

Geometry-Aware Instance Segmentation with Disparity Maps

Cho-Ying Wu¹ Xiaoyan Hu² Michael Happold² Qiangeng Xu¹ Ulrich Neumann¹

¹University of Southern California ²Argo AI

{choyingw, qiangenx, uneumann}@usc.edu {xiaoyan, mhappold}@argo.ai

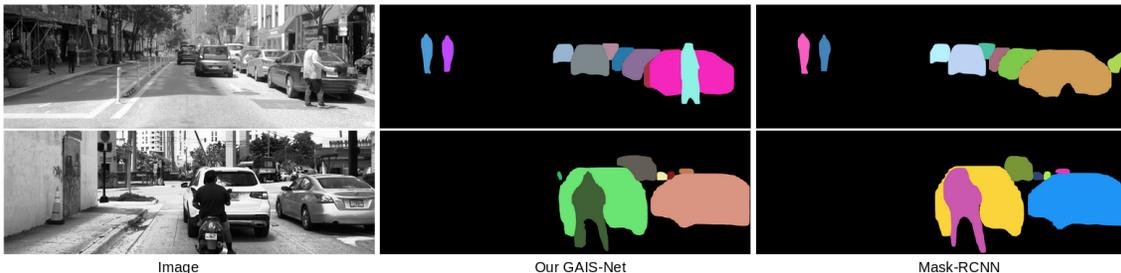


Figure 1. **GAIS-Net results on HQDS dataset.** Left column shows stereo left images with histogram equalization to enhance contrast for better visualization. Middle and right column show Mask-RCNN and GAIS-Net results, respectively. Each instance has different colors. With the aid of geometric information, GAIS-Net can segment out the person from the overlapping area in the first row example. In the second row scenario, Mask-RCNN generates distorted mask for the smoking motorcyclist because of cigarette plume and in contrast GAIS-Net displays a more robust shape control capability.

Abstract

Most previous works of outdoor instance segmentation for images only use color information. We explore a novel direction of sensor fusion to exploit stereo cameras. Geometric information from disparities helps separate overlapping objects of the same or different classes. Moreover, geometric information penalizes region proposals with unlikely 3D shapes thus suppressing false positive detections. Mask regression is based on 2D, 2.5D, and 3D ROI using the pseudo-lidar and image-based representations. These mask predictions are fused by a mask scoring process. However, public datasets only adopt stereo systems with shorter baseline and focal length, which limit measuring ranges of stereo cameras. We collect and utilize High-Quality Driving Stereo (HQDS) dataset, using much longer baseline and focal length with higher resolution. Our performance attains state of the art. Please refer to our [project page](#) for codes and data. The full paper is available [here](#).

1. Introduction

Instance segmentation, which segments out every object of interest, is an elemental task for computer vision. It is crucial for autonomous driving because it is vital to know positions for every object instance on roads. In the context of instance segmentation on images, previous approaches only operate on RGB imagery, such as Mask-RCNN [3].

However, image data could be affected by illumination, color change, shadows, or optical defects. These factors can degrade the performance of image-based instance segmentation. By utilizing another modality that provides geometric cues of scenes, and since object shapes are independent of object texture and color change, these strong priors add more robust information of the scenes. A prior work [7] that goes beyond the dominant paradigm to incorporate depth information only uses it for naive ordering rather than directly regressing masks or building an end-to-end trainable model to propagate depth information. Besides, their depth maps are predicted from monocular images, making the depth ordering unreliable.

In outdoor scenes, stereo cameras or lidar sensors are commonly used for depth acquisition. Stereo cameras are low-cost and their adjustable parameters, such as *longer baselines* (b) and *focal lengths* (f), favor stereo matching at far fields. Relationship of depth and disparity is given by

$$\text{depth} = \frac{f \times b}{\text{disparity}}. \quad (1)$$

1-disparity (the minimal pixel difference showing the ideal longest range a stereo system could detect) represents farther distance if using longer f and b . Next, longer baselines and focal lengths favor more precise geometric estimations [5], since longer baselines produce smaller triangulation error, and longer focal lengths project objects on

images with more pixels and thus enhance the robustness of stereo matching and show more complete shapes.

