



MamboNet: Adversarial Semantic Segmentation for Autonomous Driving

Jheng-Lun Liu², Augustine Tsai^{1,3}(✉), Chiou-Shann Fuh², and Fay Huang³

¹ The Institute for Information Industry, Taipei, Taiwan
atsai@iii.org.tw

² Department of CSIE, National Taiwan University, Taipei, Taiwan
fuh@csie.ntu.edu.tw

³ Department of CSIE, National Ilan University, Yilan, Taiwan
fay@niu.edu.tw

Abstract. Environment semantic maps provide essential information for autonomous vehicles to navigate in complex road scenarios. In this paper, an adversarial network to complement the conventional encoder-decoder semantic segmentation network is introduced. A newly proposed adversarial discriminator is piggybacked to the segmentation network, which is used to improve the spatial continuity and label consistency in a scene without explicitly specifying the contextual relationships. The segmentation network itself serves as a generator to produce an initial segmentation map (pixel-wise labels). The discriminator then takes the labels and compare them with the ground truth data to further update the generator in order to enhance the accuracy of the labeling result. Quantitative evaluations were conducted which show significant improvement on spatial continuity.

Keywords: Generative Adversarial Network (GAN) · Semantic segmentation · Autonomous driving

1 Introduction

Surrounding understanding is critical to the safety of autonomous vehicles. The ability to recognize the drivable areas and dynamic objects on the road enables the safe navigation. Conventionally, camera frames are used to detect pedestrians, cars, motorcycles, roads, and sidewalks in pixel-level. The goal of this task is to produce semantic segmentations by assigning each input data point, namely a pixel, a unique class label. With the advancements of LiDAR sensor technology in recent years, many commercial products can detect points beyond 200 m. In this paper, we tackled the semantic segmentation task using a rotating LiDAR scanners. Comparing to solely using camera frames, 3D point clouds obtained by LiDAR provide a richer spatial and geometry information. However, the unstructured and sparse nature of the 3D data presents another level of challenges.

The major contribution of this paper is a novel method which can efficiently improve 3D LiDAR point cloud segmentation. We complemented an end-to-end encoder decoder

segmentation pipeline with an adversarial network which is derived from Generative Adversarial Network (GAN) [1]. The network improves the spatial continuity and label consistency without explicitly specifying the contextual information. The adversarial network was only applied during model training, and was removed during the online inference stage. The complexity of the overall architecture is kept in minimum.

2 Related Work

Semantic segmentation is one of the most important deep learning applications. In 2D image segmentation, U-Net [2] pioneered the encoder-decoder CNN architecture adoption, they transferred the entire feature map from encoders to the corresponding decoders and concatenates them to up-sampled (via deconvolution) decoder feature maps. In order to reduce memory requirements, Kendall [3] proposed to store the max pooling indices instead of concatenation with fewer parameters for decoder reconstruction.

Nowadays, 360° revolving LiDAR is the most common laser scanner for autonomous driving. In order to address 3D point cloud segmentation using aforementioned 2D segmentation paradigm, common approach is to spherically project the 3D point cloud data onto 2D range image plane. Leading the online frame-rate processing for practical applications, Wu [4] proposed a light weighted model derived from SqueezeNet to process data in 2D image plane. SqueezeSegV2 [5] extended V1 with Contextual Aggregation Module (CAM) [6] to mitigate LiDAR sensor data drop out issues.

A synthetic point cloud generation using GTA-V game engine with intensity rendering was also proposed to augment the training data. Due to nonhomogeneous spatial distribution of point cloud, SqueezeSegV3 [7] proposed Spatial-Adaptive Convolutions (SAC) which may change the weights according to the input data location. Miliotos [8] extended Wu [4] 3 label classes to 19 classes and replace extended the label classes from three to nineteen, and replaced the 2D CRF to 3D GPU-based nearest neighbor search acting directly on the full, un-ordered point cloud. This last step helps the retrieval of labels for all points in the cloud, even if they are occluded in the range image.

