [7] introduced the R-CNN network, which selects region of interests from an image and conduct convolutions on these regions. Regional CNNs have supersede all previous implementations in performance and accuracy. Ross at el. [8] introduced the Fast R-CNN network that significantly faster than R-CNN, which was later improved again into Faster R-CNN [1], which improved proposal quality and performance by using a Region Proposal Network for proposal generations.

He et al. extends Mask R-CNN [2] to instance segmentation by introducing a Mask Head, which conducts segmentation on detected objects using work derived from Long at el. [9] on semantic segmentation using Fully Convolutional Networks(FCN). He also proposed RoI Align solution in place of RoI pooling which achieved pixel accurate region of interests. Chen et al. [3] incorporated Cascade techniques [4] to Instance Segmentation and improved accuracy through stage specific refinement. Ding et al. [5] extends cascade instance segmentation by taking advantage of separate stages to establish a bidirectional relation between the mask and the bounding box. By introducing a mask guided RoI align method, the bounding box and mask prediction continues to refine each other for accurate instance segmentation. Mask Scoring R-CNN [10] aims to reevaluate predicted masks by giving lower scores for masks with bad quality, and thus improve mask prediction.

2.2 Monocular Depth Estimation

Monocular depth estimation use one image to inference depth of each pixel to achieve a 2D-3D projection. Monodepth [11] proposed a self-supervised training pipeline for monocular depth estimation. They used a FCN

to predict disparities for rectified stereo images and supervise it via an image reconstruction loss. Godard et al.[12] extends the monodepth model to a video-based monocular depth estimator. Performance was improved by introducing a Per-Pixel Minimum Reprojection Loss, Auto-Masking Stationary Pixels, and Multi-scale Estimation. Fu et al. [13] directly regress the depth through a CNN and introduced a method for monocular depth estimation that relies on space-increasing discretization, differing from past implementation by its smaller architecture and faster convergence, achieved very accurate results. Lee at al. [14] utilizes planar guidance layers in multiple stages of decoding, which utilizes light and color changes to generate surface information for depth prediction.

2.3 Amodal Instance Segmentation

As work regarding modal instance segmentation begins to yield state-of-the-art results, research have been done to extend instance segmentation to occluded region of objects. Some early work involves directly extending and modifying existing R-CNN architecture to support amodal instance segmentation [15–17] Li at al. [18] used an iterative bounding box method to filter amodal instance segmentation results from a CNN. Xiao et al. [6] used a 2D-shape prior codebook that conducts feature matching to enhance amodal instance segmentation predictions. Zhang et al. [19] introduces a semantics-aware distance map that allows for pixel level amodal segmentation and occlusion order prediction. Others have tried to include abstract feature awareness to achieve more accurate amodal segmentation predictions [20] Yang et al. [21] extended amodal instance segmentation using multi-perspective images that are taken at positions determined by the network. Deng [22] used depth information to predict amodal bounding box of 3D objects with large success. This demonstrates that 3D visual features can enhance amodal predictions overall and the 2.5D visual features are correlated to 3D object sizes, locations, and orientations.

3 Method

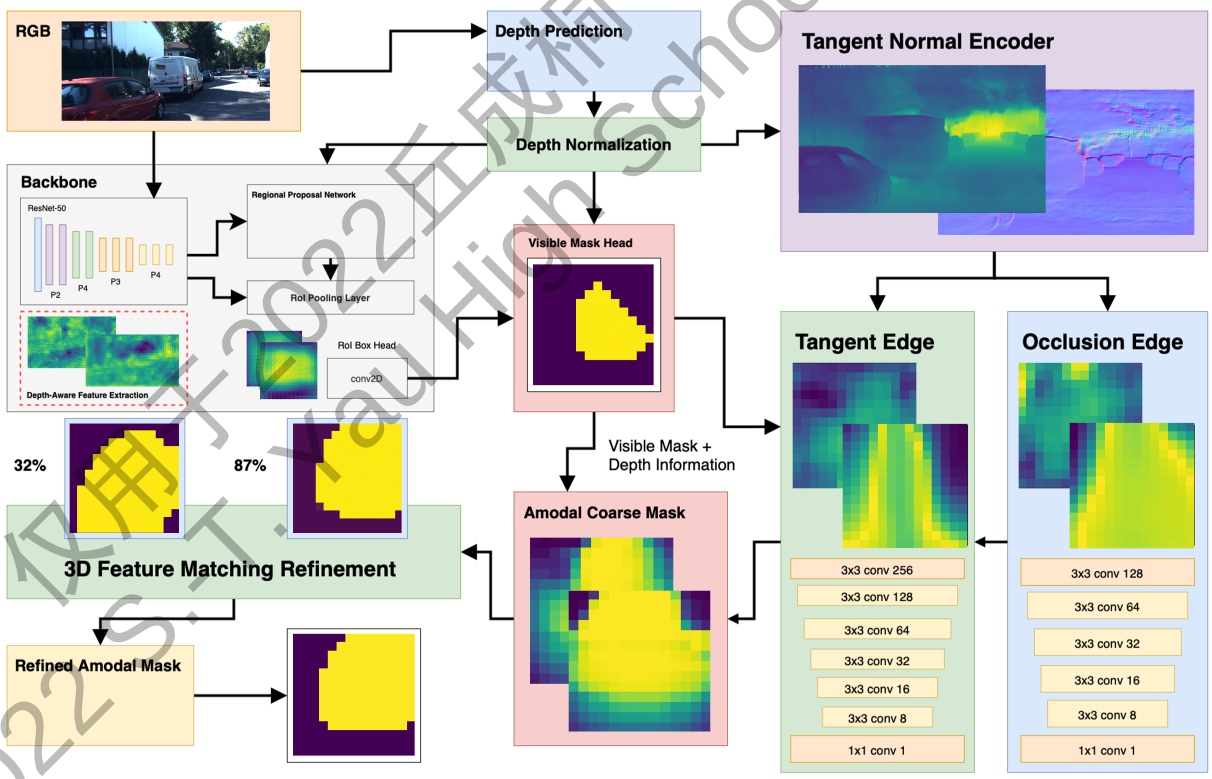

Fig. 3: Network overview

3.1 Overview

Our network is improved upon the Mask-RCNN network that was adapted for amodal instance segmentation. We first obtained depth information using a monocular depth prediction network [14]. We implemented different

methods of using depth information for amodal segmentation. Fig 3 demonstrates our complete network. Predicted depth information is first passed to the feature extraction backbone alongside RGB data.The backbone extracts feature maps and generate proposals using the Regional Proposal Network; Since depth information was included in feature extraction, this entire process is depth-awared.

Depth information is also explicitly given to the RoI pooling layer to generate pooled depth regions with the same dimension as the object proposals. Box proposals are then used to generate visible mask predictions through the visible mask head. The visible mask and depth information are both passed to the 3D-aware refinement layers, which consists of occlusion edge and 3D shape prior refinements.

The occlusion edge refinement layer uses localized depth values to deduce possible regions of occlusion, which is then given to an occlusion refinement head that refines the amodal mask using occlusion edges. The 3D shape prior is based upon 2D shape prior works done by Xiao et al. [6]. The shape prior mechanism uses deep learning to effectively encode and decode shapes. Statistical operations such as clustering can be done on encoded shapes, which is used for shape matching and refinement based on prior seen shapes. 3D shape prior uses localized depth information during encoding, which means that the stored codebook contains 3D representations of objects. After refinement through both modules, the network returns the final amodal instance segmentation results.

3.2 Synthetic Data Collection

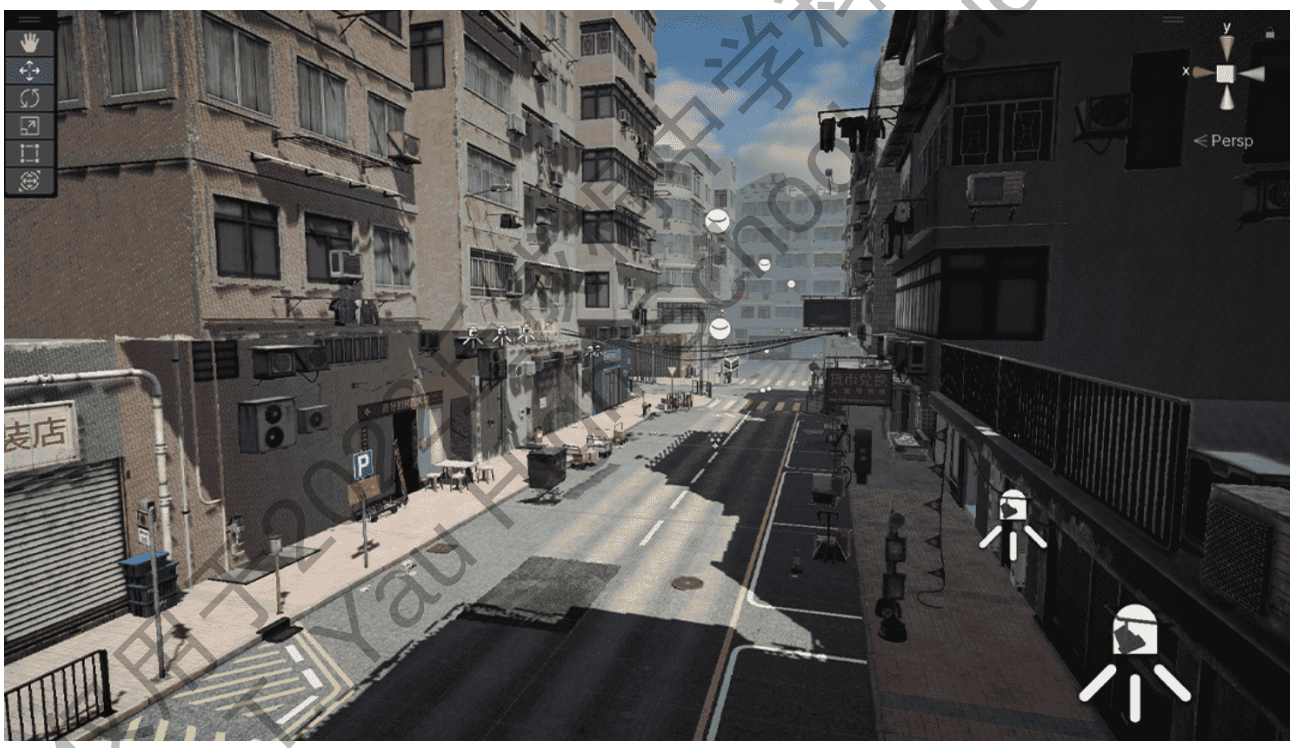

Fig. 4: High-fidelity street scene created within Unity Engine using premade assets. The camera will move through the street, and objects will be placed around it. This mimics many of the images in KITTI which was taken in a narrow street with vehicles.

