

Domain Adaptation in LiDAR Semantic Segmentation by Aligning Class Distributions

Iñigo Alonso¹, Luis Riazuelo¹, Luis Montesano^{1,2} and Ana C. Murillo¹

¹RoPeRt Group, DIIS - I3A, Universidad de Zaragoza, Spain

²Bitbrain, Zaragoza, Spain

Keywords: Robotics, Autonomous Systems, LiDAR, Deep Learning, Semantic Segmentation, Domain Adaptation.

Abstract: LiDAR semantic segmentation provides 3D semantic information about the environment, an essential cue for intelligent systems, such as autonomous vehicles, during their decision making processes. Unfortunately, the annotation process for this task is very expensive. To overcome this, it is key to find models that generalize well or adapt to additional domains where labeled data is limited. This work addresses the problem of unsupervised domain adaptation for LiDAR semantic segmentation models. We propose simple but effective strategies to reduce the domain shift by aligning the data distribution on the input space. Besides, we present a learning-based module to align the distribution of the semantic classes of the target domain to the source domain. Our approach achieves new state-of-the-art results on three different public datasets, which showcase adaptation to three different domains.

1 INTRODUCTION

3D semantic segmentation is an important computer vision task that provides useful information to each registered 3D point of the surrounding environment. It has a wide range of applications and is particularly important in robotics, since most autonomous systems require an accurate and robust perception of their environment. LiDAR (Light Detection And Ranging) is a frequently used sensor for 3D perception in autonomous vehicles, such as cars or delivery robots, that provides accurate distance measurements of the surrounding 3D space. Despite LiDAR broad adoption, recognition systems such as semantic segmentation methods that generalize and perform well for different LiDAR sensors, vehicle set-ups or environments remains an unsolved challenging problem. Existing solutions typically need significant amounts of labeled data to adapt to different domains (Mei et al., 2019).

In recent years, deep learning methods are achieving state-of-the-art performance in 3D LiDAR semantic segmentation (Milioto et al., 2019; Alonso et al., 2020) when trained and evaluated in the same domain. Nevertheless, generalizing the learned knowledge to new domains or environments not seen during training is still a frequent open problem. When applying existing models on data with a different distribution

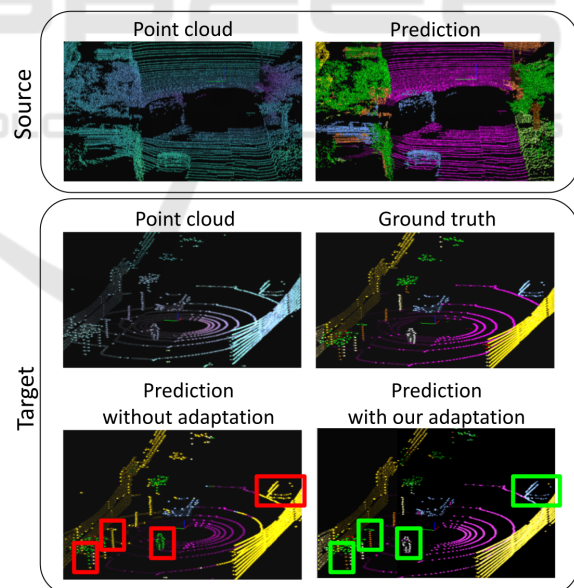


Figure 1: Domain Adaptation in LiDAR Semantic Segmentation. Top row: segmentation result of a given model train on the same source domain (Behley et al., 2019). Bottom row: segmentation of the same model on a *target domain* without adaptation and the improved result applying our proposed adaptation. Note relevant details missed if we do not adapt the model (marked with a red/green square).

than the training data, i.e., from a different domain, the performance is considerably degraded (Wang and Deng, 2018; Tzeng et al., 2017; Vu et al., 2019).

Domain adaptation techniques aim to eliminate or reduce this drop. Existing works for domain adaptation in semantic segmentation focus on RGB data (Vu et al., 2019; Zou et al., 2019; Chen et al., 2019; Li et al., 2019). Most of them try to minimize the distribution shift between two different domains. Very few approaches have tackled this problem with LiDAR data (Wu et al., 2019), which equally suffers from the domain shift. RGB data commonly suffers from variations due to light and weather conditions, while the most common variations within 3D point clouds data come from sensor resolution (i.e., sensors with more laser sweeps generate denser point clouds) and from the sensor placement (because point clouds have relative coordinates with respect to the sensor). Both sensor resolution and placement issues are common examples that change the data distribution of the captured 3D point clouds. Coping with these issues would enable the use of large existing labeled LiDAR datasets for more realistic use-cases in robotic applications, reducing the need for data labeling.