Cortinhal [9] transformed the deep network with Bayesian treatment by introducing uncertainty measures, epistemic and aleatoric noises. Luc [10] introduced an adversarial network to discriminate the predicted segmentation maps either from the ground truth or segmentation network to mitigate the higher order label inconsistencies. Souly [11] introduced a semi-supervised segmentation using weakly labelled data for the generator. In this paper, the proposed MamboNet was inspired by many of these approaches and mostly by Luc’s adversarial network.

3 Method

A. 3D to 2D Projection

The projection method as mentioned in [4, 5, 7–9] has been applied for data pre-processing. Each raw 3D point cloud in 360° surrounding is spherical projected onto a 2D grid point on a range image as illustrated in Fig. 1. A 3D point (x, y, z) with respect

to the world coordinate system originated at the sphere center is projected to the image with coordinates of $(\theta_{loc}, \phi_{loc})$, which is calculated as follows:

$$\begin{aligned}\theta &= \arcsin \frac{z}{\sqrt{x^2 + y^2 + z^2}}, \theta_{loc} = \lfloor \theta / \Delta\theta \rfloor \\ \phi &= \arcsin \frac{y}{\sqrt{x^2 + y^2}}, \phi_{loc} = \lfloor \phi / \Delta\phi \rfloor\end{aligned}\quad (1)$$

Here, $\Delta\theta$ and $\Delta\phi$ are quantization steps. Each grid point represents a five-dimensional feature vector: three for its associated 3D location (x, y, z) , one for the intensity value, and the other for the range value.

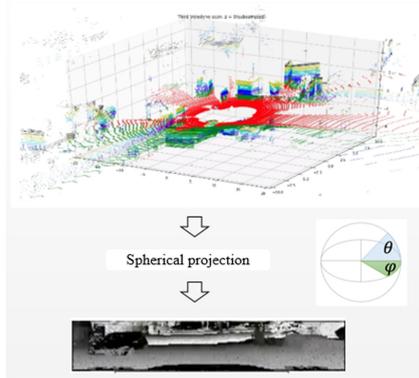


Fig. 1. Spherical projection

B. Architecture

The main objective of applying adversarial network is to enforce the spatial continuity and label consistency. Conventional encoder-decoder network [3] creates a segmentation map (pixel-wise labeling), and then follows up with conditional random field (CRF) to impose pixel grouping constraints. We replaced CRF with a discriminator which is only used during the training and can be dropped in inference to maintain minimum network complexity and it is similar to bag of freebies in [11]. Our adversarial network (shown in Fig. 2) is similar to [10], the discriminator takes two inputs, namely, predicted and ground truth maps. Both maps are concatenated with the same 2D input data. The predicted map is generated by the encoder-decoder semantic segmentation network.

A detailed version of the generator is shown in Fig. 3, each yellow block of the encoder is an Inception [13] like module with a group of mixed kernel sizes and dilation rates. Each block has three parallel convolution layers, the outputs are concatenated and then summed up with forth convolution layer. Between encoder and decoder, an Astrous Spatial Pyramid Pooling (ASPP) [2] module is inserted for exploiting multi-scale features and enlarging the receptive field. ASSP is employed to capture small street objects, such as pedestrian and cyclists. In decoder, the conventional transpose convolution layer