To create a dataset that contains both segmentation and depth data, we decided to develop a synthetic dataset using Unity Engine©. Similar to the KINS dataset used in past research, our synthetic dataset focuses on amodal instance segmentation of Pedestrians and Vehicles in street and road scenarios. For our street scenario, we setup a 3D street that spans for 100 meters long. Buildings and other items are placed for realism and to ensure our network knows how to ignore distractions. The camera moves randomly through the street within a fixed boundary.

$$C_{xyz} \begin{cases} -5 \le x \le 5 \\ 1 \le y \le 2 \\ 0 \le z \le 50 \end{cases} \quad V \begin{cases} V_n &= r(2,3) \\ V_p &= [\, r(-10,10),\ \gamma,\ z + r(10,50)\,] \\ V_r &= [\, 0,\ r(0,360),\ 0\,] \end{cases} \tag{1}$$

Where for vehicle V: V_n denotes number of vehicles in the scene, V_p denotes position vector of a specific vehicle , γ denotes a constant y value for vehicle height, V_r denotes rotation of V euler angles, $r(min, max)$ specifies a random value between two inclusive values.

We assigned bounding boxes for each of the object and use it to detect for collisions. Overlapping objects will continue to move until it found a suitable location. For pedestrians, we opted for 2D cutouts placed among the street using a very thin collision box. The ratio of vehicles and pedestrians is roughly 1:3. The directional light in the scene is also randomly changed per image, so to make sure our dataset is not prone to shadow and light variances. Post processing image effects such as bloom, which blurs bright regions of an image to mimic bright light is also used for photo realism.

After the scene is set up, the camera will render four passes. The RGB pass - objects with proper lighting and albedo textures. The Depth pass, where a fragment shader will compute linear depth between 0 and 50 meters using geometric data. Unlike other depth datasets like KITTI which uses LiDAR, our ground truth depth does not contain gaps and is pixel accurate, this significantly improves the quality of our ground truth data. The final two passes are modal and amodal masks of each generated object, which is generated by assigning an unlit color to the target geometry, while other objects only effect it through occlusion for visible mask capture.

While the street scenario focuses on Pedestrians and few cars in a complex environment, the open roads scenario generates data with lots of vehicles. A simple road network in the form of a graph is established. Vehicles will move through this road network, simulating traffic and road scenarios. Our camera follows one of these vehicles, for every random interval between 1 and 2 seconds, the camera will initiate a scene check to determine is enough vehicles are infront of it for valid data capture.

Fig. 5: The Openroads subset of RGSD, which contains RGB, depths and annotations of amodal and visible masks.

3.3 Monocular Depth Prediction

To predict amodal instance segmentation on datasets like KINS which do not provide depth data, we rely on a pretrained monocular depth prediction module. We used the BTS network [14] with a backbone trained using ResNet50 on the KITTI Eigen Split depth dataset. We also trained the BTS network [14] on our own DAIS-RGSD dataset. This monocular prediction module is inserted before the actual backbone. By doing this, our network is able to predict a depth map of our RGB input, and send it to the module. The depth information is also passed to the RoI Head module for occlusion edge and amodal proposal calculations.

Depth values are in an entire different representation than RGB values within the range of 0-255. Our model normalizes RGB values using standard deviation and mean. Notably:

$$z = \frac{x - \mu}{\sigma} \tag{2}$$

Where x denotes a pixel of a single channel, z denotes the normalized pixel of a specific channel, μ denotes the mean of $\sum_i^w \sum_j^h x_{ij}$, σ denotes the standard deviation of $\sum_i^w \sum_j^h x_{ij}$

The same normalization is applied over depth d, but its exponential depth value is first linearized to achieve linear depth:

$$d_l = log_{10}(d) \tag{3}$$

Fig. 6: Monocular depth prediction from a single RGB image from the KITTI dataset, the model used the BTS Network with a DenseNet161 backbone.

It is not possible for depths to be normalized on a per object basis, as that would require the complete segmentation of objects within the image accurately. Since we want the network to be entirely depth-aware from the beginning, other normalization methods have to be used. The approach using visible regions to normalize depth per object however, is used in the Edge Occlusion refinement head.

Although it is not possible to completely remove object distance variance entirely without object data, we can significantly reduce its impact using a Regional Depth Normalization scheme. The scene is being split into subregions occupying a slice of the full distance, the size of a subregion is called the step size or D_s. We use the floor function to acquire the subregion of any given depth:

$$D_{ir} = D_i - \left\lfloor \frac{D_i}{D_s} \right\rfloor \times D_s \tag{4}$$

Where D_i represents the depth value of a single pixel, D_ir represents the base reference distance / subregion this pixel belongs to. We obtain the results shown in Fig 7 B, although the floor function divides the image into subregions, there exists a hard edge between regions. For this, we introduce a soft floor function:

$$y = x - \frac{\sin(2\pi x)}{2\pi} \tag{5}$$

Where x is rounded down to the greatest integer less than x with a smooth interpolation between integers. We use this function to pass in D_i/D_S as x, obtaining the result in Fig 7 C.

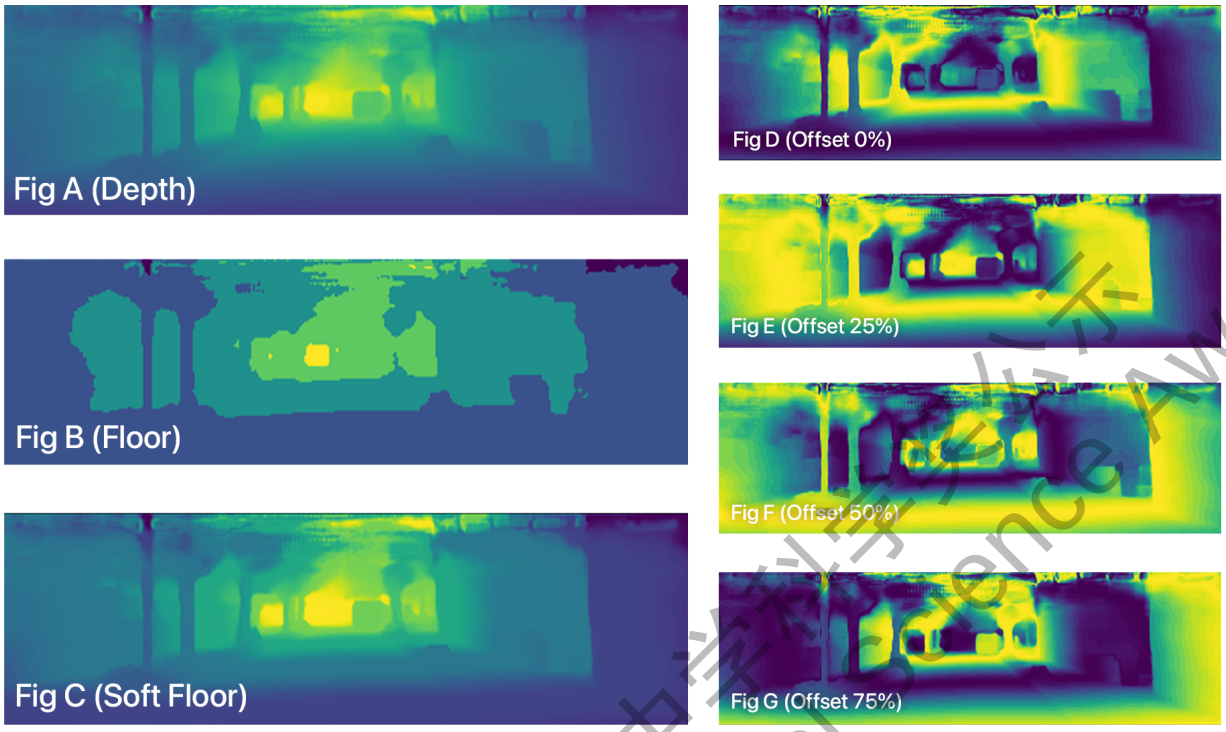

Fig. 7: Regional depth normalization pass.