This work proposes two strategies to improve unsupervised domain adaptation (UDA) in LiDAR semantic segmentation, see a sample result on Fig. 1. The first strategy addresses this problem by applying a set of simple data processing steps to align the data distribution reducing the domain gap on the input space. The second strategy proposes how to align the distribution on the output space by aligning the class distribution. These two proposed strategies can be applied in conjunction with current state-of-the-art approaches boosting their performance.

We validate our approach on different scenarios getting state-of-the-art results. We use the SemanticKitti (Behley et al., 2019) as the source domain and we adapt it to SemanticPois (Pan et al., 2020), to Paris-Lille-3D (Roynard et al., 2018) and to a new collected and released dataset.

2 RELATED WORK

This section describes the related work for the most relevant related topics for this work: 3D LiDAR segmentation and, unsupervised domain adaptation for semantic segmentation.

2.1 3D LiDAR Point Cloud Segmentation

3D LiDAR Semantic segmentation aims to assign a semantic label to every point scanned by the LiDAR sensor. Before the current trend and wide adoption of deep learning approaches, earlier methods relied on exploiting prior knowledge and geometric constraints (Xie et al., 2019). As far as deep learning methods are concerned, there are two main types of approaches to tackle the 3D LiDAR semantic segmentation problem. On one hand, there are approaches that work directly on the 3D points, i.e., the raw point cloud is taken as the input (Qi et al., 2017a; Qi et al., 2017b). On the other hand, other approaches convert this 3D point cloud into another representation (images(Alonso et al., 2020), voxels(Zhou and Tuzel, 2018), lattices(Rosu et al., 2019)) in order to have a structured input. For LiDAR semantic segmentation, the most commonly used representation is the spherical projection (Alonso et al., 2020; Milioto et al., 2019; Wu et al., 2018). Several recent works (Milioto et al., 2019; Wu et al., 2018; Alonso et al., 2020) show that point-based methods, i.e., approaches that work directly on the 3D points, are slower and tend to be less accurate than methods which project the 3D point cloud into a 2D representation and make use of convolutional layers.

2.2 Unsupervised Domain Adaptation for Semantic Segmentation

Unsupervised Domain Adaptation (UDA) aims to adapt models that have been trained on one specific domain (source domain) to be able to work on a different domain (target domain) where there is a certain lack of labeled training data. Most works follow similar ideas: the input data or features from a source-domain sample and a target-domain sample should be indistinguishable. Several works follow an adversarial training scheme to minimize the distribution shift between the target and source domains data. This approach has been shown to work properly at pixel space (Yang and Soatto, 2020), at feature space (Hoffman et al., 2018) and at output space (Vu et al., 2019). However, adversarial training schemes, such as DANN (Ganin et al., 2016), tend to present convergence problems. Alternatively, other works follow different schemes like CORL (Sun et al., 2017) or MinEnt (Vu et al., 2019). Entropy minimization methods (Vu et al., 2019; Chen et al., 2019) do not require complex training schemes. They rely on a loss function that minimizes the entropy of the unlabeled target domain output probabilities.

Regarding segmentation on LiDAR data, very few works have studied the problem of domain adaptation. SqueezeSegV2 (Wu et al., 2019) proposes how to adapt synthetic data where only coordinates are available without reflectance to real data. For this, they propose a learned intensity rendering to create the intensity for the synthetic data. A very recent work, Xmuda (Jaritz et al., 2020) focuses on combining different modalities: LiDAR and RGB for multi-modal domain adaptation. They propose to apply the KL divergence between the output probabilities of both modalities as the main loss function. Besides, they also apply previously proposed methods like entropy minimization (Vu et al., 2019). Differently, this work investigates different UDA strategies (both existing and novel) to improve UDA for the particular case of urban LiDAR semantic segmentation. The presented results show their effectiveness in reducing the domain gap.