In this paper, we propose Geometry-Aware Instance Segmentation Network (GAIS-Net) that takes the advantages of both the semantic information from image domain and geometric information from disparity maps. Our contributions are summarized as follows:

1. To our knowledge, we are the first to perform instance segmentation on imagery by fusing images and disparity information to regress object masks.

2. We collect High-Quality Driving Stereo (HQDS) dataset, with a total of 8.8K stereo pairs and with $f \times b$ **4 times larger** than the current best dataset, Cityscapes.

3. We present GAIS-Net, an aggregation of representation design for instance segmentation using images, image-based, and point cloud-based networks. We train GAIS-Net with different losses, and fuse these predictions using the mask scoring. GAIS-Net achieves the state of the art.

2. Method

Our goal is to construct an end-to-end trainable network to perform instance segmentation for autonomous driving. Our system segments out each instance and outputs confidence scores for bounding boxes and masks for each instance. To exploit geometric information, we adopt PSM-Net [1], the state-of-the-art stereo matching network, and introduce disparity information at ROI heads. The whole network design is in Fig. 2.

We build a two-stage detector with a backbone network, such as ResNet50-FPN, and a region proposal network (RPN) with non-maximum suppression. Object proposals are collected by feeding a stereo left image into the backbone network and RPN. The same as Mask-RCNN, we perform bounding box regression, class prediction for proposals, and mask prediction based on image domain features. Corresponding losses are denoted as \mathcal{L}_{box} , \mathcal{L}_{cls} , and \mathcal{L}_{2Dmask} , and are identified in [3].

2.1. Geometry-Aware Mask Prediction

2.5D ROI and 3D ROI. We use PSMNet [1] and stereo pairs to predict dense disparity maps, projected onto the left stereo frame. Next, RPN outputs region proposals. We collect proposals and crop out these areas from the disparity map. We call these cropped out disparity areas as *2.5D ROI*.

Based on the observations from pseudo-lidar work [6], which describes the advantage of back-projecting 2D grid structured data into 3D point cloud and processing with point cloud networks, we back-project the disparity map into \mathbb{R}^3 space, where for each point, the first and second components describe its 2D grid coordinates, and the third component stores its disparity value. We name this representation as *3D ROI*.

Instance Segmentation Networks. Each 3D ROI contains different number of points. To facilitate training, we uniformly sample the 3D ROI to 1024 points, and collect all the 3D ROI into a tensor. We develop a PointNet structured instance segmentation network to extract point features and perform per-point mask probability prediction. We re-project the 3D feature onto the 2D grid to calculate the mask prediction and its loss \mathcal{L}_{3Dmask} . The re-projection is efficient because we do not break the point order in the point cloud-based instance segmentation. \mathcal{L}_{3Dmask} , same as \mathcal{L}_{2Dmask} , is a cross-entropy loss between a predicted probability mask and its matched groundtruth.

To fully utilize advantages of different representations, we further do 2.5D ROI instance segmentation with an image-based CNN. Similar to instance segmentation on 2D ROI, this network extracts local features of 2.5D ROI, and later performs per-pixel mask probability prediction. The mask prediction loss is denoted as $\mathcal{L}_{2.5Dmask}$.

2.2. Mask Continuity

We sample 3D ROI to 1024 points uniformly. However, predicted masks, denoted as M_{3D} , and their outlines are sensitive to pseudo-lidar sampling strategies. An undesirable sampling is illustrated in Fig. 3. To compensate the undesirable effect, we introduce a mask continuity loss. Since objects are structured and continuous, we calculate a *mask Laplacian* as $\nabla^2 M = \frac{\partial^2 M}{\partial x^2} + \frac{\partial^2 M}{\partial y^2}$, where x and y denote the dimensions of M . Mask Laplacian computes continuity of M . Further, the mask continuity loss is calculated as $\mathcal{L}_{cont} = \|\nabla^2 M\|^2$ for penalizing discontinuities of M .