Although this method greatly decreases variance caused by difference in distances, regional normalization still means that objects will have some variance in distance as seen in Fig 7 D. This could be minimized by decreasing the step size, but it will make the data less perceivable to the backbone and feature extraction. To address this, we introduce an offset value, which creates multiple regional depth maps at different global offsets. We specify the number of offset maps as D_m, which in our setup is 4 by default. Each offset map should have a different offset value, which is obtained simply with D_s/D_m. This means we are creating D_m different regional depth maps over equal intervals between the step size D_s. This produces multiple regional depth maps as shown from Fig 7 D to Fig 7 G. A D_m of 4 is suitable as each object now has a representation of its regional depth in roughly all possible ranges. The remaining variance in distance can be ignored due to its insignificance. This makes individual objects almost invariant to distance, as it includes a normalized depth representation of itself in ranges [(0, 0.25), (0.25, 0.5), (0.5, 0.75), (0.75, 1)] approximately. This means the the backbone will receive 4 additional channels as depth inputs, for each regional depth map. Experiments on the effect of various D_m and D_s on the quality of object detection and mask prediction will be conducted.

3.4 Normalized Tangent Map

Depth information however, has extreme variance shifts by nature. Fundamentally, RGB values represent the intensity of each primary color; Depth values represents distance from the camera, this poses a problem, as two objects of the same orientation, shape and feature at two distances will produce completely different sets of depth values. Thus the problem of normalizing global depth to match the RGB input is ill-posed, as there is no perfect way to normalize the depth of every object prior to object extraction and segmentation. Which poses a problem in using depth information for backbone feature extraction.

A way to remove distance variance of depth information is to compute normal tangent maps. Normal tangent maps represent the tangent of a surface at a specific pixel in the form of a normalized vector, where its xyz component translates to RGB respectively, creating an image of 3 channels. Tangent normal maps have been used extensively in computer graphics as it is an easy way of representing 3D surfaces without actual geometry. In our case, we essentially obtained 3D features of an image with depth invariance. Tangent normal maps are also represented in RGB, which our network backbone is already designed and tested for.

To compute tangent normal maps, we must first obtain the partial derivatives of every pixel. Because depth data is discrete, we can only estimate its gradient at a given pixel. We use the Sobel–Feldman operator, a discrete

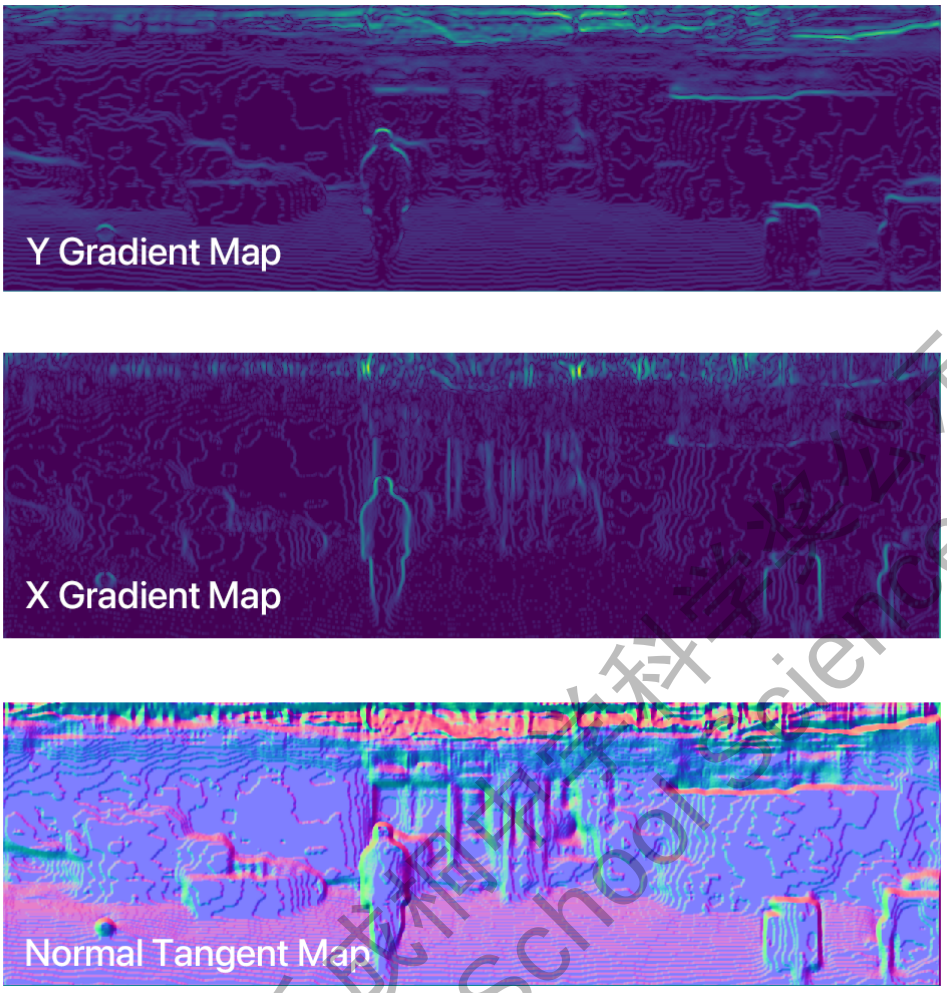

Fig. 8: Normal tangent information obtained using the same depth map in the previous figure. The gradient maps show the raw outputs of a 5x5 Sobel filter over D_x and D_y. The final normal tangent map uses a D_z of 0.7.

differentiation operator that produces approximations of gradients. For a kernel size of 3x3, where a gradient of only the adjacent pixels are accounted for, we denote adjacent pixels to origin pixel o from a to i.

$$\begin{bmatrix} a & b & c \\ d & o & f \\ g & h & i \end{bmatrix} \tag{6}$$

We obtain a gradient G between o and any pixel by dividing it's difference by distance.

$$\begin{bmatrix} \frac{a-o}{\sqrt{2}} & b-o & \frac{a-c}{\sqrt{2}} \\ c-o & 0 & f-o \\ \frac{g-o}{\sqrt{2}} & h-o & \frac{i-o}{\sqrt{2}} \end{bmatrix} \tag{7}$$

To obtain the gradient in a given direction, we must first obtain the dot product of our directional vector z_x in the x direction or z_y in the y direction with the unit vector $\hat{u}$:

$$g_x = \hat{u} \cdot z_x \tag{8}$$

$$g_y = \hat{u} \cdot z_y \tag{9}$$

From this we obtain the dot product from o to each pixel with respect to our directional vector z_x and z_y. By multiplying g_x and g_y with gradient G, we obtain the gradient in x and y direction G_x and G_y.

The unit vector in the x direction is:

$$\begin{bmatrix}1\\0\end{bmatrix} \tag{10}$$

The unit vector $\hat{c}$ of c is:

$$\begin{bmatrix}1/\sqrt{2}\\-1/\sqrt{2}\end{bmatrix} \tag{11}$$

The dot product in the x direction is:

$$\begin{bmatrix}1\ 0\end{bmatrix} \cdot \begin{bmatrix}1/\sqrt{2}\\-1/\sqrt{2}\end{bmatrix} = 1/\sqrt{2} \tag{12}$$

We obtain the dot product g_x and g_y for every pixel, and multiply it with the gradient G at that pixel, obtaining:

$$G_x = \begin{bmatrix}-\frac{a-o}{2} & 0 & \frac{a-c}{2}\\ -(c-o) & 0 & f-o\\ -\frac{g-o}{2} & 0 & \frac{i-o}{2}\end{bmatrix} \quad G_y = \begin{bmatrix}-\frac{a-o}{2} & -(b-o) & -\frac{a-c}{2}\\ 0 & 0 & 0\\ \frac{g-o}{2} & h-o & \frac{i-o}{2}\end{bmatrix} \tag{13}$$

To extract the kernel from G_x and G_y, we simply divide each value by its difference. To simply our calculations, we multiply it by a constant factor of 2 to obtain D_x and D_y:

$$D_x = \begin{bmatrix}-1 & 0 & +1\\ -2 & 0 & +2\\ -1 & 0 & +1\end{bmatrix} \cdot Z \quad D_y = \begin{bmatrix}-1 & -2 & -1\\ 0 & 0 & 0\\ +1 & +2 & +1\end{bmatrix} \cdot Z \tag{14}$$

Where D_x and D_y represent the gradient approximation in the X and Y directions for a certain pixel. Z represents the depth map, which is convolved over by a 3x3 kernel depicted by the 3x3 matrices. This produces an output image of equal size to the depth map using 1x1 padding. To produce an estimate for the tangent at a certain pixel, we first produce the following vector:

$$T_i = \begin{bmatrix}-D_{xi}\\-D_{yi}\\D_z\end{bmatrix} \tag{15}$$

Where i represents a single pixel, T_i represents the tangent at pixel i, D_{xi} and D_{yi} represent the approximate gradient in x and y direction at pixel i. D_z is a global constant that represents the magnitude of the Z component, this effects how sensitive the final tangent is with respect to change in D_z and D_x. There is no specific value for D_z because the strength of both the normal map and the depth map which it derives from is arbitrary, thus D_z is also arbitrary. The higher D_z is, the less sensitive our tangent is to change and vice versa. For our model, we used a D_z of 0.8.

T_i can be considered the raw tangent of a given pixel, the final normal tangent map is obtained by calculating the unit vector of T_i. This is done by dividing each vector by its Euclidean norm:

$$\hat{T_i} = \frac{T_i}{|T_i|} = \frac{T_i}{\sqrt{{D_{xi}}^2 + {D_{yi}}^2 + {D_z}^2}} \tag{16}$$

Where $\hat{T_i}$ is the normalized tangent vector of pixel i.

