

# SUMMIT: Source-Free Adaptation of Uni-Modal Models to Multi-Modal Targets

Cody Simons<sup>1</sup> Dripta S. Raychaudhuri<sup>1,2,\*</sup> Sk Miraj Ahmed<sup>1</sup> Suya You<sup>3</sup>  
Konstantinos Karydis<sup>1</sup> Amit K. Roy-Chowdhury<sup>1</sup>

<sup>1</sup>University of California, Riverside <sup>2</sup>AWS AI Labs <sup>3</sup>DEVCOM Army Research Laboratory

{csimo005@, drayc001@, sahme047@, kkarydis@ece., amitrc@ece.}ucr.edu suya.you.civ@army.mil

## Abstract

*Scene understanding using multi-modal data is necessary in many applications, e.g., autonomous navigation. To achieve this in a variety of situations, existing models must be able to adapt to shifting data distributions without arduous data annotation. Current approaches assume that the source data is available during adaptation and that the source consists of paired multi-modal data. Both these assumptions may be problematic for many applications. Source data may not be available due to privacy, security, or economic concerns. Assuming the existence of paired multi-modal data for training also entails significant data collection costs and fails to take advantage of widely available freely distributed pre-trained uni-modal models. In this work, we relax both of these assumptions by addressing the problem of adapting a set of models trained independently on uni-modal data to a target domain consisting of unlabeled multi-modal data, without having access to the original source dataset. Our proposed approach solves this problem through a switching framework which automatically chooses between two complementary methods of cross-modal pseudo-label fusion – agreement filtering and entropy weighting – based on the estimated domain gap. We demonstrate our work on the semantic segmentation problem. Experiments across seven challenging adaptation scenarios verify the efficacy of our approach, achieving results comparable to, and in some cases outperforming, methods which assume access to source data. Our method achieves an improvement in mIoU of up to 12% over competing baselines. Our code is publicly available at <https://github.com/csimo005/SUMMIT>.*

## 1. Introduction

There has been a recent surge of interest in autonomous vehicles which typically rely on a wide variety of sensors. This has fueled the need for machine learning models ca-

pable of processing multiple sensing modalities, commonly referred to as multi-modal models. One problem of particular interest is 3D semantic segmentation, which has received a lot of interest [12, 24, 36] driven by the introduction of new multi-modal datasets [5, 11, 4]. Many 3D semantic segmentation methods (e.g., [12, 24, 36]) fuse data across different sensing modalities, e.g., RGB images and point clouds to obtain and employ colored pointclouds [31], to increase performance and robustness. As with most learning problems, the performance of 3D semantic segmentation degrades as the input data distribution diverges from the training set distribution [16], referred to as domain shift. This is particularly true for autonomous navigation which can experience domain shifts because of lighting and weather changes throughout the day, as well as geographic changes when traveling over large distances. Many works have sought to address this domain shift in the Unsupervised Domain Adaptation (UDA) setting [14, 21, 8]. Recently, the relationship between different modalities has been leveraged to aid in the adaptation process via cross-modal UDA (xMUDA) [16].

While xMUDA [16] has made significant improvements over uni-modal UDA methods, it assumes that the source dataset used for training (i) consists of paired multi-modal data, and (ii) is available during adaptation to the target domain.<sup>1</sup> However, these conditions may be hard to satisfy in real-world scenarios:

- The first assumption is problematic because it requires collection and annotation of large volumes of paired multi-modal data for every sensor configuration (e.g., RGB and depth, RGB and IR, etc), a very time-consuming operation. Also, it fails to take advantage of the large volumes of uni-modal data and pre-trained models that are easily accessible.
- The second assumption is problematic because sharing the source data for adaptation may be impossible due to privacy, security and commercial reasons. Additionally, as datasets have grown, their transfer and storage have begun to present non-trivial engineering challenges and financial costs.

<sup>1</sup>In this work, modality refers to a specific type of input data, such as RGB images or point clouds, while domain refers to the underlying data distribution, such as cities in different continents.

\* Currently at AWS AI Labs. Work done while the author was at UCR.