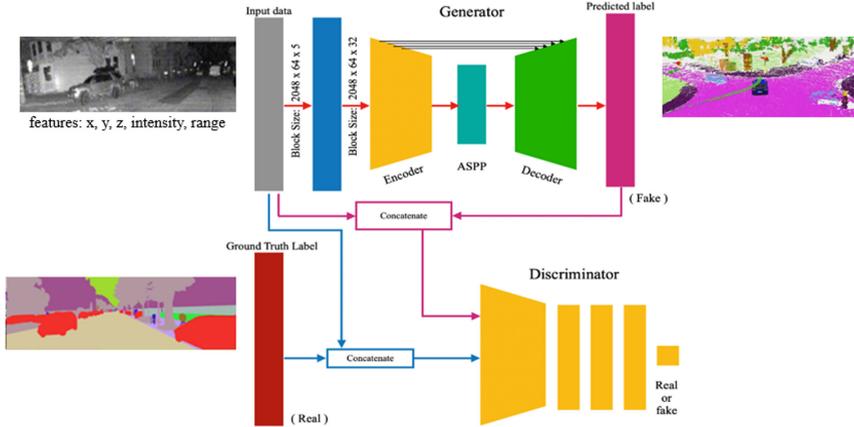


Fig. 2. Overall network architecture

is replaced with the low computation pixel-shuffle layer, similar to super resolution [14]. It can leverage low resolution feature map to generate up-sampled feature maps by converting information of the channel dimension to the spatial dimension. The operation is to convert a feature map of $(H \times W \times Cr^2)$ to $(Hr \times Wr \times C)$, where H, W, C and r are the height, width, number of channel, and up-sampling factor.

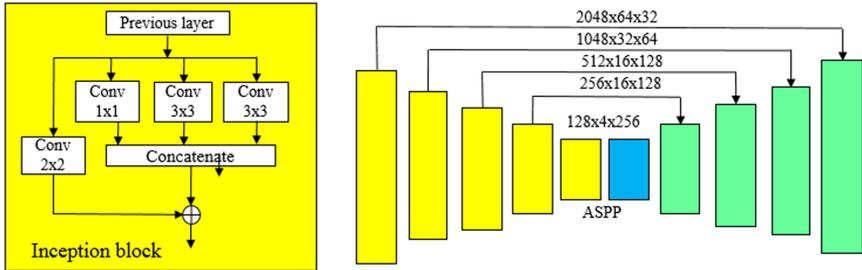


Fig. 3. Details of the generator

The discriminator is a VGG based convolutional network shown in Fig. 4. The data size is $2048 \times 64 \times 6$. The first two dimensions are the image width and height, and the third dimension includes $x, y, z, \text{intensity, range, and class label}$. Each layer uses 3×3 convolution kernel and is followed by a 2×2 max pooling except for the 1st layer. The sizes of the last three fully connected layers are 2048, 512, and 512, respectively.

C. Loss Function and Training

The training, shown in Fig. 5, is based on conditional GAN ($cGAN$) [15] architecture. The discriminator, D , learns to classify fake (predicted semantic map) and real (ground truth map). Both generator and discriminator observe the same 2D range imagery input.

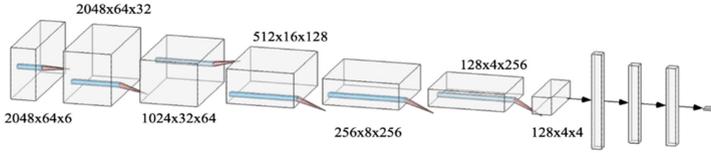


Fig. 4. Discriminator: VGG based convolutional network

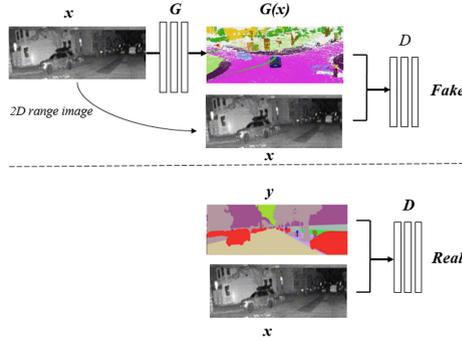


Fig. 5. Conditional GAN training: map 2D range imagery to segmentation map.