The sobel filter can be expanded to larger kernel sizes. This is useful as larger sobel filters could detect gradient changes on larger scales, where smaller 3x3 filters are prone to local variations that do not contribute to useful feature extraction. By closely observing the calculations for a 3x3 sobel filter, we note that the dot product and gradient results in the inverse of the squared magnitude with a negative sign if it is against the gradient direction and vice versa. Thus, for an arbitrary kernel size of N:

$$\begin{bmatrix} A_{11} & A_{12} & A_{13} & \dots & A_{1N} \\ A_{21} & A_{22} & A_{23} & \dots & A_{2N} \\ A_{31} & A_{32} & A_{33} & \dots & A_{3N} \\ A_{41} & A_{42} & A_{43} & \dots & A_{4N} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ A_{N1} & A_{N2} & A_{N3} & \dots & A_{NN} \end{bmatrix} \tag{17}$$

Where for A_{yx} we denote j and i as its vector from the center o:

$$j = y - \frac{N-1}{2} - 1 \tag{18}$$

$$i = x - \frac{N-1}{2} - 1 \tag{19}$$

Using i and j, we can formulate a general equation for sobel filters of any size N:

$$Gx_{ij} = \frac{i}{i^2 + j^2} \tag{20}$$

$$Gy_{ij} = \frac{j}{i^2 + j^2} \tag{21}$$

Using this equation, we obtain for instance the default 5x5 Sobel filter in the X direction used for normal tangent calculations:

$$\begin{bmatrix} -5 & -4 & 0 & 4 & 5 \\ -8 & -10 & 0 & 10 & 8 \\ -10 & -20 & 0 & 20 & 10 \\ -8 & -10 & 0 & 10 & 8 \\ -5 & -4 & 0 & 4 & 5 \end{bmatrix} \tag{22}$$

It is multiplied by a constant factor of 20 to obtain integer values.

3.5 Depth-Aware Feature Backbone

Our backbone is based on ResNet50. It received with RGB input in the dimension of (N, C_{in}, H, W), or (1, 3, 2656, 800) by default. The monocular depth prediction module produces a tensor of size (1, 4, 2656, 800) when using a subregion D_s of 4.

We simply add the depth tensor as addition channels using the by concatenation, which results in a tensor of size (1, 7, 2800, 400) for RGB + depth. To accommodate for this increase in channel count, the first layer of our network is also modified. To best preserve the pre-trained weights, we preserve the weights for the original three channels, and repeat it for the subsequent channels. The same applies for the normal tangent input, which 3 channels.

Because objects can occur at various different sizes in instance segmentation, feature maps that are either too complex or simple may have negative effects on classification and segmentation The backbone thus generates feature maps at different dimensions as implemented in Mask R-CNN. Since we require depth as an explicit information in proceeding layers, a separate set of depth maps is generated alongside the feature maps, it is scaled using Bilinear interpolation.

3.6 ROI Pooling

Region of Interest Pooling is the extraction of the bounding box portion of the feature map for evaluation. Mask R-CNN used an heuristically defined equation to specify which feature map to use:

$$f(\alpha) = \left\lfloor 4 + log_2\left(\frac{\sqrt{\alpha}}{224}\right) \right\rfloor \tag{23}$$

Where α is the area of the proposal box, $f(\alpha)$ is the feature map index.

Since depth information is required for occlusion edge calculation on a pixel to pixel accuracy, depth maps are also scaled to the same resolutions as the feature maps, and pooled alongside feature maps.

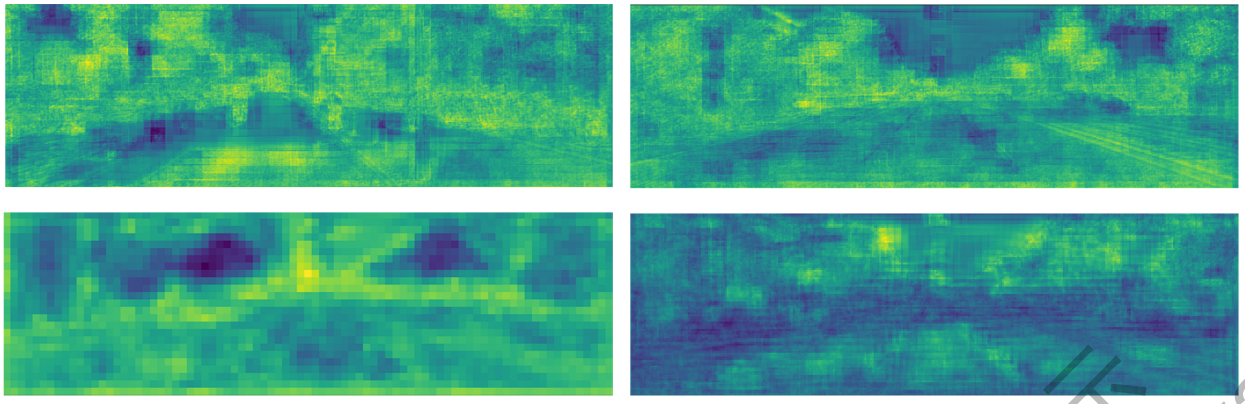

Fig. 9: The backbone extracts features from an image, usually hundreds of channels. Each channel is sensitive to one or multiple visual features. Multiple feature maps of various resolutions (extracted at different layers of the backbone) and channels are shown below.

3.7 Occlusion Edge Refinement[1]

Fig. 10: Ground truth (ideal) masks at occlusion edge. Occlusion edge is focused on inferring the unseen based on visible and depth information. The pooled depth data shows strong feature resemblance to the occlusion edge, it is easy to infer these regions are potentially occluded.

If the visible region of an object is touching another region that is in front of it, it is highly likely that the object is being occluded in this direction. The occlusion edge module predicts segments of a visible mask that is possibly being occluded instead of being its natural boundary. We use the intersecting edge between amodal and visible ground truths masks to generate the ground truth for occlusion edge. A Gaussian blur is applied over this edge for relative tolerance. Loss is calculated as:

$$\mathcal{L}_{eo} = \sum_{i=1}^{a}\sum_{j=1}^{a} \ell(\sigma(e_{ij}),\ \hat{e}_{ij}) \tag{24}$$

Where $\mathcal{L}_{oe}$ denotes the loss for occlusion edge, e_{ij} denotes the probability of pixel (i, j) being the edge of occlusion, $\hat{e}_{ij}$ denotes the ground truth of pixel (i, j) being on the edge of occlusion, α denotes the size of the pooled feature map, σ denotes the sigmoid activation function.

Notably, ℓ denotes the Binary Cross Entropy loss:

$$\ell(y, \hat{y}) \ = \ -\frac{1}{N}\sum_{i=1}^{N} y_i \ \cdot \ \log\ \hat{y}_i \ + \ (1 - y_i) \ \cdot \ \log\,(1 - \hat{y}_i) \tag{25}$$

[1]Occlusion Edge and Edge Occlusion can be used interchangeably, OE and EO stands for the same module.

The sigmoid activation function $\sigma(x)$ is:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \tag{26}$$

Where x is the predicted probability of a given value.

To use occlusion edge information for amodal mask prediction, we introduce Amodal Attention. Amodal Attention is the estimation of where the amodal mask might be, a separate mask head will use this attention information to complete the amodal mask. Amodal Attention is different from Occlusion Edge as it is already generated from other information, namely the visible mask and a rough prediction provided by Mask R-CNN's proposal network. Instead of overriding the amodal attention, we introduce an Occlusion Refinement layer, with loss $\mathcal{L}_{or}$ that uses the same loss function as $\mathcal{L}_{oe}$. This refinement layer will correct the current amodal attention using its occlusion information, regions further behind an object cannot be an object's occluded region, this can only be provided by the Occlusion Edge layer.

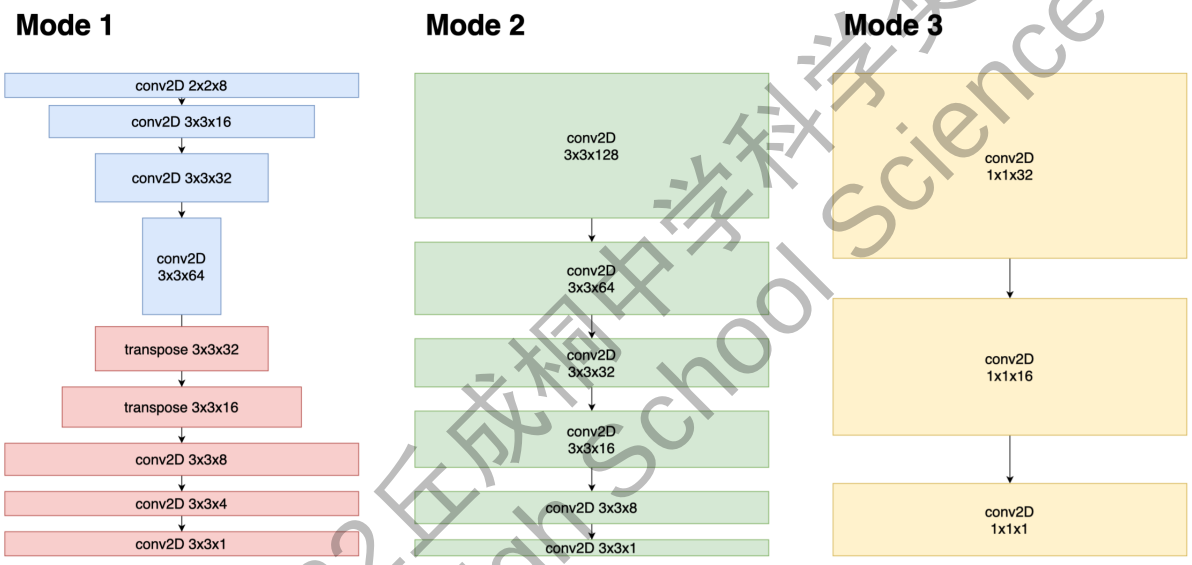

Fig. 11: Three different setups for the Edge Occlusion module.