3 UNSUPERVISED DOMAIN ADAPTATION FOR LiDAR SEMANTIC SEGMENTATION

This section describes the proposed domain adaptation approach, including the LiDAR semantic segmentation method used, the strategies proposed to reduce the domain gap (data alignment and class distribution alignment), and the formulation of the proposed learning task. Figure 2 presents an overview of our proposed approach which is further explained in the following subsections.

3.1 LiDAR Semantic Segmentation Model

We use a recent method for LiDAR semantic segmentation which achieves state-of-the-art performance on several LiDAR segmentation datasets, 3D-MiniNet (Alonso et al., 2020). This method consists of three main steps. First, it learns a 2D representation from the 3D points. Then, this representation is fed to a 2D Fully Convolutional Neural Network that produces a 2D semantic segmentation. These 2D semantic labels are re-projected back to the 3D space and enhanced through a post-processing module.

Let $\mathcal{X}_s \subset \mathbb{R}^{N \times 3}$ be a set of source domain LiDAR point clouds along with associated semantic labels, $\mathcal{Y}_s \subset (1, C)^N$. Sample \mathbf{x}_s is a point cloud of size N and $\mathbf{y}_s^{(n,c)}$ provides the label of point (n) as one-hot encoded vector. Let F be our LiDAR segmentation network which takes a point cloud \mathbf{x} and predicts a

probability distribution (size C classes) for each point of the point-cloud $F(\mathbf{x}) = \mathbf{P}_{\mathbf{x}}^{(n,c)}$.

The parameters θ_F of F are optimized to minimize cross-entropy loss:

$$\mathcal{L}_{seg}(\mathbf{x}_s, \mathbf{y}_s) = - \sum_{n=1}^N \sum_{c=1}^C \mathbf{y}_s^{(n,c)} \log \mathbf{P}_{\mathbf{x}_s}^{(n,c)}, \quad (1)$$

on source domain samples. Therefore, as the supervised semantic segmentation is concerned, the optimization problem simply reads:

$$\min_{\theta_F} \frac{1}{|\mathcal{X}_s|} \sum_{\mathbf{x}_s \in \mathcal{X}_s} \mathcal{L}_{seg}(\mathbf{x}_s, \mathbf{y}_s). \quad (2)$$

3.2 Data Alignment Strategies for LiDAR

The problem of domain adaptation, i.e., data distribution misalignment, between \mathcal{X}_s and \mathcal{X}_t (a set of target domain LiDAR point clouds), can be handled on the network weights θ_F but also modifying \mathcal{X}_s and \mathcal{X}_t in order to align the distributions at the input space.

Next, we describe the different strategies for better data alignment that we propose to improve LiDAR domain adaptation.

XYZ-shift Augmentation. One of the main causes of misalignment for LiDAR point clouds is the location of the sensor. Since the point cloud values are relative to the sensor origin, these changes cause variations affecting the whole point cloud. We perform shifts up to ± 2 meters on the Z-axis (height) and up to ± 5 meters on the Y-axis and X-axis.

Per-class Augmentation. We also propose to perform the augmentation independently per class, i.e., applying different data augmentation parameters for every class. In particular, in this work, we perform shifts up to ± 1 meters on the Z-axis (height) and up to ± 3 meters on the Y-axis and X-axis.

Same Number of Beams. Besides the sensor placement and orientation, a significant difference between sensors is the number of captured beams, which results in a more sparse or dense point cloud. We propose to match the data beams between the two domains by reducing the data from the sensor with a higher number of beams ending up with more homogeneous data within \mathcal{X}_s and \mathcal{X}_t .

Only Relative Features. Point-cloud segmentation methods commonly use both absolute and relative (between points) values of the input data. In order to

