

MoPA: Multi-Modal Prior Aided Domain Adaptation for 3D Semantic Segmentation

Haozhi Cao¹, Yuecong Xu², Jianfei Yang³, Pengyu Yin¹, Shenghai Yuan¹, Lihua Xie¹, *Fellow, IEEE*

Abstract—Multi-modal unsupervised domain adaptation (MM-UDA) for 3D semantic segmentation is a practical solution to embed semantic understanding in autonomous systems without expensive point-wise annotations. While previous MM-UDA methods can achieve overall improvement, they suffer from significant class-imbalanced performance, restricting their adoption in real applications. This imbalanced performance is mainly caused by: 1) self-training with imbalanced data and 2) the lack of pixel-wise 2D supervision signals. In this work, we propose Multi-modal Prior Aided (MoPA) domain adaptation to improve the performance of rare objects. Specifically, we develop Valid Ground-based Insertion (VGI) to rectify the imbalanced supervision signals by inserting prior rare objects collected from the wild while avoiding introducing artificial artifacts that lead to trivial solutions. Meanwhile, our SAM consistency loss leverages the 2D prior semantic masks from SAM as pixel-wise supervision signals to encourage consistent predictions for each object in the semantic mask. The knowledge learned from modal-specific prior is then shared across modalities to achieve better rare object segmentation. Extensive experiments show that our method achieves state-of-the-art performance on the challenging MM-UDA benchmark. Code will be available at <https://github.com/AronCao49/MoPA>.

I. INTRODUCTION

3D scene understanding plays an important role in various robotics tasks, such as localization [1], [2], [3], [4] and path planning [5], [6], [7]. For intelligent robotics (e.g., autonomous driving and navigation), 3D semantic information about the surrounding environments is a must and therefore has attracted increasing attention from the community. While previous methods [8], [9], [10] have achieved outstanding performance based on fully supervised learning, they require expensive annotations and suffer from performance degeneracy when used in unseen environments.

To alleviate these limitations, Unsupervised Domain Adaptation (UDA) [13], [14], [15] has been proposed to leverage the knowledge from the labeled source domain (e.g., public datasets with delicate annotations) to the unlabeled target domain (e.g., specific application environments) for 3D semantic segmentation. With an increasing number of multi-modal datasets [11], [16], [17], Multi-Modal UDA (MM-

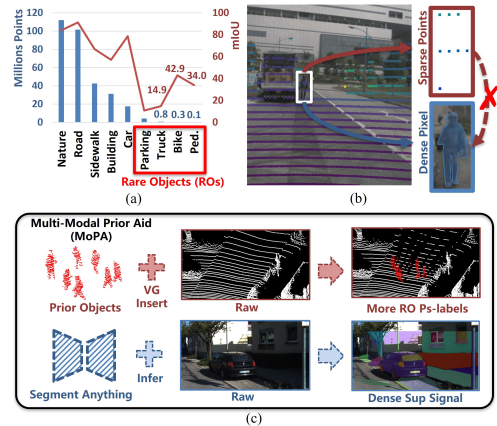


Fig. 1. Illustration of factors that lead to unsatisfactory Rare Objects (ROs) performances (a), (b) and our methods (c). (a) shows the unsatisfactory performances of ROs in a long-tailed dataset (SemanticKITTI [11]). (b) illustrates the limited supervision signals provided by point-wise pseudo-labels (only 8 points) to guide the 2D networks to identify the pedestrian in the dense image. (c) demonstrates the overall structure of our proposed MoPA, which introduces more prior ROs through VGI and dense supervision signals guided by Segment Anything Model (SAM) [12]

UDA) has been proposed to leverage the complementary information of 2D images and 3D point clouds for segmentation. As the primary work, xMUDA [18] proposes to encourage cross-modal prediction consistency. The following works further refine this framework from different perspectives (e.g., diversifying pixel-point corresponding [19], [20] and reducing domain discrepancy [21], [22]). While achieving noticeable improvement, previous works barely reveal and discuss their class-wise performance. To dig deeper into previous MM-UDA methods, we illustrate the class-wise performance of exemplary xMUDA [18] in Fig. 1(a). One can easily observe that xMUDA suffers from serious class-imbalanced performance, whose main bottleneck lies in the object classes (e.g., “Truck”, “Bike”, and “Pedestrian”). These classes, however, usually play a more important role in applications involving human-robot interaction (e.g., trajectory-aware path planning and obstacle avoidance).

Motivated by this observation, we aim to study how to alleviate the class-imbalanced performance in this work. Most existing MM-UDA methods [18], [19], [22] rely on self-training with pseudo-labels (i.e., confident predictions generated by the network itself) to compensate for the lack of ground-truth supervision signals, whereas this strategy is known to be ineffective for Rare Objects (ROs) in a long-tailed dataset as in Fig. 1(a). Without effective rectification for this class-wise imbalance, the pseudo-labels generated are biased toward the dominant classes, resulting in a scarcity of true-positive pseudo-labels for ROs. Moreover, since

¹Authors are with the Centre for Advanced Robotics Technology Innovation (CARTIN), School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. {haozhi002, pengyu001, shyuan, elhxie}@ntu.edu.sg

²Author is with the Department of Electrical and Computer Engineering, National University of Singapore. xuyu0014@ntu.edu.sg

³Author is with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. yang0478@ntu.edu.sg

This research is supported by the National Research Foundation, Singapore under its Medium Sized Center for Advanced Robotics Technology Innovation.

previous methods adopt a similar one-to-one pixel-point corresponding, the pseudo-labels generated are point-wise instead of pixel-wise. In terms of ROs with small LiDAR cross sections (e.g., “Bike” and “Pedestrian”) as shown in Fig. 1(b), the point-wise pseudo-labels are too sparse to supervise the dense pixel-wise predictions (e.g., 7 points vs. hundreds of pixels), resulting in inferior RO performances.