We also had different implementations of the occlusion edge module as seen in Fig 11. The first implementation is an micro FCN network, this network deals normalized depth information as a classification and segmentation problem. The second implementation copies the structure of the amodal mask head, with deep convolutions to extract and use features from the normalized depth map for occlusion edge prediction. The third implementation is a simple 1x1 convolution that simply infers potential areas of occlusion based on the depth value on that pixel. The setup can be selected based on training configurations, the best will be chosen during the experimentation phase.

3.8 3D Shape Prior Refinement

The fundamentals of predicting the unseen region is based o n the seen regions and features to look for similar seen shapes, or shape prior for amodal completion. This is already implicitly achieved by the amodal prediction mask head. Weights are already trained to correspond visible features to amodal shapes, this is why the Mask-RCNN network can make reasonable amodal predictions without any modification - shape memorization is implicitly achieved during training. The shape prior refinement process, first introduced by Xiao at el. [6] further exploit this principle by explicitly storing and matching seen shapes for refinement.

The shape prior refinement process can be simplified into a problem of finding and matching different seen shapes. It is much easier to find similarities between shapes if it could be quantified. This is achieved through embeddings, which converts complex and high-dimensional data into simple, low-dimensional representations. This method is widely used in machine learning, because simpler data - like numbers and vectors, can be

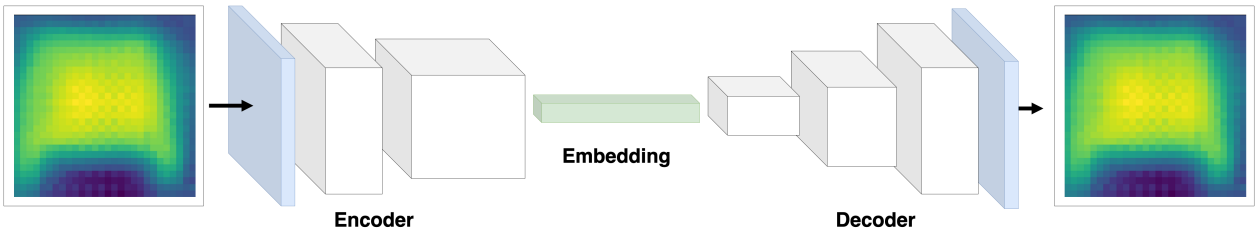

Fig. 12: The encoder-decoder network setup, embedding is obtained at the middle of the network. Training conducted by setting the loss as the difference between the output and the input, in which case they should be the same.

easily compared. It is very difficult conversely, to directly compare similarities between different shapes. A simple network consisting of multiple convolutional and deconvolutional (transposed convolutional) layers. The embedding is obtained in the middle between the convolutions (down-sampling) and the deconvolutions (up-sampling). By formulating a training method where given an input shape should result in the exact same shape on output, we are essentially training a network that is capable of encoding a low-dimensional representation of shapes, and conversely decode it back into the shapes. We can simply split the model into decoding and encoding functions for us to encode and decode embeddings.

We improve upon the shape prior refinement process by adding localized depth information as an additional input channel. By adding localized depth values into the encoding process, embeddings generated by the encoder effectively contains 3D surface features of the shape, which means that it now retains 3D memories of shapes. This is useful as it means that shape refinement based on shape prior can be significantly more specific, as a higher dimension allows for clearer distinction between different seen shapes. Shapes that appear to be similar in 2D might be vastly different when it is represented in 3D, which is something the original 2D shape refinement process is insensitive to.

Fig. 13: Embeddings of a batch of 3D shapes. Each embedding vector has a dimension of 1x392. The image shows a batch of embeddings with dimensions 32 x 392, which means it consists of 32 individual embeddings.

It is expensive to compare similarity against every single embedding, many shapes are also extremely similar. We employ K-means clustering to cluster similar embeddings. This way, we only have to compare an embedding against every cluster in the codebook to obtain the closest matching shape prior.

4 Experiment

The aim of our experiment is to evaluate and prove that incorporating depth information into amodal instance segmentation tasks can yield better results. Because we have introduced a multitude of ways in which depth information can be used, an ablation study will be carried out that evaluates the impacts of each module on performance and different per module configurations.

4.1 Datasets

KINS Dataset For training and evaluation of our model, we used the KINS dataset [23] which provides amodal and visible annotations of objects. The KINS dataset is based upon the KITTI dataset, which provides the KITTI Eigen Split depth dataset for which our monocular depth prediction module is trained upon. This ensure the quality of depth predictions while maintaining RGB as the only ground truth input for inference. There is a total of 7474 images for training and 7517 images for evaluation. Annotations in the KINS Dataset are stored in vector masks, it is rasterized into bitmasks prior to training. The KINS Dataset is widely used in amodal instance segmentation research, many models were trained and evaluated using KINS. This provides a common benchmark for comparasion between different methods.

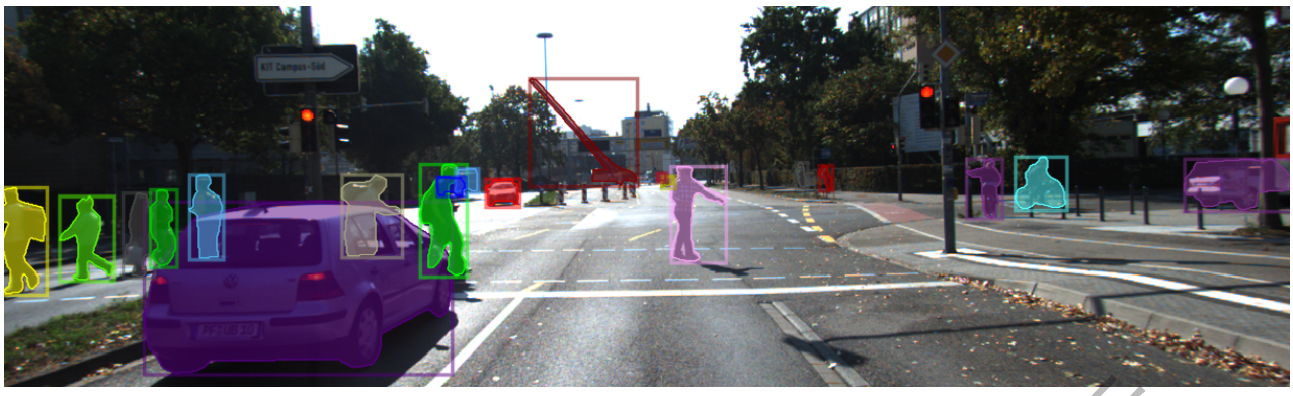

Fig. 14: A single image from the KINS dataset, complete with bounding boxes and amodal masks.

RGSD Dataset All synthetic data are captured in 1280 x 720 pixel resolution on an M1 Max using Unity Engine. The synthetic training dataset comprises of 10000 city scenes, 10000 street scenes and 20000 open road scenes. The validation dataset contains 1000 sets, where 400 are open road and 300 for both street and city. It includes the RGB image, the ground truth depth, amodal and visible masks of objects stored in a bitmask format.

4.2 Model Set-up

Depth Prediction Module Training of the Monocular Depth Prediction Module for KITTI is done using a single RTX3070ti with 46375 iterations with a batch size of 2 over 14 epochs. The synthetic dataset is trained over 19000 iterations with a batch size of 2 over 14 epochs. Both models used a learning rate of 0.0001 with 0.01 linear decay. To optimize training speed and memory usage, the depth prediction module is loaded in inference mode during amodal instance segmentation, it does not take part in training. Although depth information is inferred from RGB inputs, a single image always produces the same result. A cache mechanism is used to reload depth data from disk if it has been predicted before. This reduces training time from 10 hours to 3 hours, more than 3 times quicker.

Amodal Instance Segmentation The project is setup and implemented using a modified version of Detectron2 [24], which has a Mask-RCNN implemented using Pytorch and CUDA. Our project was trained on Ubuntu 20 using Python 3.6, Pytorch 1.4 and CUDA 10.1. Training is conducted on single Nvidia Tesla V100 instances. We used the ResNet-50 weights trained on ImageNet Classification as our pretrained weight for our feature extraction backbone, we also used weights from monocular depth prediction, as well as random weights. Depending on the configuration, data like depth and normal tangent will effect the input channel count for feature extraction. To preserve the pretrained weights, the first convolution layer will expand and repeat its three channel weights. For an input channel count of 9, the pretrained weights will be repeated three times. There are a total of 47979 iterations, the KINS dataset is further post-processed with random crop, rotation and color filters to enlarge the dataset.