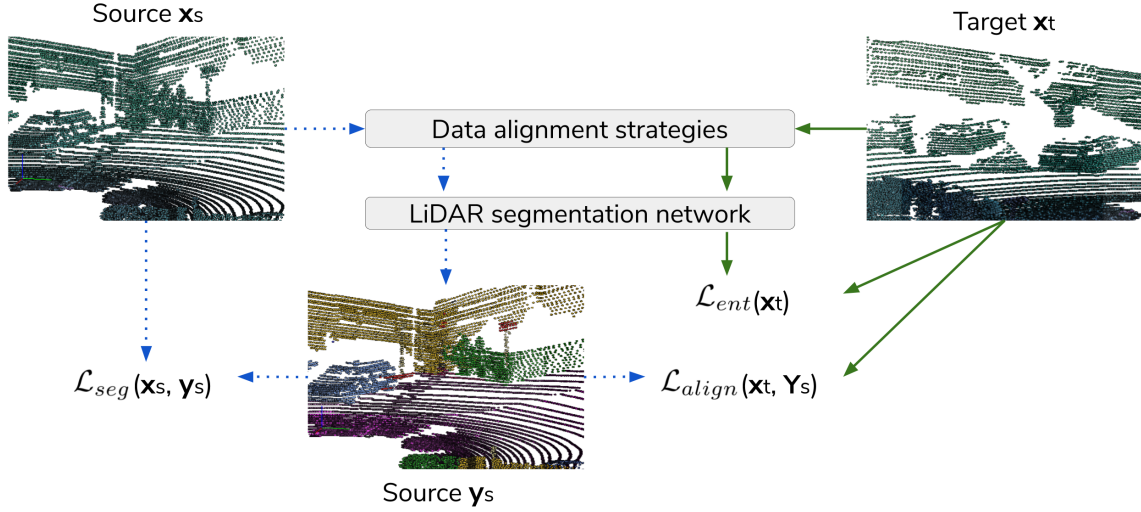


Figure 2: **Approach overview.** The figure shows our pipeline steps and optimization losses. First, we perform distribution alignment on the input space, i.e., data alignment strategies. Then, we optimize the segmentation loss for source samples where the labels are known and the class alignment and entropy losses for target data where no labels are available (See Sect. 3 for details). Green continuous arrows are used for target data and blue dotted arrows for source data.

be independent of absolute coordinates that are less robust compared to relative coordinates, we propose to use only relative features of the data. Therefore, we propose to use only relative distances of every point with respect to their neighbors.

3.3 Class Distribution Alignment

The domain shift appears due to many different factors. For example, different environments can present quite different appearances, the spatial distribution of objects may vary, the capturing set-up for different scenarios can be totally different, etc. Depending on the problem tackled and prior knowledge, we can hypothesize which of these differences can be neglected and assumed not to affect to the models we are learning. In this work, we tackle the domain adaptation problem between different urban environments, since all the datasets used are from urban scenarios. Taking this into account, although the data distribution changes between the datasets, we can assume that the class distribution is going to be very similar across these scenarios since the distribution shift in urban environments is mostly due to appearance and spatial distribution changes. For example, we can assume that if \mathbf{y}_s has a distribution of 90% road pixels and 10% car pixels, then \mathbf{y}_t will likely present a similar distribution.

Our approach learns parameters θ_F of F in such a way that the predicted class distribution $F(\mathcal{X}_t)$ matches the real class distribution of \mathcal{Y}_s , i.e., the histogram representing the frequency of occurrence of each class, previously computed in an offline fashion.

To do so, We propose to compute the KL-divergence between these two class distribution as the class alignment loss

$$\mathcal{L}_{align}(\mathbf{x}_t, \mathcal{Y}_s) = \sum_{n=1}^N \mathbf{hist}(\mathcal{Y}_s) \log \frac{\mathbf{hist}(\mathcal{Y}_s)}{\mathbf{P}_{\mathbf{x}_t}^{(n)}}. \quad (3)$$

Therefore, the optimization problem reads as:

$$\min_{\theta_F} \frac{1}{|\mathcal{X}_t|} \sum_{\mathbf{x}_t \in \mathcal{X}_t} \mathcal{L}_{align}(\mathbf{x}_t, \mathcal{Y}_s). \quad (4)$$

Equation 2 requires to compute the class distribution $\mathbf{P}_{\mathbf{x}_t}$ over the whole dataset. As this is computationally unfeasible, we compute it over the batch as an approximation.

3.4 Optimization Formulation

The entropy loss is computed as in MinEnt (Vu et al., 2019):

$$\mathcal{L}_{ent}(\mathbf{x}_t) = \frac{-1}{\log(C)} \sum_{n=1}^N \sum_{c=1}^C \mathbf{P}_{\mathbf{x}_t}^{(n,c)} \log \mathbf{P}_{\mathbf{x}_t}^{(n,c)}, \quad (5)$$

while the segmentation \mathcal{L}_{seg} and class alignment \mathcal{L}_{align} losses are computed as detailed in previous subsections.