There are three lost terms, the first term is the general cross-entropy term for segmentation network (generator), $S(\cdot)$, to predict each location (pixel-wise) of the output map with independent class label. It is a weighted cross-entropy loss as is expressed as.

$$L_{wce} = - \sum_{c \in C} \frac{1}{\sqrt{f_c}} Y_c \log(S_c) \tag{2}$$

where Y and S are the one-hot vector maps for ground truth and predicted label, respectively. Due to the imbalance data nature of the street scene, pedestrians and cyclists are less seen compared to other cars, the way to mitigate the network biases toward the classes with higher frequency of occurrence is to add a weighted factor f . The second term is the *Lovász -Softmax* loss [16]. The loss is used to improve the intersection-of-union (IoU) or *Jaccard index*. The convex *Lovász* extension of submodular losses relaxes the IoU hypercube constraint where each vertex is a plausible combination of the class labels. Therefore, IoU score can be defined anywhere inside the hypercube. This term is expressed as

$$L_{ls} = \frac{1}{|C|} \sum_{c \in C} \Delta J_C(m(c)) \tag{3}$$

$$m_i(c) = 1 - x_i(c) \text{ if } c = y_i(c) \\ = x_i(c), \text{ otherwise}$$

where $x_i(c) \in [0,1]$ is the pixel-wise predicted probability, $y_i(c)$ is the predicted label. The loss will penalize the wrong prediction.

The third term is the adversarial loss which can be expressed as.

$$L_{adv}(G, D) = E_{x, P_{gt}}[\log D(x, y)] + E_{x, P_p}[\log(1 - D(x, G(x, z))), \quad (4)$$

where D is the discriminator which produces *Real* and *Fake* binary outputs, and G generates the predicted label, x is the 2D range image, z is the optional random noise input, P_{gt} is the distribution of ground truth label, y , and P_p is the distribution of the predicted label.

D tries to maximize the *Jensen-Shannon* divergence [1] between P_{gt} and P_p . On the contrary, G tries to minimize the same distribution divergence in order to make P_p indistinguishable from P_{gt} . The final objective is a mix-max optimization of the loss summation of cross entropy, *Lovász -Softmax* and adversarial terms as shown in Eq. (5)

$$G^* = \arg \min_G \max_D L_{adv} + L_{ls} + L_{ce} \quad (5)$$

4 Experiments

Semantic KITTI data set [17] was used for algorithm evaluation. The dataset contains 28 classes including classes of non-moving and moving objects. The scanned sequences of 0–10 except 8 were used for training, and sequence 8 was used for validation. Sequences 11–21 was used for testing, however, the annotations for the testing sequence are not available to the general public. In order to evaluate the performance, the labeled data were submitted to Semantic KITTI official server for test results. The evaluation metric is based on *Jaccard Index* or mean Intersection-over-Union (IoU) metric as shown in the Eq. (6).

$$mIoU = \frac{1}{C} \sum_{c=1}^C \frac{TP_c}{TP_c + FP_c + FN_c} \quad (6)$$

where TP , FP , and FN correspond to the number of true positive, false positive, and false negative predictions for class c , and C is the number of classes.

A. Quantitative Results

In Table 1, our method not only outperforms most of the 3D point-wise methods [18–21], but also is superior to other projection based methods, especially in small object segmentation, such as person, bicyclist, and motor-cyclist categories. We compare our method with two other networks, the first one is the SalsaNext baseline [9], and the second one is the SalsaNext augmented with a discriminator. The discriminator is a VGG-based convolutional network. In the beginning, we trained the SalsaNext baseline using their open source Github repository, and the test result of mIOU is 57.2, which is a little lower than the published result (59.5) [9]. The discrepancy can be due to the limited batch size [15] in our single board training configuration. In Table 1, SalsaNext with discriminator outperforms baseline in 15 out of 19 categories, and the mIOU of 57.9 is slightly improved. Our method, MamboNet, achieves over one percent mOUT improvement of 58.5.

Table 1. Quantitative results comparison on SemanticKITTI testing set (Sequence 11–21).