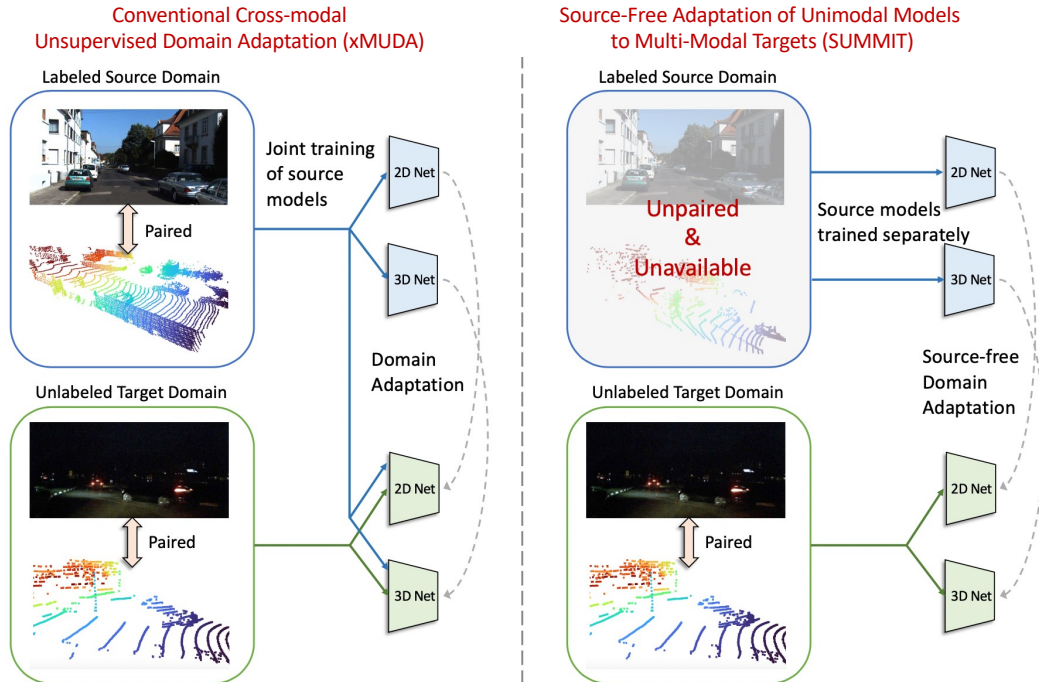


Figure 1. **Problem setup.** Our goal is to adapt a pair of uni-modal models, which have been trained independently on a source domain, to a target domain consisting of unlabeled, paired, multi-modal data, without access to the original source dataset. In contrast to conventional cross-modal UDA [16] (left panel), we do not assume that the source dataset used for training (i) consists of paired multi-modal data, and (ii) is available during adaptation to the target domain.

Relaxing the first assumption of paired modalities in the source domain remains unaddressed in cross-modal UDA. The second assumption has been relaxed in the source-free UDA setting [18, 2, 39, 37], but to the authors’ best knowledge, no such prior works deal with multi-modal data.

In this paper, we propose *Source-free adaptation of Uni-modal Models to Multi-modal Targets* (SUMMIT), relaxing both the above assumptions (see Figure 1). Relaxation of the two assumptions in conventional cross-modal UDA makes our problem substantially more challenging. In the conventional cross-modal UDA setting, the approach in [16] relies on labeled paired multi-modal data in the source domain to learn correlations across modalities. Learned correlations are then exploited to improve transfer to the target domain. In our setting, the correlations between modalities must be learned on the unlabeled target data, as we are working with uni-modal source models and do not assume access to such labeled pairs. The lack of source data already makes the alignment of the source and target distributions a challenging problem. Combining it with uni-modal models on the source side makes the overall problem even more challenging.

To address these challenges, we propose a new adaptation framework built on pseudo-label fusion across modalities. First, we utilize the trained uni-modal models to generate pseudo-labels on the target data, separately for each modality. Second, these pseudo-labels are fused together across

modalities to filter out noisy predictions. We introduce a data-driven switching method that automatically chooses between two complementary approaches for cross-modal pseudo-label fusion – *agreement filtering* and *entropy weighting* – based on the estimated domain gap. The fused pseudo-labels are used to supervise the process of learning the correlations across modalities allowing for cross-modal learning to take place. Optimizing for explicit cross-modal objectives, our framework learns the correlations across modalities even without the presence of the source data and improves transfer beyond standard uni-modal adaptation.

**Main contributions.** Our primary contributions can be summarized as follows.

- We address the problem of adapting a set of models trained independently on uni-modal data to a target domain consisting of unlabeled, paired, multi-modal data, without access to the original source dataset. This setting of great practical importance as explained above.
- We propose a new cross-modal, source-free UDA framework which fuses pseudo-labels across modalities using information-theoretic and hypothesis testing approaches. This helps increase robustness of the predictions and implicitly allows for cross-modal correlations to be learned on the target domain without access to source data.
- We perform extensive experiments on seven challenging benchmarks which demonstrate that our method provides an

improvement of up 12% over competing baselines.

## 2. Related Work

**Unsupervised Domain Adaptation.** UDA methods have been applied to a wide variety of computer vision tasks, including image classification [34], semantic segmentation [33] and object detection [15], in an effort to address the data distribution shift. Most approaches try to align the source and target data distributions, using techniques such as maximum mean discrepancy [20] and adversarial learning [10, 34, 25]. A separate line of work uses image translation techniques to perform adaptation by translating the source images to the target domain [14, 32]. In the case of semantic segmentation, existing UDA methods fall primarily into three categories: output space alignment, pixel-space alignment, and pseudo-labeling. The first category of methods aligns the output or feature distributions [33, 3]. The pixel alignment approaches use image translation techniques similar to image classification [14, 7, 6, 38], while the pseudo-labeling techniques generate pixel-wise pseudo-labels to fine-tune the source model [28, 40, 21, 8]. Some approaches also try to combine these strategies to perform adaptation [23, 22, 30]. Our method is a UDA approach but in a multi-modal and source-free setting.