2.3. Representation Correspondence

We use the point cloud-based network and the image-based network to extract features and regress M_{3D} and $M_{2.5D}$. These two masks should be similar because they are from the same disparity map. To evaluate the similarity, cross-entropy is calculated between M_{3D} and $M_{2.5D}$, and serves as a *self-supervised* correspondence loss \mathcal{L}_{corr} . Minimizing this term lets the networks of different representations supervise each other to extract more descriptive features for mask regressing, resulting in similar probability distribution between $M_{2.5D}$ and M_{3D} . Mask-RCNN uses a 14×14 feature grid after ROI pooling to regress masks. We also use this size at the mask heads.

2.4. Mask Scores and Mask Fusion

MS-RCNN [4] introduces mask scoring to directly regresses MaskIoU score based on a predicted mask and its associated matched groundtruth, showing quality of the mask prediction. However, their scores are not adopted at inference time to help manipulate mask shapes.

We adopt mask scoring and further exploit MaskIoU scores to fuse mask predictions from different represen-

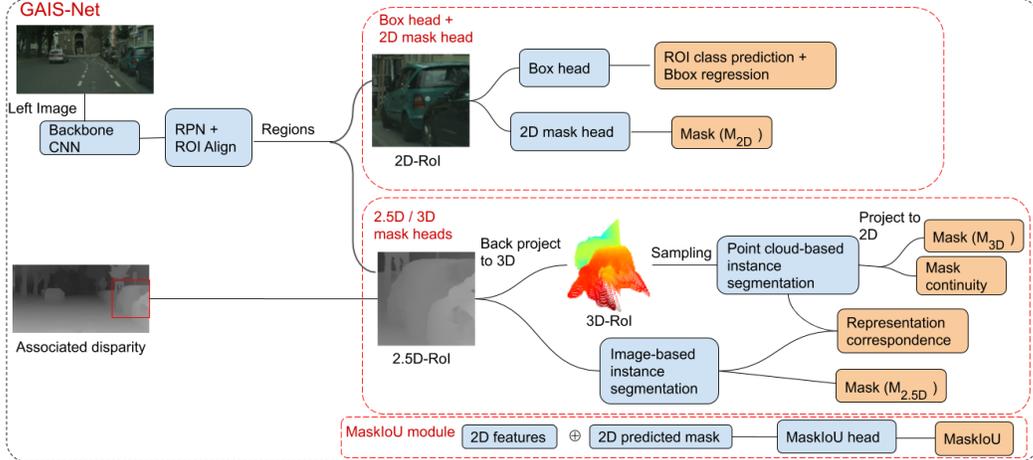


Figure 2. **Network design of our GAIS-Net.** Bbox is for bounding box. We color modules in blue and outputs or loss parts in orange. In the MaskIoU module, the 2D features and 2D predicted mask are from the 2D mask head. They are fed into MaskIoU head to regress MaskIoU scores. We draw the MaskIoU head separately for clear visualization. \oplus stands for concatenation.

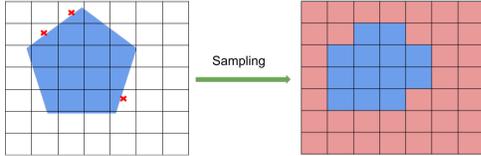


Figure 3. **Undesirable sampling example.** The blue areas represent foreground. Suppose we uniformly sample every grid center point in the left figure, resulting in the point cloud showing in the occupancy grid on the right. Red crosses are undesirable sampling points, which just lie outside the foreground object, making the shape after sampling different from the original one.

