LiDAR-UDA: Self-ensembling Through Time for Unsupervised LiDAR Domain Adaptation
Supplementary Materials

Amirreza Shaban* JoonHo Lee* Sanghun Jung* Xiangyun Meng Byron Boots
University of Washington

1. LiDAR Intensity

Source	Target (no DA)	source model	Intensity Used	mIoU (%) ↑
KITTI	KITTI	MinkowskiNet	✓	66.72
			✗	60.24
		SalsaNext	✓	58.49
			✗	55.81
	USL	MinkowskiNet	✓	38.56
			✗	41.51
		SalsaNext	✓	23.62
			✗	36.95

Table 1. Generalization performance of MinkowskiNet and SalsaNext trained with and without intensity values as input. We train the source models on SemanticKITTI (KITTI) and evaluate their performance on SemanticUSL (USL). Dropping intensity leads to significantly improved generalization performance on the target domain.

Although LiDAR intensity provides additional information for distinguishing geometrically similar objects, we find each LiDAR sensor has significantly different intensity ranges and distribution from others. While prior works apply various techniques to utilize the intensity as an additional input [2, 4], we view intensity matching to remain a non-trivial problem on its own. As we show in Table 1, using intensity results in degraded generalization performance of the source model - an average drop of 8.1% mIoU between MinkowskiNet and SalsaNext.

Furthermore, as our framework *re-introduces* intensity of the target domain as an additional input during the student model training, the proposed LiDAR-UDA method is still able to utilize the intensity information of the target domain in a transferable manner.

2. Classwise mIoU Analysis on SemanticKITTI and nuScenes

We present classwise mIoU results of our method and the source model for the SemanticKITTI ↔ nuScenes DA scenario in Table 2. Our method achieves significant improvements over the source model in most classes, indicating its robust generalization capability. For example, in SemanticKITTI→nuScenes, our method improves over the source model by 36.06% IoU in Pedestrian and 21.21% IoU in Bus classes. In nuScenes→SemanticKITTI, our method improves by 31.92% IoU in Drivable and 28.21% IoU in Car classes.

3. LAM Analysis: Weight Distributions

To gain better insight of the weighting dynamics behind LAM, we collected statistics of the weights predicted by LAM in the SemanticKITTI→nuScenes experiment on various slices of the target dataset, visualized as histograms in Figures 1 to 3.

- **Temporal offset:** Figure 1 shows weight histograms for different temporal distances from the current frame t. We find that points with a shorter temporal distance to t (3rd histogram) have relatively higher weights.
- **Distance from Sensor:** Figure 2 shows that LAM favors points closer to the sensor origin in their corresponding LiDAR scan (first row), giving them higher weights. This matches the intuition because the model's accuracy typically decreases for objects farther away from the sensor due to the sparsity of LiDAR point cloud. To illustrate the significance of this feature, we point out that the aggregated point cloud combines 36 seconds of LiDAR data so if the vehicle speed is as low as 10 miles per hour, the distance variation of points within an ϵ-ball to their sensor location can exceed 150 meters.
- **Distance from center:** Figure 3 reports on the distance from the center of ϵ-ball ($||\mathbf{p} - \mathbf{p}'||_2$), which shows only a slight decay in the assigned weights as the distance increases, suggesting that LAM depends more on other factors.

Our analysis provides a limited view of predicted weights within each slice, and we overall observe that LAM does not soley rely on only one of the input features above but the combination of all input features to make predictions.

*Equal Contribution.