During training, we jointly optimize the supervised segmentation loss \mathcal{L}_{seg} on source samples and the class alignment loss \mathcal{L}_{align} and entropy loss \mathcal{L}_{ent} on target samples. The final optimization problem is formulated as follows:

$$\begin{aligned} \min_{\theta_F} & \frac{1}{|\mathcal{X}_s|} \sum_{\mathbf{x}_s} \mathcal{L}_{seg}(\mathbf{x}_s, \mathbf{y}_s) + \frac{\lambda_{align}}{|\mathcal{X}_t|} \sum_{\mathbf{x}_t} \mathcal{L}_{align}(\mathbf{x}_t, \mathcal{Y}_s) \\ & + \frac{\lambda_{ent}}{|\mathcal{X}_t|} \sum_{\mathbf{x}_t} \mathcal{L}_{ent}(\mathbf{x}_t), \end{aligned} \quad (6)$$

with λ_{ent} and λ_{align} as the weighting factors of the alignment and entropy terms.

4 EXPERIMENTAL SETUP

This section details the setup used in our evaluation. This includes the datasets used in the evaluation and the training protocol we followed.

4.1 Datasets

We use four different datasets for the evaluation. They were collected in four different geographical areas, with four different LiDAR sensors, and with four different set-ups. We take the well known SemanticKITTI dataset (Behley et al., 2019) as the source domain dataset and the other three datasets as target domain data.

SemanticKITTI. The SemanticKITTI dataset (Behley et al., 2019) is a recent large-scale dataset that provides dense point-wise annotations. The dataset consists of over 43000 LiDAR scans from which over 21000 are available for training. The dataset distinguishes 22 different semantic classes. The capturing sensor is a Velodyne HDL-64E mounted on a car.

Paris-Lille-3D. Paris-Lille-3D (Roynard et al., 2018) is a medium-size dataset that provides three aggregated point clouds. It is collected with a tilted rear-mounted Velodyne HDL-32E placed on a vehicle. Following PolarNet work (Zhang et al., 2020), we extract individual scans from the registered point clouds thanks to the scanner trajectory and points' timestamps. Each scan is made of points within +/-100m. We take the Lille-1 point cloud for the domain adaptation and Lille-2 for validation. We use the following intersecting semantic classes with the SemanticKitti: car, person, road, sidewalk, building, vegetation, pole, and traffic light.

SemanticPoss. The SemanticPoss (Pan et al., 2020) is a medium-size dataset which contains 5 different sequences from urban scenarios providing 3000 LiDAR scans. The sensor used is a 40-line Pandora

mounted on a vehicle. We take the three first sequences for applying the adaptation methods and the last two sequences for validation. We use the following intersecting classes with the SemanticKitti: car, person, trunk, vegetation, traffic sign, pole, fence, building, rider, bike, and ground (combines road and sidewalk).

I3A. We have captured a small dataset to test our approach in a different scenario. In contrast to the three previous datasets, this dataset is not captured from a vehicle but from a TurtleBot, having the sensor at a lower height than in the other set-ups. The capturing sensor is the Velodyne VLP-16. We capture an urban environment similar to previous datasets. The dataset contains two sequences, one for training and another for validation. We use the intersecting semantic classes with the SemanticKitti: car, person, road, sidewalk, building, vegetation, trunk, pole, and traffic light.

4.2 Training Protocol

As we mentioned in Sec. 3.1, we use 3D-MiniNet (Alonso et al., 2020) as the base LiDAR semantic segmentation method. In particular, we use the available 3D-MiniNet-small version because of memory issues. For computing the relative coordinates and features, we follow 3D-MiniNet approach extracting them from the N neighbors of each 3D point where N is set to 16.

For all the experiments we train this architecture for 700K iterations with a batch size of 8. We use Stochastic Gradient Descent (SGD) optimizer with an initial learning rate of 0.005 and a polynomial learning rate decay schedule with a power set to 0.9.