To this end, we propose an MM-UDA method called **Multi-modal Prior Aided (MoPA)** that boosts the ROs segmentation performance by addressing the aforementioned limitations as shown in Fig. 1(c). To rectify the imbalanced pseudo-labels, MoPA randomly inserts prior ROs with accurate labels collected from the wild (i.e., 3D prior knowledge) in each raw scan. Considering that directly inserting ROs can compromise the raw scan and lead to trivial solutions, Valid Ground-based Insertion (VGI) is developed to ensure the effectiveness of inserted prior ROs through validity checking and ray-tracing-based style translation. To compensate for the lack of dense 2D supervision signals, MoPA leverages the semantic masks from the foundation model Segment Anything Model (SAM) [12] as the 2D prior knowledge with reliable pixel-wise supervision signals. Combined with cross-modal learning, the multi-modal prior knowledge can be propagated across modalities to achieve better RO segmentation.

In summary, our contributions are listed as follows:

- We propose an MM-UDA method called MoPA, leveraging multi-modal prior to improve RO segmentation.
- VGI is developed to introduce more reliable RO pseudo-labels with validity guarantee and style translation. Meanwhile, SAM consistency is proposed to compensate for the lack of 2D dense supervision signals.
- Extensive experiments show that our MoPA achieves state-of-the-art performance on the challenging MM-UDA benchmark for 3D semantic segmentation.

II. RELATED WORK

Uni-Modal Domain Adaptation for Segmentation. Previous UDA methods for semantic segmentation can be divided as follows: (i) domain adversarial discriminative approaches [23], [24], (ii) generative approaches [25], [26], (iii) self-training approaches [27], [28], and (iv) uncertainty minimization approaches [29], [30]. In terms of UDA for 3D segmentation, besides those mostly inspired by the 2D UDA methods [31], [32], some previous works develop UDA methods to address some specific point-cloud-only discrepancy (e.g., the difference of point density [33], [34]). However, these methods are designed for uni-modal input and thus can not fully leverage the multi-modal information.

Multi-Modal Domain Adaptation for Segmentation. Thanks to the emerging adaptation of multi-modal sensors and multi-modal datasets [11], [35], [16], [17], domain adaptation with multi-modal input [18], [19], [21], [22], [20], [36] has attracted increasing attention. xMUDA [18] proposes the first MM-UDA method by encouraging cross-modal prediction consistency. DsCML [19] interprets the point-pixel-correspondence in a deformable manner and extends the cross-modal learning to an inter-domain perspective. The

following works follow this cross-modal-consistency strategy and further improve the segmentation performance by introducing additional components (e.g., adversarial learning [21], knowledge distillation [22], and contrastive learning [20]). Despite their consistent improvement in overall performance, none of them discuss the imbalanced class-wise performance of ROs, whose detection, however, can play an important role in other downstream tasks such as navigation and autonomous driving. To fulfill this limitation, we specifically study how to improve the performance of ROs for MM-UDA.

Mixing-based Strategies. Sample mixing has been widely adopted in the image domain [37], [38], [39], which has shown its effectiveness in reducing domain discrepancy. Inspired by this idea, some recent works have proposed different mixing strategies for point cloud classification [40], [41] or segmentation [42], [43], [44]. For segmentation, CosMix [42] concatenates the confident semantic point cloud from the source domain to the target domain and vice versa. Hasecke *et al.* [43] proposes to regenerate point clouds to mitigate the domain gap caused by different sensor setups. By projecting point clouds to range images, ConDA [44] constructs an intermediate domain through the concatenation of range images. While performance improvement can be achieved, they also introduce some artificial artifacts during the sample mixing of point clouds. In this work, we propose a Valid Ground-based Insertion method that aims to improve the RO class-wise performance for MM-UDA without introducing artificial artifacts.

III. METHODOLOGY

A. Preliminary and Overview

In MM-UDA, the labeled source domain is denoted as $\mathcal{S} = \{(\mathbf{x}_{\mathcal{S},i}^{2D}, \mathbf{x}_{\mathcal{S},i}^{3D}, \mathbf{y}_{\mathcal{S},i})\}_{i=1}^{|\mathcal{S}|}$, where $\mathbf{x}_{\mathcal{S},i}^{2D} \in \mathbb{R}^{H \times W \times 3}$ and $\mathbf{x}_{\mathcal{S},i}^{3D} \in \mathbb{R}^{N \times 3}$ denote the 2D RGB image and 3D point cloud (i.e., input channels include x, y, and z). $\mathbf{y}_{\mathcal{S},i} \in \mathbb{R}^{N \times C}$ is the one-hot point-wise semantic label, where i and C denotes the sample index and the number of classes, respectively. The goal is to leverage the knowledge learned from \mathcal{S} to an unlabeled target domain $\mathcal{T} = \{(\mathbf{x}_{\mathcal{T},j}^{2D}, \mathbf{x}_{\mathcal{T},j}^{3D})\}_{j=1}^{|\mathcal{T}|}$. Following previous methods [18], [19], we introduce modal-specific networks $\phi_{\theta}^m(\cdot)$, $m \in \{2D, 3D\}$ to process 2D and 3D input independently. Given the model-specific prediction $\mathbf{p}_{\mathcal{D}}^m = \phi_{\theta}^m(\mathbf{x}_{\mathcal{D}}^m)$, $\mathcal{D} \in \{\mathcal{S}, \mathcal{T}\}$, the cross-modal prediction $\mathbf{p}_{\mathcal{D}}^{\text{XM}}$ is computed as their average following the same point-to-pixel projection protocol as in [18], [19]. In the following sections, the subscript of the sample index is omitted for clarity.