Source	Target	Method	Drivable	Sidewalk	Terrain	Pedestrian	Vegetation	Bicycle	Bus	Car	Motorcycle	Truck	mIoU
KITTI	KITTI	Source	87.66	68.75	65.17	18.63	90.12	0.31	15.53	90.17	3.70	17.99	45.80
	NUS	Source	79.19	32.22	21.72	4.72	73.97	0.07	4.46	55.32	3.35	2.46	27.75
		SqueezeSegV2* [4]	-	-	-	-	-	-	-	-	-	-	10.10
		SWD* [3]	-	-	-	-	-	-	-	-	-	-	27.70
		Complete & Label [5]	-	-	-	-	-	-	-	-	-	-	31.60
		Graph Matching [1]	-	-	-	-	-	-	-	-	-	-	37.30
		LiDAR-UDA (Ours)	**87.44**	**42.31**	**47.88**	**40.78**	**83.22**	**0.85**	**25.67**	**73.48**	**15.86**	0.90	**41.84**
NUS	NUS	Source	91.44	52.3	58.33	48.05	86.42	3.00	24.23	80.39	28.88	34.18	50.72
	KITTI	Source	33.77	2.81	30.05	12.56	80.94	0.45	4.95	57.98	4.40	3.82	23.17
		SqueezeSegV2* [4]	-	-	-	-	-	-	-	-	-	-	13.40
		SWD* [3]	-	-	-	-	-	-	-	-	-	-	24.50
		Complete & Label [5]	-	-	-	-	-	-	-	-	-	-	33.70
		LiDAR-UDA (Ours)	**65.69**	**6.07**	**54.05**	**16.49**	**85.65**	0.00	3.15	**86.19**	**13.87**	**9.30**	**34.04**

Table 2. Classwise mIoU (%) of methods on KITTI to Nuscnes and nuScenes to KITTI adaptation scenarios.