Learning Rate A base learning rate of 0.0025 was chosen for training, this works particularly well with the pretrained ImageNet backbone. However, for randomly initialized weights, the loss becomes infinite or not defined (due to division by zero) after a few hundred to thousand iterations depending on the configured complexity. A lower learning rate solves this problem, but it also means the model might not be trained to its best performance. Learning rate warmup is a typical approach for solving high loss training. Constant warmup specifies a lower learning rate for a given number of iterations. Linear warmup increases the learning rate linearly from zero for a given number of iterations. We observe that during training, both constant and linear warmup methods are not suitable for our scenario; Loss either converge too slow, or the initial learning rate remains too high. We introduce a smooth warmup method, which increases the learning rate at a given power:

$$lr_w = \left(\frac{i}{t}\right)^p \times f \times lr_f \tag{27}$$

Where lr_w is the warmup learning rate, lr_f is the final learning rate, i is the current iteration, t is the warmup iteration count, p is the warmup power, and f is the constant warmup factor.

For random weights, we used $lr_f = 1.0$, $p = 1.3$, $t = 15000$.

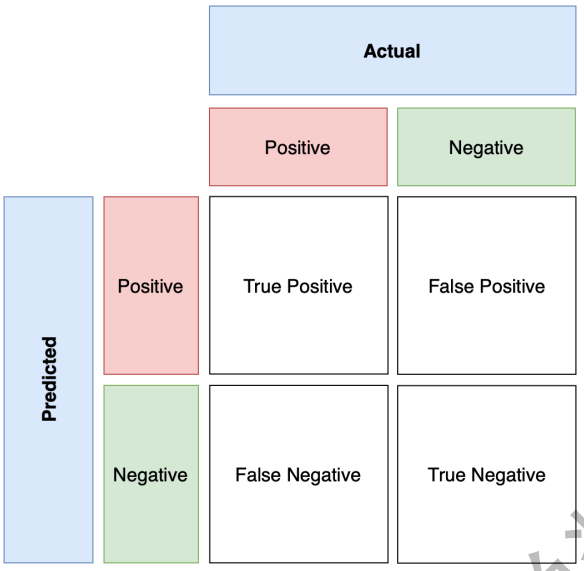

Fig. 15: Confusion Matrix

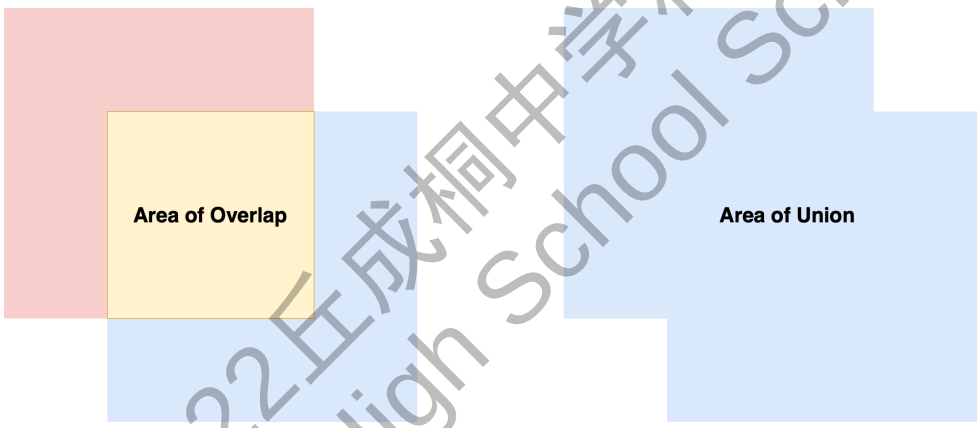

Fig. 16: Area of Overlap vs Area of Union

Post Processing Post processing of training data is commonly used to increase volume of the training set without actually acquiring additional data. The KINS dataset provides 7000 images for training, which is significantly lower than datasets used for instance segmentation and object detection for a task that is overall more difficult. A post processing layer creates variants of the same image, such as flipping the image horizontally, changing the contrast, and randomly crop the image. Post Processing is however computationally expensive, it drastically increases training time. For the ablation study, post processing will be disabled for all tests. We notice a 15% drop in performance with post processing disabled.

4.3 Metrics

Average Precision, or AP, is a method for evaluating performance and accuracy of object detection and segmentation problems. For our evaluation, we used Average Precision, Mean Average Precision and Average Recall to evaluate performance of different models.

The confusion matrix maps out all possible combinations of results as shown in Fig 15. If a result is predicted, it is considered positive; If a result is not predicted, it is considered negative. If the prediction is correct, it is considered true, otherwise it is considered false. Thus, we have True Positive TP, False Positive FP, True Negative TN and False Negative FN.

An important metric in determining Bounding Box precision is the Intersection over Union, also known as IoU:

$$IoU = \frac{|A \cap B|}{|A \cup B|} \tag{28}$$

Where A, B describes the predicted and ground truth bounding box.

IoU is used to determine whether a prediction is True or False. This depends on an IoU threshold. If the predicted IoU is above this threshold, the result is considered true. Using the confusion matrix, we can determine Precision and Recall:

$$Precision = \frac{TP}{TP + FP} \quad Recall = \frac{TP}{TP + FN} \tag{29}$$

Precision depicts how many predictions are correct across all predictions. It evaluates how many correct predictions are made among all predictions. Recall depicts how many correct predictions are made. It evaluates how many correct predictions are made among all correct results.

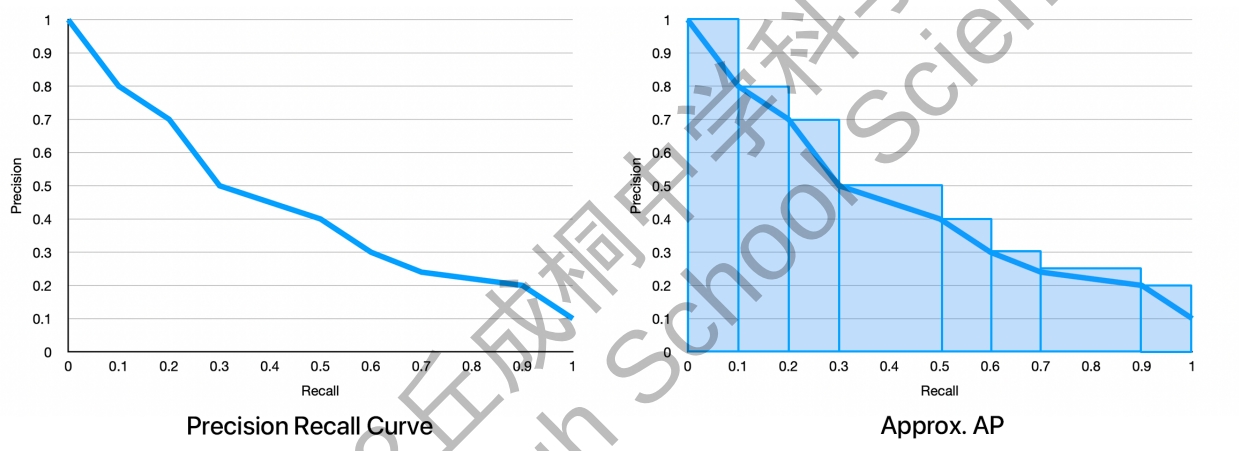

Fig. 17: Average Precision is the area under a Precision Recall Curve, which is precision over recall.

We use the area under the precision recall curve to calculate Average Precision. However, the precision recall curve is not a continuous function, so we can only obtain approximations of the area. One method is to use rectangles to estimate area under the curve as shown in Fig 17.

Calculation of Average Precision is usually for a single class of objects. For the representation of all classes, we use Mean Average Precision (mAP):

$$mAP = \frac{1}{N} \sum_{i}^{N} AP_i \tag{30}$$

Where N is the number of classes, AP_i is the average precision of a single class.

Similarly, Average Recall (AR) is twice the area under a Recall over IoU curve. Similar approximation methods are used to calculate the area.

For both AR and AP, the number behind states the IoU threshold for determining true and false results. AR_{50} represents an average recall with 0.5 IoU threshold or 50%.

Comparisons To further evaluate our network's performance, we will compare our best setup with other commonly used methods such as Mask R-CNN and PANet. All networks used the KINS dataset for evaluation. Post processing will be applied to our final setup during training.

5 Evaluation and Ablation Study

Fig. 18: Amodal instance segmentation results from ImageNet weights using RGB and occlusion edge refinement.

We conducted an ablation study on our network by testing out different setups and combinations of features. The aim of our study is to evaluate how depth information improves amodal instance segmentation performance, and what are the most effective methods of using the depth information. The complete table of results is shown in Table 2. R, D, N O and S stands for RGB, Depth, Normalm Occlusion Edge and 3D shape prior respectively. A checkmark ✓symbol denote this component is enabled for the given row, other special characters represent the use of a variation of this component. Depth with N for example, means depth without normalization. Abbreviations will also be used to describe configurations, the order of components follows that of the table. For instance, random weights with RGB, depth without normalization, and normal tangent with medium sized filter is described as Random R+DN+NM.

5.1 Backbone

Our feature extraction backbone used the ResNet50 architecture, by using a standard backbone, we can use transfer learning to improve our network's performance. A separate and complete ablation study is done for both weights. We are using random weights because we want to evaluate the pure performance improvements from using depth information. A backbone with pretrained weights that are obtained from other tasks may effect the results of our study negatively. Conversely, most amodal instance segmentation networks utilize a pretrained backbone feature extractor, which results in significant performance boost from using random weights. We will evaluate how well depth information copes with pretrained weights that is not depth aware.