As shown in Fig 1(a), a significant bottleneck of current MM-UDA methods is the unsatisfactory performance of ROs due to (i) limited accurate pseudo-labels from RO classes and (ii) the lack of effective RO supervision signals on the 2D image. Motivated by such observation, our MoPA is designed to address the aforementioned gaps by fully leveraging multi-modal prior knowledge. Given an RO pool as 3D prior knowledge, we propose to insert more prior ROs with accurate labels through Valid Ground-based Insertion (VGI). (Sec. III-B). Meanwhile, we leverage the 2D prior




Fig. 2. Overview of our proposed MoPA operating on the target domain. Given a point cloud scan as input, prior objects are randomly sampled from the prior object pool and inserted into the raw input scan (Sec. III-B) to introduce more ROs with accurate pseudo-labels. Meanwhile, the prior 2D semantic mask is generated from the foundation model SAM [12] to serve as our dense 2D supervision signals (Sec. III-C). Additionally, the EMA module generates online pseudo-labels and facilitates cross-modal learning by introducing cross-modal pseudo-labels. Combined with xMUDA [18] multi-modal learning pipeline, our MoPA can effectively improve the segmentation performance on ROs by propagating the prior knowledge in a cross-modal manner.

knowledge from foundation models to compensate for the lack of dense ROs supervision signals (Sec. III-C). Fig. 2 further illustrates the overall pipeline of MoPA (Sec. III-D).

B. Valid Ground-based Insertion (3D)

Point cloud mixing [42], [44] has been known as an effective method for domain adaptation, which mitigates the domain discrepancy by constructing an intermediate domain. Different from previous works, our goal is to introduce more prior ROs with accurate pseudo-labels and propagate the knowledge learned from prior ROs to those in the original scan. To achieve this goal, VGI is designed to insert prior ROs in the input scan with validity checking and style translation to avoid trivial solutions.

1) *Preparation of the Object Pool*: Before insertion, an object pool is collected from the wild. Given an arbitrary scan of the object source, we preserve the points with the semantic labels of interest and utilize DBSCAN [45] to cluster the object instances, leading to an object pool denoted as:

$$\mathcal{Q} = \{(\mathbf{x}_{q,k}, \mathbf{y}_{q,k})\}_{k=1}^{|\mathcal{Q}|}, \quad (1)$$

where $\mathbf{x}_{q,k}$ and $\mathbf{y}_{q,k}$ indicate the 3D points and the semantic label of instance k .

2) *Overlap Checking and Grounding*: Given any raw input scan $\hat{\mathbf{x}}_{\mathcal{T}}^{3D}$, VGI randomly samples object instances from the object pool and inserts them in the target scan. Unlike previous point cloud concatenation methods [42], [44], which directly mix point clouds without validity checking, our insertion process follows two main steps to avoid introducing artificial artifacts: (i) voxel-based overlap checking and (ii) object grounding. For voxel-based overlap checking, the valid location candidates are generated as the inserted object center locations that do not compromise the raw input scan. Specifically, given N_q voxelized prior objects sampled from \mathcal{Q} denoted as $\hat{\mathcal{Q}}' = \{\hat{\mathbf{x}}_{q,n}, \hat{\mathbf{y}}_{q,n}\}_{n=1}^{N_q}$, the query of overlap checking $\hat{\mathbf{q}} = [\hat{e}_x, \hat{e}_y, \hat{e}_z]$ is the 3D extent of the minimum axis-aligned bounding box that includes all prior objects $\hat{\mathbf{x}}_{q,n} \in \hat{\mathcal{Q}}'$. Subsequently, the overlap checking is conducted as the grid search between the voxel grids of the query $\hat{\mathbf{q}}$ and the searching area in the voxelized input scan $\hat{\mathbf{x}}_{\mathcal{T}}^{3D}$:

$$\hat{\mathbf{Q}} \in \mathbb{R}^{\hat{e}_x \times \hat{e}_y \times \hat{e}_z}, \hat{\mathbf{Q}}_{ijk} = 1, \quad (2)$$

$$\hat{\mathbf{K}} \in \mathbb{R}^{\Delta x \times \Delta y \times \Delta z}, \hat{\mathbf{K}}_{ijk} = \begin{cases} 1, & \text{if occupied by } \hat{\mathbf{x}}_{\mathcal{T}}^{3D} \\ 0, & \text{otherwise} \end{cases}, \quad (3)$$

$$\hat{\mathbf{V}} = \{(\hat{x}_i, \hat{y}_i, \hat{z}_i) | (\hat{\mathbf{Q}} * \hat{\mathbf{K}})_{\hat{x}_i \hat{y}_i \hat{z}_i} = 0\}, \quad (4)$$

where $\Delta x, \Delta y, \Delta z$ are the corresponding axis-along extents of the searching area pre-defined by its corner coordinates (x_1, y_1, z_1) and (x_2, y_2, z_2) (i.e., $\Delta x = |x_1 - x_2|$, likewise for Δy and Δz). $\hat{\mathbf{Q}}$ and $\hat{\mathbf{K}}$ denote the voxel grids of the query and the searching area, while $*$ represents the 3D convolution computation. Based on the overlap checking results from Eq. 4, the voxelized valid location candidates $\hat{\mathbf{V}} = \{(\hat{x}_i^v, \hat{y}_i^v, \hat{z}_i^v)\}_{i=1}^{|\hat{\mathbf{V}}|}$ are then computed by mapping $\hat{\mathbf{V}}$ back to the input scan coordinate, at each of which the object insertion won't compromise the raw input scan.

Subsequently, the valid location candidates are filtered by only preserving those above the ground as shown in Fig. 3(a). This is based on the fact that objects are mostly grounded in urban environments. Such a filtration process can simultaneously discard the unreasonable valid location candidates introduced during the grid search (e.g., those behind the building walls). Specifically, given the raw input scan $\hat{\mathbf{x}}_{\mathcal{T}}^{3D}$, we leverage the geometric-based ground detection algorithm Patchwork++ [46] as our ground prior thanks to its real-time efficiency, whose detected ground points are similarly voxelized as $\hat{\mathcal{G}} = \{(\hat{x}_i^g, \hat{y}_i^g, \hat{z}_i^g)\}_{i=1}^{|\hat{\mathcal{G}}|}$. Given $\hat{\mathcal{G}}$ and $\hat{\mathbf{V}}$, the valid locations $\hat{\mathbf{V}}_g$ is preserved as follows:

$$\hat{\mathbf{V}}_g = \{(\hat{x}_i^v, \hat{y}_i^v, \hat{z}_i^v) \in \hat{\mathbf{V}} | \exists j, (\hat{x}_i^v, \hat{y}_i^v, \hat{z}_i^v - \hat{e}_z/2) = (\hat{x}_j^g, \hat{y}_j^g, \hat{z}_j^g) \in \hat{\mathcal{G}}\}, \quad (5)$$

where each insertion location of $\hat{\mathbf{V}}_g$ can be grounded without compromising the raw input scan. N insertion locations, denoted as $\mathcal{C}_{\mathcal{Q}'} = \{(x_i^c, y_i^c, z_i^c)\}_{i=1}^{N_q}$ are then sampled from $\hat{\mathbf{V}}_g$ for each object in \mathcal{Q}' . In practice, the altitude of each insertion location is further refined after de-voxelization to ensure the inserted object is grounded in the input scan.