We set λ_{ent} to 0.001 as suggested in MinEnt (Vu et al., 2019) and λ_{align} to 0.001. We empirically noticed that the performance is very similar when these two hyper-parameters are set between 10^{-5} and 10^{-2} . The two main conditions for them to work properly are: (1) be greater than 0 and, (2) do not be higher than the supervised loss.

One thing to take into account is that, as explained in 4.1, the Paris-Lille-3D has a very limited field of view. Therefore in order to make MinEnt (Vu et al., 2019) work in this dataset, we had to simulate the same field of view on the source dataset.

5 RESULTS

This section presents the experimental validation of our approach compared to different baselines. The

Table 1: Ablation study of our domain adaption pipeline for semantic segmentation. Source dataset: SemanticKitti (Behley et al., 2019).

Target dataset	mIoU on I3A	mIoU on ParisLille	mIoU on SemanticPoss
Base model	15.9	19.2	14.0
+ XYZ-shift augmentation	25.1	28.9	16.8
+ Per-class augmentation	27.0	30.1	18.2
+ Same number of beams	42.0	35.4	19.7
+ Only relative features	47.1	—	21.8
+ MinEnt (Vu et al., 2019)	50.3	41.5	26.2
+ Class distrib. alignment	52.5	42.7	27.0

— Not used because there was no performance gain.

proposed approach achieves better results than the other baselines in the three different scenarios for unsupervised domain adaptation in LiDAR Semantic Segmentation. In all the experiments we use the SemanticKITTI dataset (Behley et al., 2019) as the source data distribution and perform the adaptation on the other three datasets.

5.1 Ablation Study

The experiments in this subsection show how the different data alignment steps and the proposed learning losses affect the final performance of our approach. Table 1 summarizes the ablation study performed on three different scenarios. The results show how all the steps proposed contribute towards the final performance. The main insights observed in the ablation study are discussed next.

Performing strong *XYZ-shifts* results in a boost on the performance, meaning that the domain gap is considerably reduced. The distribution gap reduced by this step is the one caused by the fact of using different LiDAR sensor set-ups (such as different acquisition sensor height). Besides, in these autonomous driving set-ups, the distance between the car and the objects depends on how wide are the streets or on which lane is the data capturing source. Therefore, this is an essential and really easy data transformation to perform which gives an average of 7.2% mIoU gain.

The *per-class data augmentation* also boosts the performance. This data augmentation method tries to reduce the domain gap by adding different relative distances between different classes gaining an average of 1.3% mIoU gain.

Another interesting and straightforward technique to perform is to *match the number of LiDAR beams* of the source and target data, i.e., match LiDAR point-cloud resolution. This helps the data alignment especially for having the same point density on the 3D point-cloud and similar relative distances between the points. We show that this method gives an improvement similar to the XYZ-shift data augmentation, hugely reducing the domain gap. The higher the

initial difference in the number of beams, the more improvement we can get: the i3A LiDAR has 16 beams, the ParisLille 32, and the SemanticPoss 40, compared to the 64 of the source data (SemanticKitti).

The use of *relative features only* does not always help to reduce the domain gap, it was only beneficial on the i3A and SemanticPoss datasets. Removing the absolute features and only learning from relative features helps especially when the relative distances between the 3D points have less domain shift than the absolute coordinates. This will depend on the dataset, but the stronger the differences between capturing sensors, the more likely that the use of relative features will help.

Besides the data alignment steps, our approach includes two *learning losses* to the pipeline to help to reduce the domain gap. The first one is the entropy minimization loss proposed in MinEnt (Vu et al., 2019). The second one is our proposed class distribution alignment loss introduced in this work. We show that these two losses can be combined for the domain adaptation problem and that, although less significantly with respect to previously discussed steps, they also improve on the three different set-ups, contributing to achieving state-of-the-art performance.

5.2 Comparison with Other Baselines

Table 2: Results on the three different LiDAR semantic segmentation datasets using different domain adaptation methods. The source dataset is the SemanticKitti dataset (Behley et al., 2019).

	mIoU on I3A	mIoU on ParisLille	mIoU on SemanticPoss
Baseline	15.9	19.2	13.4
MinEnt (Vu et al., 2019)	28.4	23.2	19.6
AdvEnt (Vu et al., 2019)	21.0	20.7	19.5
MaxSquare (Chen et al., 2019)	28.4	22.8	19.3
Data alignment (ours)*	47.1	36.2	19.0
Full approach (ours)	52.5	42.7	27.0

* Only data alignment strategies from Sect. 3.2

Table 2 shows the comparison of our pipeline (composed of all the steps discussed in the ablation study) with other existing methods for domain adaptation. We select MinEnt, Advent, and MaxSquare as the baselines because they are leading the state-of-the-art for unsupervised domain adaptation. We use the available authors' code for replication.