As seen in Table 2, there is around a 30% improvement in accuracy from random weights to ImageNet weights. Imagenet R+N has an AP of 22.83, while Random R+N has an AP of 16.52 only. This demonstrates that even for weights that were not trained on depth information still yields better results than random weights.

5.2 Depth

The most basic approach towards incorporating depth in amodal instance segmentation is by feeding it through the feature extraction backbone. By doing this, the entire network will conduct instance segmentation and amodal segmentation tasks using depth-aware feature maps.

Table 1: Results on Amodal Instance Segmentation
R D N O S stands for RGB, Depth, Normal, Occlusion Edge and 3D Shape Prior
N stands for raw depth without normalization; S1 stands for regional normalization with subregion of 1; O1 stands for regional normalization with offset interval of 1; M stands for medium sized sobel filter; L stands for large sized sobel filter; V stands for visible only depth

Weights	R	D	O	R	S	*AP*	*AP*50	*AP*75	AP_s	AP_m	AP_l	*AR*1	*AR*10	*AR*100
Random	✓					13.20	28.02	10.20	18.21	17.86	14.48	10.51	25.10	28.76
Random	✓	✓				16.52	31.49	15.60	22.26	21.28	16.87	11.91	27.72	30.14
Random	✓	N				16.42	31.64	15.50	21.68	20.52	17.22	11.96	27.72	30.14
Random		✓				10.84	20.71	10.52	14.00	13.09	11.68	7.98	18.32	20.05
Random	✓		✓			9.55	20.10	8.02	12.20	11.74	10.08	6.72	16.98	19.03
Random			✓			11.26	21.38	10.83	14.66	13.45	11.74	8.39	18.89	20.52
Random	✓	✓	✓			16.13	30.71	15.26	21.56	20.52	17.59	11.16	26.77	29.21
Random	✓			✓		**17.19**	32.62	16.38	23.02	21.94	18.09	12.38	28.29	30.61
Random	✓	✓		✓		17.17	23.57	16.30	23.58	22.44	18.15	12.68	28.59	30.91
Random	✓		✓	✓		9.74	20.33	8.38	12.53	12.25	10.41	6.65	17.50	19.66
Random	✓	✓	✓	✓		10.11	21.34	8.50	12.99	12.37	10.61	7.30	18.46	20.59
ImageNet	✓	✓				22.83	40.77	22.97	28.54	26.53	22.17	15.94	33.41	35.36
ImageNet	✓	O1				22.19	39.90	22.26	27.97	26.22	22.36	15.76	32.67	34.64
ImageNet	✓	S1				22.35	39.93	22.49	28.20	26.39	22.42	16.00	33.04	33.04
ImageNet	✓	N				22.57	40.47	22.51	28.28	26.52	21.82	15.69	32.79	34.86
ImageNet	✓		✓			24.37	44.02	24.48	30.05	29.08	24.45	16.77	35.56	37.98
ImageNet	✓		M			17.35	32.55	16.32	22.32	20.82	16.96	12.84	26.76	28.50
ImageNet	✓		L			16.25	31.04	15.12	21.19	19.79	16.27	12.11	25.30	27.05
ImageNet	✓				✓	24.62	45.63	24.13	30.41	27.67	23.83	16.84	32.60	36.63
ImageNet	✓				V	24.01	45.58	22.92	30.14	28.49	23.58	16.88	32.55	36.62
ImageNet	✓	✓	✓			23.11	41.22	23.44	29.11	27.36	22.74	16.39	33.90	35.93
ImageNet	✓			✓		**25.41**	44.98	25.64	30.73	29.27	24.59	17.45	35.81	37.89
ImageNet	✓	✓		✓		23.00	40.80	23.48	28.59	27.02	22.60	16.23	33.65	28.15
ImageNet	✓		✓	✓		17.13	32.32	16.05	22.26	20.67	16.93	12.46	26.33	28.15
ImageNet	✓	✓	✓	✓		16.93	31.72	16.16	22.27	2.23	17.53	12.73	26.64	28.43

Usefulness By concatenating normalized depth information consisting of 4 channels by default to the RGB input, there is a 25% increase in AP from Random R to Random R+D. This demonstrates that using depth information does result in significant improvements for amodal instance segmentation. To further examine whether depth information can actually provide features for amodal instance segmentation, the network is trained on depth without using any RGB information. Random D is 18% lower in AP than Random R, and 35% lower in AP than Random R+D. An AP of 10.84 means that the network is still able to infer reasonable results soley from depth data. We can also conclude that depth information cannot substitute RGB data, which is logical as RGB information provides surface features that depth is insensitive to.

Impact of Regional Depth Normalization As mentioned in the methods section, a regional depth normalization pass was done over the depth information before feeding it to the feature extractor. To evaluate the impact of said normalization, we tested the model without normalization, denoted by N. For both Random R+DN and ImageNet R+DN, there was only marginal improvements in AP when using regional normalization, still this means that the model performed worse without regional normalization. Evaluation on the impact of step size and subregion intervals were also conducted. ImageNet R+D1, which represents regional depth normalization with a subregion count of 1, essentially not creating subregions. This results in a small 1% decrease in AP and other metrics. ImageNet R+DS1 represents normalization with a step size of 1, which is 2/3 less than the default step size of 3. This means depth information is far more localized and distance specific. As hypothesized in the method section, a smaller step size does lead to a drop in AP of 2%, likely caused by over-fitted local variations.

5.3 Normal Tangent

Depth information can be further processed into normal tangent maps. Normal tangent maps describe the surface tangent at a given pixel. This is different from depth information as normal tangent maps provide an explicit representation of 3D surface features, while depth only provide it implicitly to the network.

Usefulness We have to first evaluate if normal tangent information can be perceived well by the feature extraction backbone. Normal tangent information is different from depth as depth information can be represented well within in one channel. Normal tangents are represented in 3 channels, representing the x, y and z components of the unit vector at a given pixel. There is no implicit separation between RGB and xyz channels, which might also negatively effect results. ImageNet R+N yields an AP of 24.37, which is 6% higher than that of ImageNet R+D. This shows that normal information is more useful than depth information for amodal instance segmentation.

Accuracy and Sensitivity The normal map deduced from the depth information is only an estimate in surface tangents, it is not the physical geometrical tangents. Since the normal map information is not completely accurate, how well do different parameters effect estimations is also important. As described in the method section, different kernel sizes for the sobel filter effect how large of a gradient the normal map is sensitive to. Larger kernel sizes reflect a broader change in gradient, at the cost of less details. We experimented with ImageNet R+NM, ImageNet R+NL, which stands for medium and large kernel sizes respectively. The medium kernel has a size of 21x21, while the large kernel has a size of 71x71. Results show that there is a 29% drop in AP for the medium filter, and a further 34% drop in AP for the large filter. This means that the improvements in overall tangent accuracy does not compensate for the loss in detail. Furthermore, the backbone network ResNet50 uses convolutional kernels larger than 1x1, which increases the perceptive field of a given pixel. This may implicitly create a representation of more comprehensive gradient changes without sacrificing detail.

5.4 Occlusion Edge Refinement

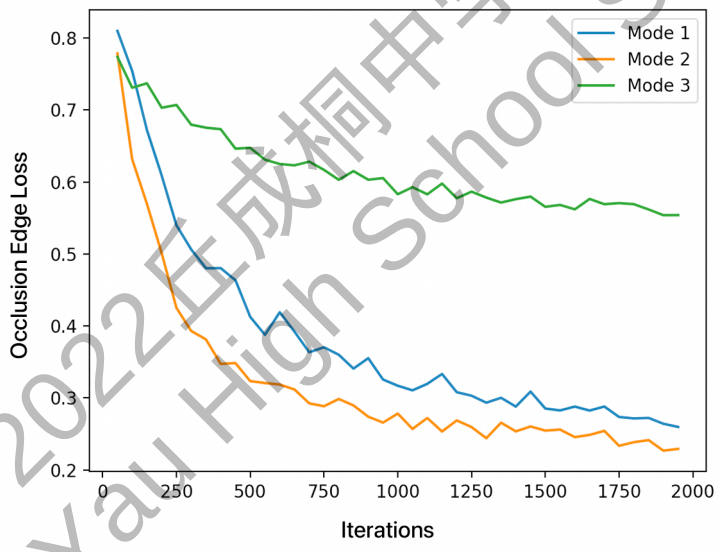

Fig. 19: Loss of different network setups for EO Refinement Layer

Occlusion edge refinement differs from other components in that it does not affect the overall network in feature extraction, object detection and visible segmentation. As described in the method section, there were multiple network setups for the occlusion edge refinement module. The networks were first evaluated with how well loss decreased before using the best network for amodal instance segmentation.

Network Setup Results indicate that Mode 2, which mimics the network setup of the amodal mask head, reaches the lowest loss using the least time. The eo loss of different modes over iteration is shown in Fig.19. Mode 3 yields the highest loss. Although it is hypothesized that Mode 1 which uses a micro FCN implementation should perform best, this is now the case. We believe this is due to either insufficient information, over-complicated network or improper configurations. Mode 3 performs the worst because the network is too simple, it only infers occlusion edge based on local depth values, and do not take in account of more general depth patterns.