tations at the inference time. The mask fusion process is illustrated in Fig. 4. During the inference time, we concatenate features and predicted masks of different representations respectively as inputs to the MaskIoU head. Scores of s_{2D} , $s_{2.5D}$, and s_{3D} are outputs from the MaskIoU head. We fuse mask predictions using their corresponding mask scores. We first linearly combine $(M_{2.5D}, s_{2.5D})$ and (M_{3D}, s_{3D}) to obtain (M_D, s_D) for the disparity. The formulation is as follows.

$$M_D = M_{2.5D} \times \frac{s_{2.5D}}{s_{2.5D} + s_{3D}} + M_{3D} \times \frac{s_{3D}}{s_{2.5D} + s_{3D}}, \quad (2)$$

$$s_D = s_{2.5D} \times \frac{s_{2.5D}}{s_{2.5D} + s_{3D}} + s_{3D} \times \frac{s_{3D}}{s_{2.5D} + s_{3D}}. \quad (3)$$

Later, we linearly fuse (M_{2D}, s_{2D}) and (M_D, s_D) likewise to obtain the final probability mask M_f and its corresponding final mask score. The inferred mask is created by binarizing M_f .

The mask scoring process should not be different for each representation. We only use 2D image features and M_{2D} to train a single MaskIoU head instead of constructing 3 MaskIoU heads for each representation. In this way, the MaskIoU module would not add much more memory use and the training is also effective. The MaskIoU loss is denoted as \mathcal{L}_{miou} .

Table 1. **Comparison between collected HQDS and other public datasets for instance segmentation with stereo data.** Stereo pairs # means number of training stereo pairs. Stereo camera baseline (b) is in meters. f_x is for horizontal focal length.

Dataset	Resolution (megapixels)	Stereo Pairs #	b (m)	f_x (pixels)
Cityscapes	2.09	2.7K	0.2	2.2K
KITTI	0.71	0.2K	0.5	0.7K
HQDS	3.15	6K	0.5	3.3K

3. Experiments

3.1. HQDS Dataset

Outdoor RGBD scene understanding is still less explored since much longer range sensing is required to align information from images and depth. Such as vehicles at distances showing in images but undetected by a depth sensor would bring ambiguity into RGBD methods. To conduct exploration of outdoor RGBD methods, and provide high quality data to reveal advantages of sensor fusion, we collect High-Quality Driving Stereo (HQDS) dataset in urban environments. Table 3.1 shows a comparison with other public datasets for instance segmentation. Image resolution of HQDS is 1024×3072 . From the table and Eq. 1. HQDS has the largest $f \times b$. Measuring range by the configuration is up to **1650 meters** with 1-pixel disparity, which is only 440 and 350 for Cityscapes and KITTI. Note that produced disparity maps are computed by stereo matching methods so actual working distances are associated with methods' robustness and image noise. However, *longer baselines and focal lengths* still favor *far-field stereo matching* since the former could show *better geometry and more complete shapes* for objects at distances [5].

HQDS contains 6K/1.2K stereo pairs for training/testing. We follow a half-automation process to annotate data with

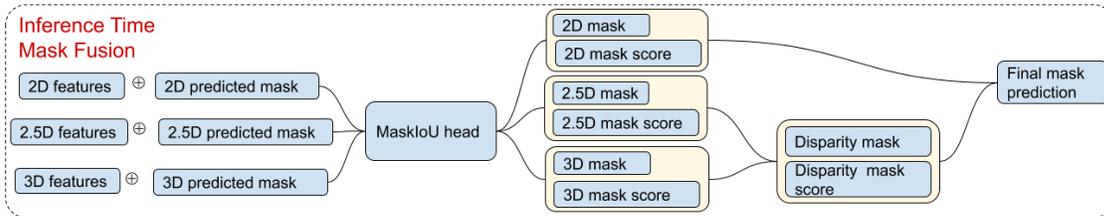


Figure 4. **Inference time mask fusion of predictions from different representations.** We fuse the 2.5D mask and 3D mask first because they are from the same source. We then fuse the mask predictions from the image domain and disparity. \oplus represents concatenation.

a group of supervised annotators. Our internal large-scale labeling system produces preliminaries, and the annotators adjust yielded bounding boxes and mask shapes or filter out false predictions to produce HQDS groundtruth.