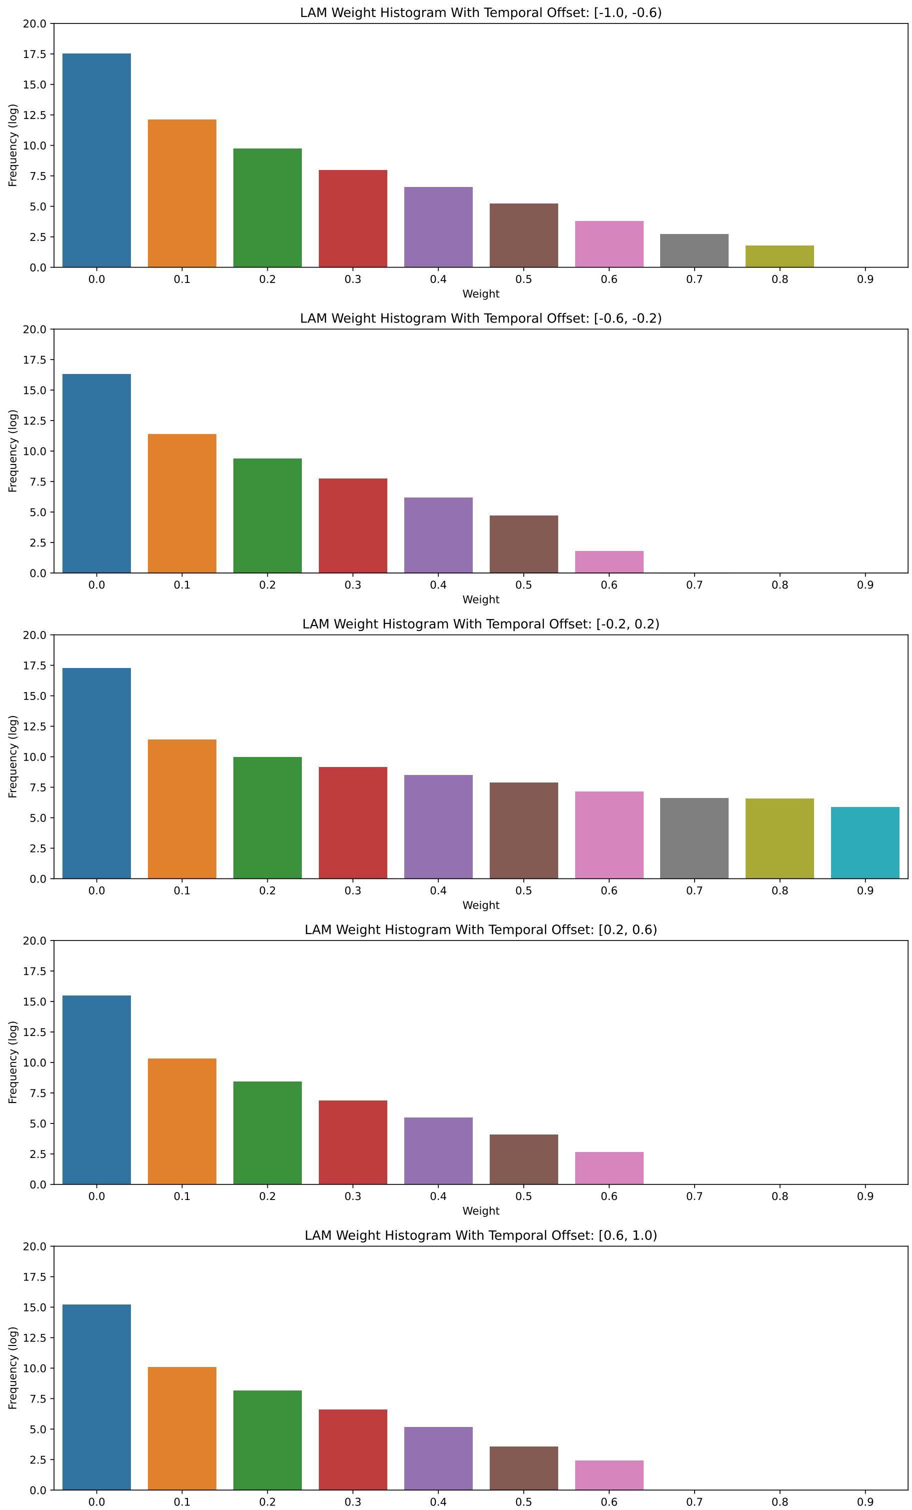

Figure 1. Histogram visualization of LAM weights applied to $\mathcal{B}_\epsilon$ neighbors based on their temporal distance of current frame t normalized between $(-1, +1)$.