3) *Insertion and Style-Translation*: After overlap checking and grounding, the prior object instances are inserted in the raw scan and post-processed for further refinement. Considering the shape of object instances captured by LiDAR sensors is orientation-specific, we preserve the orientation of the prior object instances to further avoid introducing




Fig. 3. The detailed steps in VGI. Given a prior RO $\mathbf{x}_{q,n}$ and the raw scan \mathbf{x}_T^{3D} , the valid insertion locations (highlighted in yellow) are first computed through (a) overlap checking and grounding, where the green points are the detected ground. (b) $\mathbf{x}_{q,n}$ is then inserted in the raw scan, whose style is translated by (c) occlusion removal and adaptive downsampling. Blue and red points are the preserved and discarded points in step (c), respectively. Best when zoomed in and viewed in color.

artificial artifacts during the insertion. Given the bounding box center of a prior object instance \mathbf{x}_q and its corresponding insertion location denoted as $\mathbf{o}_q = (x_q, y_q, z_q)$ and $\mathbf{o}_c = (x^c, y^c, z^c)$, $\mathbf{o}_c \in \mathcal{C}_{Q'}$, the radial translation and z-rotation matrices are formulated as:

$$\Delta\phi = \phi_c - \phi_q, \Delta\rho = \rho_c - \rho_q, \quad (6)$$

$$\mathbf{t} = [\Delta\rho \cos \phi_q \quad \Delta\rho \sin \phi_q \quad z_c - z_q], \mathbf{T} = \begin{bmatrix} \mathbf{I} & \mathbf{t}^T \\ \mathbf{0} & 1 \end{bmatrix}, \quad (7)$$

$$\mathbf{r} = \begin{bmatrix} \cos \Delta\phi & -\sin \Delta\phi \\ \sin \Delta\phi & \cos \Delta\phi \end{bmatrix}, \mathbf{R} = \begin{bmatrix} \mathbf{r} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (8)$$

where ϕ_q and ρ_q denote the azimuth and radial of the projected point of \mathbf{o}_q on the xy-plane (likewise for ϕ_c and ρ_c). With the radial translation matrix $\mathbf{T} \in \mathbb{R}^{4 \times 4}$ and z-rotation matrix $\mathbf{R} \in \mathbb{R}^{4 \times 4}$, the transformed prior object $\tilde{\mathbf{x}}_q = (\mathbf{R} \times \mathbf{T} \times \mathbf{x}_q^T)^T$ is inserted at \mathbf{o}_c .

Subsequently, our style-translation stage removes the insertion-induced artificial artifacts that could lead to trivial solutions. This post-processing process adaptively downsamples the point density of $\tilde{\mathbf{x}}_q$ to a similar level of its surroundings and synthesizes the occlusion introduced by the $\tilde{\mathbf{x}}_q$. Given the raw input scan \mathbf{x}_T^{3D} and any inserted prior object $\tilde{\mathbf{x}}_q$, their Range View (RV) image coordinates $\mathbf{x}_T^{RV} \in \mathbb{R}^{N \times 2}$ and $\tilde{\mathbf{x}}_q^{RV} \in \mathbb{R}^{M \times 2}$ are first computed by the Eq. (1) in [47]. For each unique RV coordinate and its corresponding 3D points, the style-translation process acts as an RV-based ray tracing operation, preserving the 3D point with the minimum range while discarding the others as:

$$(\mathbf{x}^{3D}, \mathbf{x}^{RV}, \tilde{\mathbf{y}}^{3D}) = (\mathbf{x}_T^{3D} \oplus \tilde{\mathbf{x}}_q, \mathbf{x}_T^{RV} \oplus \tilde{\mathbf{x}}_q^{RV}, \tilde{\mathbf{y}}_T^{3D} \oplus \tilde{\mathbf{y}}_q), \quad (9)$$

$$\mathbf{M}^{3D} = \Psi(\mathbf{x}^{3D}, \mathbf{x}^{RV}), \mathbf{M}^{3D} \in \mathbb{R}^{(N+M)}, \quad (10)$$

$$(\mathbf{x}_p^{3D}, \tilde{\mathbf{y}}_p^{3D}) = (\mathbf{M}^{3D} \odot \mathbf{x}^{3D}, \mathbf{M}^{3D} \odot \tilde{\mathbf{y}}^{3D}), \quad (11)$$

where \oplus denotes the concatenation along the point axis and $\tilde{\mathbf{y}}_T^{3D}$ denotes the pseudo-labels of the raw scan. Ψ represents the RV-based filtering operation, which preserves only one 3D point in \mathbf{x}^{3D} for each pixel in \mathbf{x}_p^{RV} and generates a valid mask \mathbf{M}^{3D} . \odot denotes the point-wise selection based on the valid mask. The generated $(\mathbf{x}_p^{3D}, \tilde{\mathbf{y}}_p^{3D})$ is then utilized as the additional input pair for the subsequent 3D training process.