Approach	car	bicycle	motor-cycle	truck	other-vehicle	person	bicyclist	motor-cyclist	road	parking	sidewalk	Other-ground	building	fence	vegetation	trunk	terrain	Pole	Traffic-sign	mean-IOU	
Point-wise	Pointnet	46.3	1.3	0.3	0.1	0.8	0.2	0.0	61.6	15.8	35.7	1.4	41.4	12.9	31.0	4.6	17.6	2.4	3.7	14.4	
	Pointnet++	53.7	1.9	0.2	0.9	0.2	0.9	1.0	72.0	18.7	41.8	5.6	62.3	16.9	46.5	13.8	30.0	6.0	8.9	20.1	
	RandLA-Net	94.2	26.0	25.8	40.1	38.9	49.2	48.2	7.2	90.7	60.3	73.7	20.4	86.9	56.3	81.4	61.3	66.8	49.2	47.1	53.9
	LatticeNet	88.6	12.0	20.8	43.3	24.8	34.2	39.9	60.9	88.8	64.6	73.8	25.5	86.9	55.2	76.4	67.9	54.7	41.5	42.7	52.2
Projection-based	SqueezeSegV2	81.8	18.5	17.9	13.4	14.0	20.1	25.1	3.9	88.6	45.8	67.6	17.7	73.7	41.1	71.8	35.8	60.2	20.2	36.3	39.7
	SqueezeSegV2-CRF	82.7	21.0	22.6	14.5	15.9	20.2	24.3	2.9	88.5	42.4	65.5	18.7	73.8	41.0	68.5	36.9	58.9	12.9	41.0	39.6
	RangeNet++	91.4	25.7	34.4	25.7	23.0	38.3	38.8	4.8	91.8	65.0	75.2	27.8	87.4	58.6	80.5	55.1	64.6	47.9	55.9	52.2
	SqueezeSegV3	92.5	38.7	36.5	29.6	33.0	45.6	46.2	20.1	91.7	63.4	74.8	26.4	89.0	59.4	82.0	58.7	65.4	49.6	58.9	55.9
	SalsaNext (train by ourself)	91.1	43.7	34.3	37.5	29.9	59.1	53.6	30.4	90.9	60.4	73.2	23.6	87.6	56.6	79.5	58.6	63.9	51.3	62.0	57.2
	SalsaNext +Discriminator	91.1	46.6	32.0	37.1	31.8	57.5	55.0	26.2	91.0	61.8	74.7	25.7	87.9	57.6	80.1	61.1	66.2	54.7	62.6	57.9
	MamboNet	92.0	47.4	39.0	25.6	34.6	59.0	57.6	27.8	91.8	64.9	75.0	21.3	88.8	60.5	81.3	64.7	64.7	54.6	61.2	58.5

B. Qualitative Results

In Fig. 6, four blocks of segmented map results are shown for visual examination. Each block has three maps, the top is the SalsaNext baseline, the middle one is our method with adversarial discriminator, and the bottom one is the ground true for comparison.

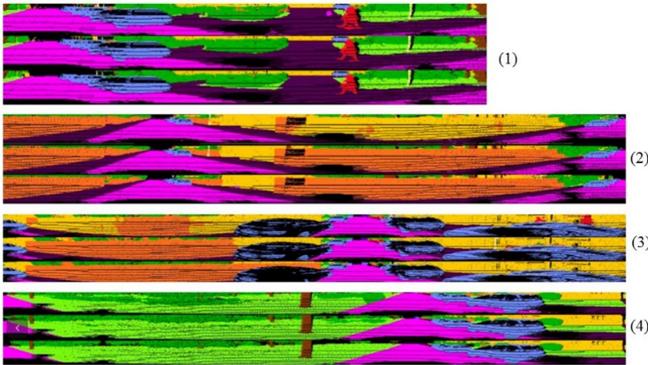


Fig. 6. Qualitative evaluation with four examples: the top strip of each example is the result without adversary training, the middle strip is with adversary training, and the bottom strip is the ground-truth label.