Usefulness The EO refinement layer has the single most improvement to amodal instance segmentation compared to other configurations. Random R+O is 30% higher in AP than Random R. ImageNet R+O is 11%

Table 2: Results on KINS dataset, average precision of amodal segmentation.
We compared our model's results against other methods.
The DAISnet result was trained using ImageNet + RGB + Edge Refinement + Post Processing

Model	*AP*
Mask R-CNN [25]	24.93
PANet [23]	27.39
BCNet [26]	28.87
DAISnet (Ours)	**31.48**

higher in AP than ImageNet R+D. We believe this is because the EO refinement layer has the most explicit use of depth information. Since EO refinement occurs after visible mask segmentation, it can use the visible regions of an object to conduct per object depth normalization. This is extremely useful, as any region with depth value above that of the visible maximum is guaranteed to not to be occluded, and any region with depth value lower than that of the visible minimum is likely to be occluded. Even an algorithm can infer plausible results using this information.

5.5 3D Shape Prior Refinement

The 3D shape prior refinement is denoted by S in Table 2. 3D Shape Prior Refinement combines depth information during embedding encoding and decoding. The original shape prior implementation [6] used coarse amodal masks from Mask-RCNN to conduct shape prior matching. Depth information however, is not entirely reflective of the actual object because it contains occluded regions. Depth of occluded regions of an object are not the actual depth of these objects. We differentiate these depths by denoting depth in visible regions as object-only depth. ImageNet R+S results in the second highest AP after Occlusion Edge Refinement. There is a 7% improvement in AP from ImageNet R+D baseline. ImageNet R+SV, which only uses the object only depth, yields an AP of 24.01, which is 2% less than using all the depth. We can conclude that depth in occluded regions does in fact contribute to 3D shape prior refinement. This might be because the shape prior refinement also implicitly memorize occlusion scenarios. If the same region was occluded and has a higher depth value, it is reasonable to assume the same for similar cases.

5.6 Combinations

We also tested different combinations of components to see if there is additional performance improvements, such as using depth and normal information together. It is logical to assume that if both components improved performance, then the combination of said components should further improve performance. But according to our experimental results, this is not the case. We first evaluate the performance of combining depth and normal information. Random R+D+N is 3% lower in AP than Random R+D. ImageNet R+D+N is 1% better than ImageNet R+D, but 5% lower than ImageNet R+N. This shows that there is no guaranteed improvement when adding depth and normal data together. We conclude that this is because normal tangents already provide the same 3D features that depth provide, and in a more explicit way without the issues distance variation poses.

We also experimented on combining depth-aware feature extraction and occlusion edge refinement. There was marginal difference in adding depth features to occlusion edge, Random R+D+O is only 0.2% less in AP than Random R+O, while ImageNet R+D+O is 10% less in AP than ImageNet R+O. We believe this is due to potentially similar roles both components play. The depth features extracted from R+D is implicitly used for amodal attention. Since the occlusion edge refinement layer uses a similar architecture to the amodal mask head, it is likely that both modules have used depth in a similar way to predict the amodal mask. However, since the occlusion edge refinement layer uses per object normalization, it should perform better, which explains why the occlusion edge component outperforms depth feature extraction. The use of 3D shape prior and occlusion edge proves that explicit 3D features at region of interest levels improve amodal segmentation results. Implicit uses of 3D features may require further investigation for improvements in results.

5.7 Limitation and Future Works

Evaluation results show that although DAISnet has utilized depth-aware 3D features to improve amodal instance segmentation performance, there are still some limitations. Computed tangent normals are not the actual geometric normals of objects, a method of predicting geometric normals using neural network should be made. Depth-based 3D features can be further incorporates into the model by combining the monocular depth prediction module with DAISnet, which provides additional 3D-aware features during monocular depth prediction

that might be useful. Current region of interest and mask logit resolutions remain low at 14x14, optimization techniques on increasing RoI resolution is worth investigating.

6 Conclusion

We proposed a novel depth-aware amodal instance segmentation network (DAISnet) that used depth information to infer 3D features and occlusion information. This model contains regional depth normalization, normal tangent computation and occlusion edge refinement methods that implicitly and explicitly utilize the depth information. We also developed a method of obtaining a synthetic dataset that contains depth and amodal mask annotations to address the difficulty of acquiring depth and amodal ground truths. Our experiments demonstrated that depth information and subsequent 3D features extracted implicitly and explicitly brings notable improvements to amodal instance segmentation performance.

7 References

[1] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Proc. NIPS*, 2015.

[2] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *Proc. ICCV*, 2017.

[3] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. Hybrid task cascade for instance segmentation. In *Proc. CVPR*, 2019.

[4] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: delving into high quality object detection. In *Proc. CVPR*, 2018.

[5] Hao Ding, Siyuan Qiao, Alan L. Yuille, and Wei Shen. Deeply shape-guided cascade for instance segmentation. In *Proc. CVPR*, 2021.

[6] Yuting Xiao, Yanyu Xu, Ziming Zhong, Weixin Luo, Jiawei Li, and Shenghua Gao. Amodal segmentation based on visible region segmentation and shape prior. In *Proc. AAAI*, 2021.

[7] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. 2013.

[8] Ross Girshick. Fast r-cnn. 2015.

[9] Fully convolutional networks for semantic segmentation. 2014.

[10] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring R-CNN. In *Proc. CVPR*, 2019.

[11] Clément Godard, Oisin Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proc. CVPR*, 2017.

[12] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. 2019.

[13] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proc. CVPR 2018*, 2018.

[14] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.

[15] Patrick Follmann, Rebecca König, Philipp Härtinger, Michael Klostermann, and Tobias Böttger. Learning to see the invisible: End-to-end trainable amodal instance segmentation. In *Proc. WACV*, 2019.

[16] Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with KINS dataset. In *Proc. CVPR*, 2019.

[17] Xunli Zeng and Jianqin Yin. Amodal segmentation just like doing a jigsaw. *CoRR*, abs/2107.07464, 2021.

[18] Ke Li and Jitendra Malik. Amodal instance segmentation. In *Proc. ECCV*, 2016.

[19] Ziheng Zhang, Anpei Chen, Ling Xie, Jingyi Yu, and Shenghua Gao. Learning semantics-aware distance map with semantics layering network for amodal instance segmentation. In *Proceedings of the 27th ACM International Conference on Multimedia*, 2019.

[20] Yanfeng Liu, Eric Psota, and Lance C. Pérez. Layered embeddings for amodal instance segmentation. *CoRR*, abs/2002.06264, 2020.

[21] Jianwei Yang, Zhile Ren, Mingze Xu, Xinlei Chen, David J. Crandall, Devi Parikh, and Dhruv Batra. Embodied amodal recognition: Learning to move to perceive objects. In *Proc. ICCV*, 2019.

[22] Zhuo Deng and Longin Jan Latecki. Amodal detection of 3d objects: Inferring 3d bounding boxes from 2d ones in rgb-depth images. In *Proc. CVPR*, 2017.

[23] Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with kins dataset. 2019.

[24] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. https://github.com/facebookresearch/detectron2, 2019.

[25] Patrick Follmann, Rebecca König, Philipp Härtinger, Michael Klostermann, and Tobias Böttger. Learning to see the invisible: End-to-end trainable amodal instance segmentation. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7-11, 2019*, pages 1328–1336. IEEE, 2019.

[26] Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Deep occlusion-aware instance segmentation with overlapping bilayers. In *CVPR*, 2021.

8 Acknowledgement

The utilization of depth information in amodal segmentation tasks drew a lot of inspirations and ideas from my past experiences in computer graphics. I developed a hobby in many areas of computer science from web development to robotics. The area that interested me the most was game development. I have been developing video games since I was Grade 7. It was fascinating because you could simulate and create an entire world of color and interactions. Throughout this time, I learnt a lot of spacial calculations in 3D and computer graphics. I taught myself how to write shaders, process mesh data, and produce stunning visuals for the final product. It was because of this experience that inspired me to incorporate depth information to amodal segmentation.

I have achieved a lot throughout my exploration in computer science, from being on USACO Gold Division to running my own CS club in school that developed applications used by thousands of students. The only field I that never attempted for was machine learning beyond following what sporadic news of new advancements. Whether it is because of my lack of guidance or belief that it is too difficult for a high school student. When I set out my mind to conduct a research project, I was finally determined to research into AI and step out of my comfort zone. Under the guidance and help provided by Mr. Ding, I finally understood how deep learning worked. When I saw the gradual decrease in loss during training for one of my starter networks, I felt very accomplished because not only did I manage to make it but also understood how and why it worked.

Amodal instance segmentation is a task that has not yet being explored much, and it is also a computer vision task - which interests me due to my past experiences in computer graphics. It was not long after I began my research did I think about using depths in amodal instance segmentation. Depth information is vital to 3D computer graphics, it provides the fundamental information to calculation occlusion (z-buffer and z-write) and distance based visual effects. Thus I began developing the RGSD synthetic database, which used the same 3D engine that I developed games using. DAISnet uses normal tangent maps, which is also an important component in 3D graphics. Normal maps represent the surface tangents of geometries using pixels without actually defining mesh vertex. It is much easier to create normal tangent maps than create actual geometry, it also significantly more efficient. I am very proud that I was able to utilize past experiences in game development to help with my research project on a relatively new field to me - and achieve great improvements in results.

I want to thank Dr. Hao Ding from Johns Hopkins University for providing guidance and feedback on my research. He has given feedback and guidance on the validation of my ideas and its implementations. I am grateful to finally gain a thorough knowledge in deep learning and neural network.

I want to thank the support provided by the computer science department in my school, namely Mr. Shawn Wu who has supported my research on behalf of my school. I want to thank my parents who has made it possible for me to access a dozen Nvidia Tesla V100 instances through Google Cloud Platform, I would not have been able to iterate and complete such a thorough ablation study in time without it.