C. SAM Mask Consistency (2D)

For ROs with relatively small LiDAR cross sections as in Fig. 1(b), the sparse supervision signals provided by the

3D point cloud cannot fully leverage the advantage of dense 2D input. To address this issue, we leverage the semantic mask from Segment Anything Model (SAM) [12] as 2D prior knowledge to regularize the dense 2D predictions. As a powerful fundamental model with zero-shot generalization, SAM is able to provide semantic masks for unseen images without re-training as shown in Fig. 2, albeit with unknown semantic categories. To leverage 2D prior masks from SAM, we propose a SAM mask consistency loss to encourage consistent pixel-wise predictions within each mask. Meanwhile, the entropy of the mask-wise mean prediction is minimized to encourage mask-wise confident predictions. Given a set of SAM masks $\mathbf{M}^{2D} = \{\mathbf{m}_i^{2D} \in \mathbb{R}^{H \times W}\}_{i=1}^{|\mathbf{M}^{2D}|}$ and pixel-wise 2D semantic predictions $\tilde{\mathbf{p}}^{2D} \in \mathbb{R}^{H \times W \times C}$ before the pixel-to-point projection, our SAM consistency loss is defined as:

$$\mathcal{L}_{SC}^T = \frac{1}{|\mathbf{M}^{2D}|} \sum_{i=1}^{|\mathbf{M}^{2D}|} (\text{MSE}(\tilde{\mathbf{p}}_{\mathbf{m}_i}^{2D}) + \sum_C \tilde{\mathbf{p}}_{\mathbf{m}_i}^{2D} \log \tilde{\mathbf{p}}_{\mathbf{m}_i}^{2D}), \quad (12)$$

where $\tilde{\mathbf{p}}_{\mathbf{m}_i}^{2D} = \mathbf{m}_i^{2D} \odot \tilde{\mathbf{p}}^{2D}$ are the pixels located in the mask \mathbf{m}_i^{2D} while $\text{MSE}(\cdot)$ denotes the mean-square-error computation. $\tilde{\mathbf{p}}_{\mathbf{m}_i}^{2D}$ is the mean predictions of pixels in the mask, which is regularized by the Shannon entropy minimization. In practice, \mathcal{L}_{SC} is computed for each image in the batch level and summarized before optimization.

D. Overall Pipeline

The overall pipeline of MoPA is illustrated in Fig. 2. While MoPA can be easily integrated with all existing MM-UDA methods, we leverage xMUDA [18] as our baseline due to its reproducibility. Specifically, when training on the source domain sample $(\mathbf{x}_S^{2D}, \mathbf{x}_S^{3D}, \mathbf{y}_S)$, our optimization objective follows xMUDA [18], formulated as:

$$\mathcal{L}_{CE}^S = \mathcal{F}(\mathbf{p}_S^{2D}, \mathbf{y}_S) + \mathcal{F}(\mathbf{p}_S^{3D}, \mathbf{y}_S), \quad (13)$$

$$\mathcal{L}_{xM}^S = D_{KL}(\mathbf{p}_S^{3D} \parallel \hat{\mathbf{p}}_S^{2D}) + D_{KL}(\mathbf{p}_S^{2D} \parallel \hat{\mathbf{p}}_S^{3D}), \quad (14)$$

where $\mathcal{F}(\cdot)$ denotes the point-wise cross-entropy function. $\hat{\mathbf{p}}_S^{3D}$ and $\hat{\mathbf{p}}_S^{2D}$ denote the prediction from the auxiliary 3D and 2D classifiers as in xMUDA, respectively. While training on the target domain, the optimization objectives of xMUDA are also preserved, including \mathcal{L}_{CE}^T and \mathcal{L}_{xM}^T similarly computed as Eq. 13 and 14 with \mathbf{y}_S replaced by pseudo-labels $\tilde{\mathbf{y}}_T$. Further, given the object-inserted sample $(\mathbf{x}_p^{3D}, \tilde{\mathbf{y}}_p^{3D})$ and the SAM consistency loss, the overall loss \mathcal{L} is formulated as:

$$\mathcal{L}_{V-CE}^T = \mathcal{F}(\phi_{\theta}^{3D}(\mathbf{x}_p^{3D}), \tilde{\mathbf{y}}_p^{3D}), \quad (15)$$

$$\mathcal{L} = \underbrace{\mathcal{L}_{CE}^S + \lambda_{xM}^S \mathcal{L}_{xM}^S + \mathcal{L}_{CE}^T + \lambda_{xM}^T \mathcal{L}_{xM}^T}_{\text{xMUDA}} + \underbrace{\lambda_{V-CE}^T \mathcal{L}_{V-CE}^T + \lambda_{SC}^T \mathcal{L}_{SC}^T}_{\text{MoPA}}, \quad (16)$$

where $\lambda_{xM}^S, \lambda_{xM}^T, \lambda_{V-CE}^T, \lambda_{SC}^T$ are the pre-defined coefficient for each loss components.

Considering the inferior RO segmentation of offline pseudo-labels, an exponential moving average (EMA) is introduced, which utilizes a slowly updated teacher to generate the pseudo-labels in an online manner. When enabling EMA, the teacher network $\phi_{\theta'}^m$ is initialized from the student network ϕ_{θ}^m and updated as the moving average of the

TABLE I

PERFORMANCE (mIoU) COMPARISON OF MoPA. METHODS WITH † ARE RE-IMPLEMENTATION BASED ON THEIR OFFICIAL CODE. THE BEST PERFORMANCE IS IN BOLD WHILE THE SECOND BEST IS UNDERLINED.