In the top strip of the first example, there is a small mis-classified pink circle inside the dark purple region (road). The middle strip of the same example, the circle disappears due to the discriminator power of enforcing regional consistency. The same rectification can be observed in the second and third examples, all middle strips correctly identify the fence region (brown), while the top strip mis-classify part of the fence to be the building regions (yellow). Finally in the fourth example, the top strip also misclassifies portion of light green (terrain) to be dark green (vegetation), however, the middle strip correctly identifies the terrain area.

5 Conclusion

We augmented an encoder-decoder segmentation network with an adversarial network to improve the semantic segmentation performance. Adversarial network can implicitly enforce the regional contextual continuity. Unlike conventional CRF and KNN post processing techniques, the adversarial is learnt only during the offline training and is not active during the test. Therefore, the online computation is greatly reduced and yet the comparable results are still attainable.

References

1. Goodfellow, I.J., et al.: Generative adversarial networks. In: NIPS (2014)
2. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (2015)
3. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian SegNet: model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In: Proceedings of the British Machine Vision Conference (BMVC), pp. 57.1–57.12. BMVA Press (2017)
4. Wu, B., Wan, A., Yue, X., Keutzer, K.: Squeezeseg: convolutional neural nets with recurrent CRF for real-time road-object segmentation from 3D lidar point cloud. In: ICRA (2018)
5. Wu, B., Zhou, X., Zhao, S., Yue, X., Keutzer, K.: Squeezesegv2: improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In: ICRA (2019)
6. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: ICLR (2016)
7. Xu, C., et al.: Squeezesegv3: spatially-adaptive convolution for efficient point-cloud segmentation, [arXiv:2004.01803](https://arxiv.org/abs/2004.01803) (2020)
8. Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: RangeNet++: fast and accurate LiDAR semantic segmentation. In: IROS (2019)
9. Cortinhal, T., Tzelepis, G., Aksoy, E.E., SalsaNext: fast, uncertainty-aware semantic segmentation of LiDAR point clouds for autonomous driving, [arXiv:2003.03653](https://arxiv.org/abs/2003.03653) (2020)
10. Luc, P., Couprie, C., Chintala, S., Verbeek, J.: Semantic segmentation using adversarial network. In: 2016, Workshop on Adversarial Training, in NIPS (2016)
11. Souly, N., Spampinato, C., Shah, M.: Semi supervised semantic segmentation using generative adversarial network. In: ICCV (2017)
12. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: YOLOv4: Optimal Speed and Accuracy of Object Detection, [arXiv:2004.10934](https://arxiv.org/abs/2004.10934) (2020)
13. Szegedy, C., et al.: Going deeper with convolutions. In: CVPR (2015)
14. Shi, W., et al.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: CVPR (2016)
15. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-Image translation with conditional adversarial networks. In CVPR (2017)
16. Berman, M., Triki, A.R., Blaschko, M.B.: The Lovász-softmax loss: a tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In: CVPR (2018)
17. Behley, J., et al.: SemanticKITTI: a dataset for semantic scene understanding of LiDAR sequences. In: Proceedings of International Conference on Computer Vision, Seoul, Korea, p. 17 (2019)
18. Hu, Q., et al.: Randla-net: efficient semantic segmentation of largescale point clouds. In: CVPR (2020)
19. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: deep learning on point sets for 3D classification and segmentation. In: CVPR (2017)

20. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: deep hierarchical feature learning on point sets in a metric space. In: NIPS (2017)
21. Rosu, R.A., Schütt, P., Quenzel, J., Behnke, S.: LatticeNet: Fast Point Cloud Segmentation Using Permutohedral Lattices, [arXiv:1912.05905](https://arxiv.org/abs/1912.05905) (2019)