We apply the different domain adaptation methods of the three different set-ups without our data alignment steps. This comparison shows that good pre-processing of the data can obtain better results than just applying out-of-the-box methods for domain adaptation. It also shows that our complete pipeline outperforms these previous methods on LiDAR domain adaptation. Our results demonstrate that com-

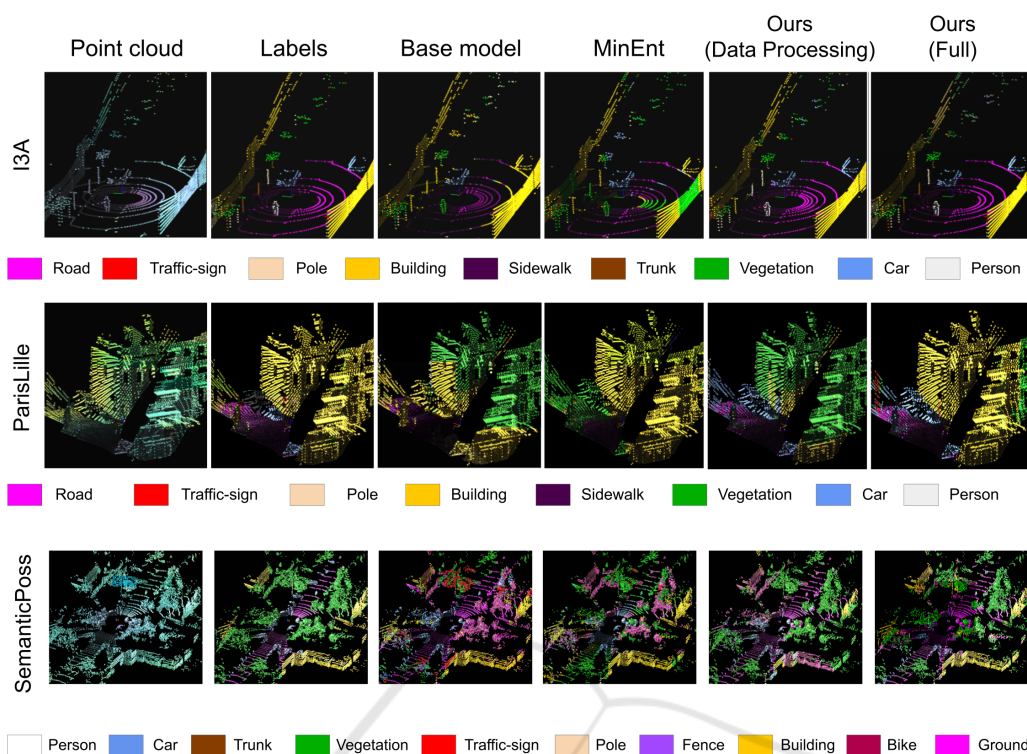


Figure 3: LiDAR domain adaptation results for different methods and datasets: I3A dataset first row, ParisLille dataset second row, and SemanticPois last row. From left to right: Input point cloud, ground truth labels, baseline with no adaptation (trained on SemanticKitti), MinEnt (Vu et al., 2019), our adaptation only with data processing strategies, our full adaptation pipeline. Best viewed in color. For better and more detail (video result) go to https://youtu.be/EpkJ_UH1F-o.

binning proper data processing with learning methods for domain adaptation gives an average of more than $\times 2$ boost on the performance.