Modality	Method	Code	Δ Params	USA \rightarrow Singapore			Day \rightarrow Night			A2D2 \rightarrow KITTI		
				2D	3D	xM	2D	3D	xM	2D	3D	xM
-	Source only	-	+0.0	53.4	46.5	61.3	42.2	41.2	47.8	36.0	36.6	41.8
Uni-Modal	Deep logCORAL [48]	✓	+2.8M	52.6	47.1	59.1	41.4	42.8	51.8	35.8	39.3	40.3
	MinEnt [29]	✓	+28.3M	53.4	47.0	59.7	44.9	43.5	51.3	38.8	38.0	42.7
	PL [26]	✓	+0.0	55.5	51.8	61.5	43.7	45.1	48.6	37.4	44.8	47.7
	CLAN [49]	✓	+2.8M	57.8	51.2	62.5	45.6	43.7	49.2	39.2	44.7	44.5
Multi-Modal	xMUDA† [18]	✓	+0.8K	58.5	51.2	61.0	47.7	45.1	50.1	42.8	45.4	48.1
	xMUDA+PL† [18]	✓	+0.8K	60.0	51.9	62.3	47.9	45.8	50.9	45.4	50.1	50.4
	DsCML† [19]	✓	+0.6M	39.0	38.8	45.2	44.7	36.3	49.7	26.3	37.0	35.6
	DsCML+PL† [19]	✓	+0.6M	39.8	41.7	46.9	46.3	34.7	46.7	26.6	37.2	36.0
	AUDA [21]	✗	+0.2M	59.8	52.0	62.7	46.2	44.2	50.0	38.3	46.0	44.0
	AUDA+PL [21]	✗	+0.2M	61.9	54.8	<u>65.6</u>	50.3	<u>49.7</u>	52.6	46.8	48.1	50.6
	CMCL [20]	✗	+34.1K	62.0	54.2	64.5	49.0	46.6	51.6	42.8	47.7	48.0
	CMCL+PL [20]	✗	+34.1K	63.3	<u>57.1</u>	66.7	50.6	49.9	54.5	47.6	51.3	52.8
	MoPA+PL (Ours)	✓	+0.8K	60.6	57.0	63.5	<u>51.4</u>	47.5	53.4	<u>48.8</u>	<u>53.4</u>	53.5
	MoPA+PLx2 (Ours)	✓	+0.8K	61.7	57.2	64.1	52.2	48.8	<u>54.2</u>	49.3	55.1	54.3

student one for each iteration ($\theta'_i = \alpha\theta'_{i-1} + \theta_i$ where i indicate the training iteration). To facilitate cross-modal learning, the pseudo-labels from EMA for each modality are occasionally swapped to the cross-modal one \tilde{y}_T^{xM} derived from p_T^{xM} (i.e., $\tilde{y}_T^n = \tilde{y}_T^{xM}$ at a percentage of p_{xM}).

IV. EXPERIMENTAL RESULTS

A. Benchmarks and Settings

Benchmarks: following xMUDA [18], we evaluate our MoPA on 3 MM-UDA benchmarks: (i) USA-to-Singapore (U-to-S), (ii) Day-to-Night (D-to-N), and (iii) A2D2-to-SemanticKITTI (A-to-S). For U-to-S and D-to-N, nuScenes [16] is leveraged as their data source. Specifically, U-to-S differs in layout and infrastructure between the source and target domains while D-to-N contains images recorded in different illumination levels. Following prior works [18], points within each bounding box are assigned with their corresponding object label, where the others are labeled as background. For A-to-S, A2D2 [35] and SemanticKITTI [11] are utilized as the source and the target domains, where the domain discrepancy lies in the pattern and density differences of 3D point clouds. For a fair comparison, we strictly follow the organization of all benchmarks in prior MM-UDA methods [18], [19]. We utilize mean Intersection over Union (mIoU) as our evaluation metric same as [18], [19].

Implementation Details: following previous works [18], [19], [22], ResNet34 [50] and U-Net-based SparseConvNet [51] are leveraged as our 2D and 3D backbones. Since we utilize xMUDA [18] as our baseline, we strictly follow most hyper-parameters settings as xMUDA, including the optimizer, learning scheduler, and xMUDA-related coefficient settings. For VGI, we collect our prior objects from Waymo dataset [17] thanks to its data variety, where “Pedestrian”, “Bicycle”, and “Motorcycle” are the prior objects of interest for all benchmarks. Each class of interest in the object pool contains 1000 object instances at most and we randomly insert 1 prior object instance for each raw input scan, where a voxel size of 50 cm is used for overlap checking and grounding. For SAM masks, the ViT-h variant is utilized to generate 2D prior masks in an offline manner, where masks larger than 1/10 image size are filtered out for mask




Fig. 4. Processing time of MoPA per batch (left) and efficiency comparison with xMUDA and DsCML (right). Postfix “+PL” is omitted for clarity.

refinement. For MoPA coefficients, we empirically set λ_{V-CE}^T and λ_{SC}^T as 0.1 and 0.01. The online pseudo-labels are enabled at the 70k iteration (100k in total) and p_{xM} is set to 0.7. All experiments are conducted with PyTorch [52] on a single NVIDIA RTX 3090.

B. Comparison with SOTA Methods

Overall Performance. Table. I shows the performance comparison with previous MM-UDA methods. Leveraging the pseudo-labels generated by xMUDA, MoPA+PL achieves a noticeable improvement compared to the baseline xMUDA+PL on U-to-S and D-to-N. On A-to-S, this improvement significantly expands to more than 3% on all predictions on all single-modal (i.e., 2D or 3D) and cross-modal predictions (i.e., xM). Considering the biased pseudo-labels generated by xMUDA, we introduce the additional MoPA+PL $\times 2$ which leverages the pseudo-labels generated by MoPA+PL instead. This enables our MoPA+PL $\times 2$ to achieve competitive performance on U-to-S and D-to-N compared to the previous SOTA method CMCL+PL [20], with a minor relative drop of less than 2.4% but introducing 97.7% fewer trainable parameters. In terms of the more challenging A-to-S with double classes as others, our MoPA+PL $\times 2$ outperforms all previous methods by more than 1.5% on the cross-modal prediction. Please note that *we report the re-implementation results of both xMUDA [18] and DsCML [19] by their official code.*

Training Efficiency. To justify the efficiency of MoPA, we present the processing time and GPU RAM consumption of MoPA on A-to-S in Fig. 4. Specifically, despite the utilization of grid-search-based validity checking, thanks to voxelization and GPU acceleration, VGI and the addition network forward process lead to a minor 0.24s of time overhead where the VGI itself only takes about 0.1s. The main computational




Fig. 5. Qualitative results of cross-modal predictions from MoPA+PLx2 and xMUDA+PL. Here different colors represent different semantic classes. Our MoPA+PLx2 obviously outperforms xMUDA+PL when segmenting the challenging RO.

TABLE II
RO PERFORMANCE COMPARISON ON A-TO-S.