There are 60K instances in the training set and 11K in the testing set. We adopt 3 instance classes: human, bicycle/motorcycle, and vehicle. Although other datasets on driving adopt more, such as Cityscapes use 8 classes, from MaskRCNN’s study [3] they suffer from much inter-class ambiguity which leads to biased results.

Associated number of instances in the training and testing sets are (5.5K, 1.5K, 52.8K) and (2.4K, 1K, 8.4K) respectively. Most of non-synthetic datasets encounter class-imbalanced issue. To remedy the imbalance, we adopt COCO dataset (instance segmentation for common objects) pretrained weights with class pruning in our implementations and comparison methods.

Evaluation and Metrics. We fairly compare with recent state-of-the-art methods validated on large-scale COCO dataset, including Mask-RCNN [3], MS-RCNN [4], Cascade Mask RCNN [2], and HTC [2] (w/o semantics), by using their publicly released codes and their COCO pretrained weights. We follow their training procedures to conduct comparison experiments.

We report numerical results in the **standard COCO-style**. Average precision (AP) averages across *different IoU levels*, from 0.5 to 0.95 with 0.05 as an interval. AP_{50} and AP_{75} are 2 typical IoU levels. The units are %. Table 2 shows the comparison with others. The proposed GAIS-Net attains the state of the art. We exceed Mask-RCNN using the same backbone by 9.7% and 6.8% for bounding box and mask AP, respectively.

3.2. Cityscapes Dataset

We also conduct experiments on Cityscapes dataset. However, its baseline and focal length are shorter than HQDS, and the maximal measuring distance is only **1/4** of HQDS. Much shorter focal length and baseline **limit the working distance** of stereo matching and produce disparity maps only focusing at near fields with *poor shapes and geometry* [5]. From Table 3, performance of GAIS-Net is still better than Mask-RCNN. The improvement gap between HQDS and Cityscapes is mainly caused by the latter’s shorter baseline and focal length.

Table 2. **Quantitative comparison on HQDS testing set.** The first table is for bounding box evaluation. The second table is for mask evaluation.

Bbox Evaluation	AP	AP_{50}	AP_{75}	AP_S	AP_L
Mask-RCNN	36.3	57.4	38.8	19.1	51.9
MS-RCNN	42.2	65.1	46.6	20.8	59.6
Cas. Mask-RCNN	37.4	55.8	38.9	18.0	54.7
HTC	39.4	58.3	43.1	18.5	57.9
GAIS-Net	46.0	67.7	53.3	23.6	66.2

Mask Evaluation	AP	AP_{50}	AP_{75}	AP_S	AP_L
Mask-RCNN	33.9	53.2	35.5	14.4	49.7
MS-RCNN	39.2	61.3	40.4	18.8	56.4
Cas. Mask-RCNN	33.4	54.4	34.8	11.7	49.5
HTC	34.5	56.9	36.7	11.6	52.0
GAIS-Net	40.7	65.9	43.5	18.3	59.2

Table 3. **Instance segmentation results on Cityscapes dataset.**

Evaluation	Training data	Mask AP
Mask-RCNN [3]	fine only	31.5
Our GAIS-Net	fine only	32.5
Mask-RCNN [3]	fine + COCO	36.4
Our GAIS-Net	fine + COCO	37.1

References

- [1] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, 2018. 2
- [2] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, 2019. 4
- [3] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 1, 2, 4
- [4] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring r-cnn. In *CVPR*, 2019. 2, 4
- [5] Masatoshi Okutomi and Takeo Kanade. A multiple-baseline stereo. *IEEE Transactions on pattern analysis and machine intelligence (TPAMI)*, 1993. 1, 3, 4
- [6] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *CVPR*, 2019. 2
- [7] Linwei Ye, Zhi Liu, and Yang Wang. Depth-aware object instance segmentation. In *ICIP*, 2017. 1