Figure 3 includes a few examples of the segmentation obtained with a baseline with no domain adaptation, using the MinEnt (Vu et al., 2019) approach only, with our approach using only the data pre-processing steps, and with our approach including all steps proposed. We can appreciate in figure 3 how data processing helps on certain semantic classes, such as road, person, car, or vegetation, while MinEnt usually improves at different ones like building. This suggests the good complementary of both strategies, and indeed combining them provides the best results. More detail in additional video results can be seen at https://youtu.be/EpkJ_UH1F-o.

with different simple strategies combined with learning losses on the semantic segmentation process that also force the data distribution alignment. Our results show that a proper data alignment on the input space can produce better domain adaptation results that just using out-of-the-box state-of-the-art learning methods. Besides, we show that combining these data alignment methods with learning methods, like the one proposed in this work to align the class distributions of the data, can reduce even more the domain gap getting new state-of-the-art results. Our approach is validated on three different scenarios, from different datasets, as the target domain, where we show that our full pipeline improves previous methods on all three scenarios.

6 CONCLUSIONS

In this work, we introduce a novel pipeline that addresses the task of unsupervised domain adaptation for LiDAR semantic segmentation. Our pipeline consists of aligning data distributions on the data space

ACKNOWLEDGMENTS

This project was partially funded by projects FEDER/ Ministerio de Ciencia, Innovación y Universidades/ Agencia Estatal de Investigación/RTC-2017-6421-7, PGC2018-098817-A-I00 and PID2019-105390RB-I00, Aragón regional government (DGA T45

17R/FSE) and the Office of Naval Research Global project ONRG-NICOP-N62909-19-1-2027.

REFERENCES

- Alonso, I., Riazuelo, L., Montesano, L., and Murillo, A. C. (2020). 3d-mininet: Learning a 2d representation from point clouds for fast and efficient 3d lidar semantic segmentation. *IEEE Robotics and Automation Society, RAL*.
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., and Gall, J. (2019). SemanticKITTI: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Chen, M., Xue, H., and Cai, D. (2019). Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2090–2099.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2016). Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., and Darrell, T. (2018). Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998.
- Jaritz, M., Vu, T.-H., Charette, R. d., Wirbel, E., and Pérez, P. (2020). xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12605–12614.
- Li, Y., Yuan, L., and Vasconcelos, N. (2019). Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6936–6945.
- Mei, J., Gao, B., Xu, D., Yao, W., Zhao, X., and Zhao, H. (2019). Semantic segmentation of 3d lidar data in dynamic scene using semi-supervised learning. *IEEE Transactions on Intelligent Transportation Systems*, 21(6):2496–2509.
- Milioto, A., Vizzo, I., Behley, J., and Stachniss, C. (2019). Rangenet++: Fast and accurate lidar semantic segmentation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Pan, Y., Gao, B., Mei, J., Geng, S., Li, C., and Zhao, H. (2020). Semanticpos: A point cloud dataset with large quantity of dynamic instances. *arXiv preprint arXiv:2002.09147*.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2017a). Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660.
- Qi, C. R., Yi, L., Su, H., and Guibas, L. J. (2017b). Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108.
- Rosu, R. A., Schütt, P., Quenzel, J., and Behnke, S. (2019). Latticenet: Fast point cloud segmentation using permutohedral lattices. *arXiv preprint arXiv:1912.05905*.
- Roynard, X., Deschaud, J.-E., and Goulette, F. (2018). Paris-lille-3d: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification. *The International Journal of Robotics Research*, 37(6):545–557.
- Sun, B., Feng, J., and Saenko, K. (2017). Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 153–171. Springer.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. (2017). Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176.
- Vu, T.-H., Jain, H., Bucher, M., Cord, M., and Pérez, P. (2019). Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2517–2526.
- Wang, M. and Deng, W. (2018). Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153.
- Wu, B., Wan, A., Yue, X., and Keutzer, K. (2018). Squeeze-seg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1887–1893. IEEE.
- Wu, B., Zhou, X., Zhao, S., Yue, X., and Keutzer, K. (2019). Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4376–4382. IEEE.
- Xie, Y., Tian, J., and Zhu, X. X. (2019). A review of point cloud semantic segmentation. *arXiv preprint arXiv:1908.08854*.
- Yang, Y. and Soatto, S. (2020). Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095.
- Zhang, Y., Zhou, Z., David, P., Yue, X., Xi, Z., Gong, B., and Foroosh, H. (2020). Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9610.
- Zhou, Y. and Tuzel, O. (2018). Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499.
- Zou, Y., Yu, Z., Liu, X., Kumar, B., and Wang, J. (2019). Confidence regularized self-training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5982–5991.