Method	Person			Bike		
	2D	3D	xM	2D	3D	xM
xMUDA	13.0	25.4	21.8	24.4	42.9	36.2
xMUDA+PL	25.2	29.8	34.0	32.7	47.1	42.9
MoPA+PL	26.3	33.6	41.1	40.9	53.1	54.0
MoPA+PLx2	42.2	44.3	53.8	44.4	57.6	55.2

TABLE III
ABLATION STUDY OF MoPA COMPONENTS.

Method	EMA	SAM	VGI	2D	3D	xM
xMUDA+PL	-	-	-	45.4	50.1	50.4
MoPA+PL	✓			42.7	51.6	47.8
		✓		45.6	52.1	51.3
			✓	45.3	51.3	51.1
	✓	✓		43.7	50.5	49.6
	✓		✓	45.2	51.9	51.5
		✓	✓	46.7	52.4	52.7
	✓	✓	✓	48.8	53.4	53.5

overhead comes from the online pseudo-labels with EMA, while we only enable it for the last 30% of the training process. Meanwhile, MoPA is more efficient in terms of GPU consumption compared to the existing method DsCML [19], which consumes 33.5% less GPU RAM yet achieves more than 35.8% improvement.

C. Ablation Studies and Analysis

Effect of MoPA on ROs. MoPA is designed to improve RO performance and we justify this by comparing MoPA with xMUDA and xMUDA+PL in Table. II. While both leveraging the pseudo-labels generated by xMUDA, our MoPA+PL noticeably outperforms xMUDA+PL on both ROs of our interest (note that “Motorcycle” is merged in “Bicycle” in A-to-S) on all single-modal or multi-modal predictions. Utilizing the rectified pseudo-labels from MoPA+PL, MoPA+PLx2 further expands this performance gap to more than 10% on all outputs for both classes of interest, which justifies the effectiveness of MoPA in improving RO performance.

Effect of Each Component of MoPA. Table. III illustrates our validation of each MoPA component on A-to-S. Specifically, utilizing only SAM or VGI results in an improvement of 0.9% and 0.7%, respectively, and their combination further enlarges the gap to 2.3%. Interestingly, utilizing only EMA update or EMA update with SAM results in an inferior performance compared to the baseline (−2.6 and −0.8%). This is mainly because utilizing the EMA update without effective regularization on pseudo-labels would result in a more biased performance. With VGI which regularizes pseudo-labels by inserting RO objects, the utilization of EMA update boosts the performance (+0.4% on pure VGI




Fig. 6. Comparison of insertion with VGI, without VGI, and with M-VGI (left) and the inserted RO training accuracy with or without VGI (right).

and +0.8% on VGI+SAM).

RO Insertion through VGI. Fig. 6 presents the comparison among insertion with VGI (w VGI), without VGI (w/o VGI), and with VGI but multiple objects (M-VGI). Compared with w/o VGI, w VGI achieves a notable improvement of more than 1.2% on all predictions. Based on the training accuracy of inserted RO from w VGI and w/o VGI, it can be seen that the insertion without VGI leads to more early saturation of recognizing inserted RO due to trivial solutions. This justifies our design of VGI to avoid introducing artificial artifacts during insertion. On the other hand, M-VGI might not be a better alternative for front-view-only cases compared to VGI since it further hampers the limited valid insertion locations, leading to less insertion variance for learning.

Qualitative results. Fig. 5 illustrates the qualitative results of our MoPA+PLx2 compared with our baseline method xMUDA+PL. Specifically, we visualize the cross-modal predictions of our MoPA+PLx2 and xMUDA+PL on samples of three different benchmarks. For D-to-N and U-to-S where xMUDA+PL struggles to segment out the motorcycle, our MoPA+PLx2 generates more accurate predictions with fewer true-negatives. For the more challenging scan from A-to-S, our MoPA+PLx2 clearly outperforms xMUDA+PL in segmenting the pedestrian and bicycle located far from the sensor. This justifies the effectiveness of our MoPA, improving the performance on challenging RO.

V. CONCLUSIONS

In this paper, we present a novel method called MoPA that leverages multi-modal prior knowledge for MM-UDA. By introducing more prior ROs with accurate pseudo-labels and learning from foundation models, MoPA can effectively improve RO performance. Extensive experiments justify the effectiveness of MoPA. The utilization of MoPA can potentially improve the performance of classes of interest in a flexible manner, which can bridge the gap between performance benchmarking and real-world application.