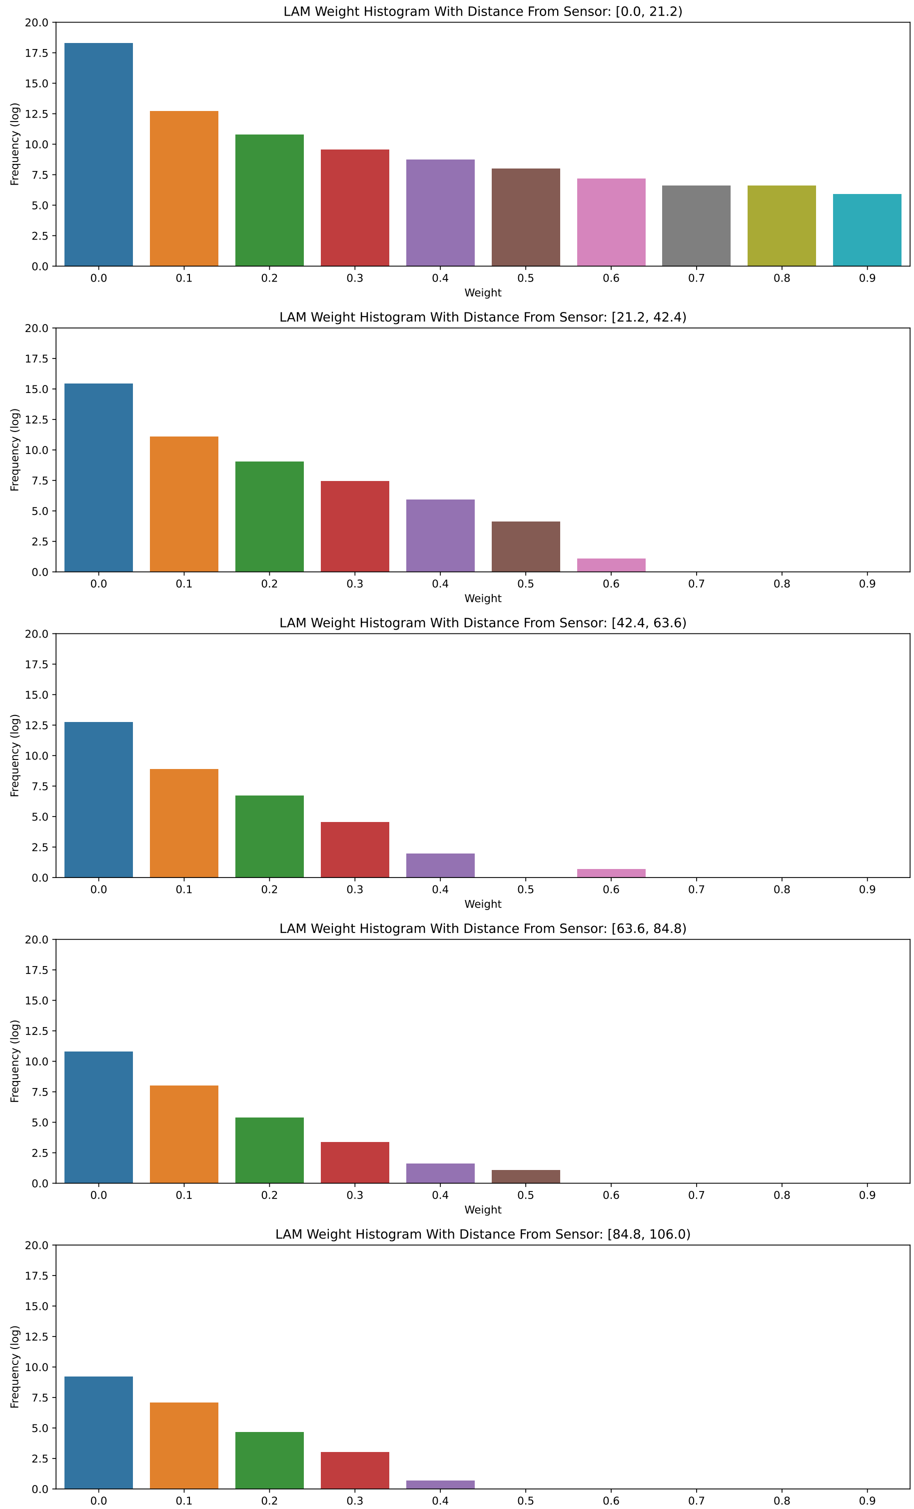

Figure 2. Histogram visualization of LAM weights applied to $\mathcal{B}_\epsilon$ neighbors based their distance to the sensor location in their corresponding LiDAR scan.

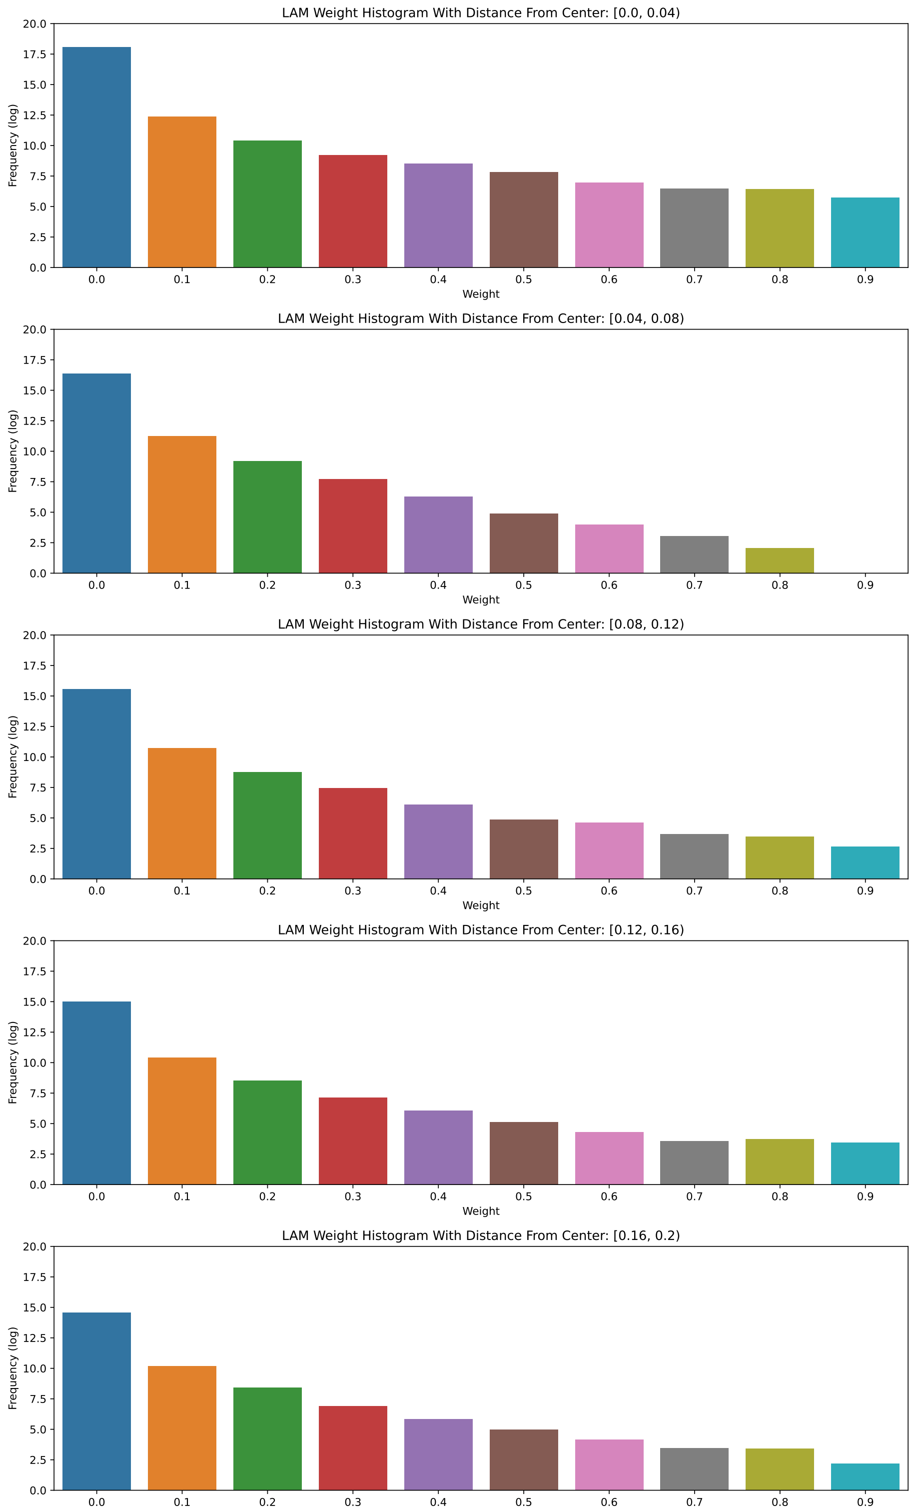

Figure 3. Histogram visualization of LAM weights applied to $\mathcal{B}_\epsilon$ neighbors based on their distance.

References

[1] Yikai Bian, Le Hui, Jianjun Qian, and Jin Xie. Unsupervised domain adaptation for point cloud semantic segmentation via graph matching. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9899–9904. IEEE, 2022. 2

[2] Peng Jiang and Srikanth Saripalli. Lidarnet: A boundary-aware domain adaptation model for point cloud semantic segmentation, 2020. 1

[3] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10285–10295, 2019. 2

[4] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *ICRA*, 2019. 1, 2

[5] Li Yi, Boqing Gong, and Thomas Funkhouser. Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15363–15373, June 2021. 2