REFERENCES

- [1] T.-M. Nguyen, A. H. Zaini, C. Wang, K. Guo, and L. Xie, "Robust target-relative localization with ultra-wideband ranging and communication," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2312–2319.
- [2] X. Chen, A. Milioto, E. Palazzolo, P. Giguere, J. Behley, and C. Stachniss, "Suma++: Efficient lidar-based semantic slam," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4530–4537.
- [3] T. Ji, C. Wang, and L. Xie, "Towards real-time semantic rgb-d slam in dynamic environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 175–11 181.
- [4] P. Yin, S. Yuan, H. Cao, X. Ji, S. Zhang, and L. Xie, "Segregator: Global point cloud registration with semantic and geometric cues," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 2848–2854.
- [5] L. Bartolomei, L. Teixeira, and M. Chli, "Perception-aware path planning for uavs using semantic segmentation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5808–5815.
- [6] F. Stache, J. Westheider, F. Magistri, C. Stachniss, and M. Popović, "Adaptive path planning for uavs for multi-resolution semantic segmentation," *Robotics and Autonomous Systems*, vol. 159, p. 104288, 2023.
- [7] F. Liu, S. Yuan, W. Meng, R. Su, and L. Xie, "Non-cooperative stochastic target encirclement by anti-synchronization control via range-only measurement," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5480–5485.
- [8] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han, "Searching efficient 3d architectures with sparse point-voxel convolution," in *European Conference on Computer Vision*. Springer, 2020, pp. 685–702.
- [9] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin, "Cylindrical and asymmetrical 3d convolution networks for lidar segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9939–9948.
- [10] X. Yan, J. Gao, C. Zheng, C. Zheng, R. Zhang, S. Cui, and Z. Li, "2dpass: 2d priors assisted semantic segmentation on lidar point clouds," in *European Conference on Computer Vision*. Springer, 2022, pp. 677–695.
- [11] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9297–9307.
- [12] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.
- [13] P. Jiang and S. Saripalli, "Lidarnet: A boundary-aware domain adaptation model for point cloud semantic segmentation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2457–2464.
- [14] M. Rochan, S. Aich, E. R. Corral-Soto, A. Nabatchian, and B. Liu, "Unsupervised domain adaptation in lidar semantic segmentation with self-supervision and gated adapters," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2649–2655.
- [15] D. Tsai, J. S. Berrio, M. Shan, E. Nebot, and S. Worrall, "Viewer-centred surface completion for unsupervised domain adaptation in 3d object detection," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9346–9353.
- [16] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 621–11 631.
- [17] P. Sun, H. Kretzschmar, X. Dotiwala, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2446–2454.
- [18] M. Jaritz, T.-H. Vu, R. d. Charette, E. Wirbel, and P. Perez, "xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2020.
- [19] D. Peng, Y. Lei, W. Li, P. Zhang, and Y. Guo, "Sparse-to-dense feature matching: Intra and inter domain cross-modal learning in domain adaptation for 3d semantic segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7108–7117.
- [20] B. Xing, X. Ying, R. Wang, J. Yang, and T. Chen, "Cross-modal contrastive learning for domain adaptation in 3d semantic segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 3, 2023, pp. 2974–2982.
- [21] W. Liu, Z. Luo, Y. Cai, Y. Yu, Y. Ke, J. M. Junior, W. N. Gonçalves, and J. Li, "Adversarial unsupervised domain adaptation for 3d semantic segmentation with multi-modal learning," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 176, pp. 211–221, 2021.
- [22] M. Li, Y. Zhang, Y. Xie, Z. Gao, C. Li, Z. Zhang, and Y. Qu, "Cross-domain and cross-modal knowledge distillation in domain adaptation for 3d semantic segmentation," in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 3829–3837.
- [23] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *International Conference on Machine Learning*. PMLR, 2015, pp. 1180–1189.
- [24] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7167–7176.
- [25] H. Huang, Q. Huang, and P. Krahenbuhl, "Domain transfer through deep activation matching," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 590–605.
- [26] Y. Li, L. Yuan, and N. Vasconcelos, "Bidirectional learning for domain adaptation of semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6936–6945.
- [27] Y. Zou, Z. Yu, B. Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 289–305.
- [28] Y. Zou, Z. Yu, X. Liu, B. Kumar, and J. Wang, "Confidence regularized self-training," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 5982–5991.
- [29] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2517–2526.
- [30] M. Chen, H. Xue, and D. Cai, "Domain adaptation for semantic segmentation with maximum squares loss," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2090–2099.
- [31] B. Wu, X. Zhou, S. Zhao, X. Yue, and K. Keutzer, "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)*. IEEE, 2019, pp. 4376–4382.
- [32] S. Zhao, Y. Wang, B. Li, B. Wu, Y. Gao, P. Xu, T. Darrell, and K. Keutzer, "epointda: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 4, 2021, pp. 3500–3509.
- [33] A. Elhadidy, M. Afifi, M. Hassoubah, Y. Ali, and M. ElHelw, "Improved semantic segmentation of low-resolution 3d point clouds using supervised domain adaptation," in *2020 2nd Novel Intelligent and Leading Emerging Sciences Conference (NILES)*. IEEE, 2020, pp. 588–593.
- [34] L. Yi, B. Gong, and T. 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*, 2021, pp. 15 363–15 373.
- [35] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn, *et al.*, "A2d2: Audi autonomous driving dataset," *arXiv preprint arXiv:2004.06320*, 2020.
- [36] H. Cao, Y. Xu, J. Yang, P. Yin, S. Yuan, and L. Xie, "Multi-modal continual test-time adaptation for 3d semantic segmentation," *arXiv preprint arXiv:2303.10457*, 2023.
- [37] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.

- [38] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "Mixmatch: A holistic approach to semi-supervised learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [39] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6023–6032.
- [40] Y. Chen, V. T. Hu, E. Gavves, T. Mensink, P. Mettes, P. Yang, and C. G. Snoek, "Pointmixup: Augmentation for point clouds," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 330–345.
- [41] J. Zhang, L. Chen, B. Ouyang, B. Liu, J. Zhu, Y. Chen, Y. Meng, and D. Wu, "Pointcutmix: Regularization strategy for point cloud classification," *Neurocomputing*, vol. 505, pp. 58–67, 2022.
- [42] C. Saltori, F. Galasso, G. Fiameni, N. Sebe, E. Ricci, and F. Poiesi, "Cosmix: Compositional semantic mix for domain adaptation in 3d lidar segmentation," in *European Conference on Computer Vision*. Springer, 2022, pp. 586–602.
- [43] F. Hasecke, P. Colling, and A. Kummert, "Fake it, mix it, segment it: Bridging the domain gap between lidar sensors," *arXiv preprint arXiv:2212.09517*, 2022.
- [44] L. Kong, N. Quader, and V. E. Liong, "Conda: Unsupervised domain adaptation for lidar segmentation via regularized domain concatenation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9338–9345.
- [45] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, ser. KDD'96. AAAI Press, 1996, p. 226–231.
- [46] S. Lee, H. Lim, and H. Myung, "Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 276–13 283.
- [47] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4213–4220.
- [48] P. Morerio, J. Cavazza, and V. Murino, "Minimal-entropy correlation alignment for unsupervised deep domain adaptation," *arXiv preprint arXiv:1711.10288*, 2017.
- [49] Y. Luo, L. Zheng, T. Guan, J. Yu, and Y. Yang, "Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 2507–2516.
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [51] B. Graham, M. Engelcke, and L. Van Der Maaten, "3d semantic segmentation with submanifold sparse convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9224–9232.
- [52] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in Neural Information Processing Systems*, vol. 32, 2019.