**Source-free Domain Adaptation.** While methods mentioned in the previous section utilize the source data during adaptation, there has been a surge of interest in adaptation using only a pre-trained source model. Recent approaches in source-free adaptation have primarily focused on the image classification task. These include techniques such as information maximization [18, 2], pseudo-labeling [39] and self-supervision [37]. In the context of semantic segmentation, [19] proposed an algorithm that combines ideas from the above methods by implementing a curriculum learning scheme. Similar ideas are explored in [35], where the authors additionally explore the concept of negative learning [17]. Finally, in [27], pseudo-labeling and uncertainty estimation via dropout are used to tackle the problem. In contrast to all these approaches which approach source-free adaptation in a *uni-modal* setting, our framework performs source-free adaptation in a *multi-modal* setting by exploiting the intrinsic semantic relationships between the different modalities.

**Cross-modal Domain Adaptation.** Despite the keen interest in multi-modal analysis, there have been few works that attempt this in the domain adaptive scenario. Recently, a cross-modal UDA (xMUDA) setting was proposed [16], where modalities learn from each other through mutual cycle-consistency to prevent the stronger modality from adopting false predictions from the weaker one. A more recent work [1], motivated by the fact that the availability of multi-modal data and models is usually limited, considered adapting source model(s) trained on some source modality to a differ-

ent target modality without access to the source data. Our problem setting also considers that uni-modal models are more widely available than multi-modal ones, but addresses a different problem: adapting uni-modal models of different modalities to a multi-modal target without source data. Critically different from [16], we do not assume the existence of paired multi-modal training data and do not need access to the source data during adaptation.

## 3. Method

### 3.1. Problem Setting

We address the problem of adapting a set of models trained independently on uni-modal data to a target domain consisting of unlabeled, multi-modal data, without access to the original source dataset. We consider a set of source models, each of a unique modality, trained in a supervised manner on a set of  $K$  categories. Once these models have been trained, the source data is discarded. We are then given a new target domain containing unlabeled, paired, multi-modal data, corresponding to each of the modalities present in the source datasets. We aim to adapt the source models to this new domain by exploiting the semantic relationships that inherently exist between the modalities.

Throughout this work we will consider two input modalities: 2D image data  $\mathcal{X}^{2D}$ , of dimension  $H \times W \times 3$ , and 3D point cloud data  $\mathcal{X}^{3D}$ , of dimension  $N \times 3$  ( $N$  is the number of points), each of which has a corresponding label space  $\mathcal{Y}^{2D}$  and  $\mathcal{Y}^{3D}$ . The source domain  $\mathcal{D}_S = \{\{x_i^{2D}, y_i^{2D}\}_{i=1}^{N_{S,2D}}, \{x_i^{3D}, y_i^{3D}\}_{i=1}^{N_{S,3D}}\}$  consists of samples  $x^{2D} \in \mathcal{X}^{2D}$  and  $x^{3D} \in \mathcal{X}^{3D}$  and their corresponding labels  $y^{2D} \in \mathcal{Y}^{2D}$  and  $y^{3D} \in \mathcal{Y}^{3D}$ , with no known association between  $x^{2D}$  and  $x^{3D}$ . We assume that each source model,  $\mathcal{M}^{2D} : \mathcal{X}^{2D} \rightarrow \mathcal{Y}^{2D}$  and  $\mathcal{M}^{3D} : \mathcal{X}^{3D} \rightarrow \mathcal{Y}^{3D}$ , is trained with access to only a single input modality. The models can be decomposed as  $\mathcal{M}^{2D} = f^{2D} \circ g^{2D}$  and  $\mathcal{M}^{3D} = f^{3D} \circ g^{3D}$ , where  $f^{2D}, f^{3D}$  are the respective feature encoders and  $g^{2D}, g^{3D}$  are the corresponding linear classifiers. We adapt these source models to the target domain  $\mathcal{D}_T = \{x_i^{2D}, x_i^{3D}\}_{i=1}^{N_T}$  consisting of paired samples  $x_i^{2D} \in \mathcal{X}^{2D}$  and  $x_i^{3D} \in \mathcal{X}^{3D}$ , but no label information.

### 3.2. Framework Overview

We adapt the source models via a pseudo-labeling approach, where predictions from the source models are fused across modalities to supervise the training on the unlabeled target dataset. Fusing pseudo-labels allows us to create more accurate labels for adaptation, and implicitly allows the usage of cross-modal correlations as part of the process.

We propose a switching framework which automatically determines how to change between two complementary fusion methods: (i) **agreement filtering**, and (ii) **entropy weighting**. Agreement filtering utilizes the consensus be-

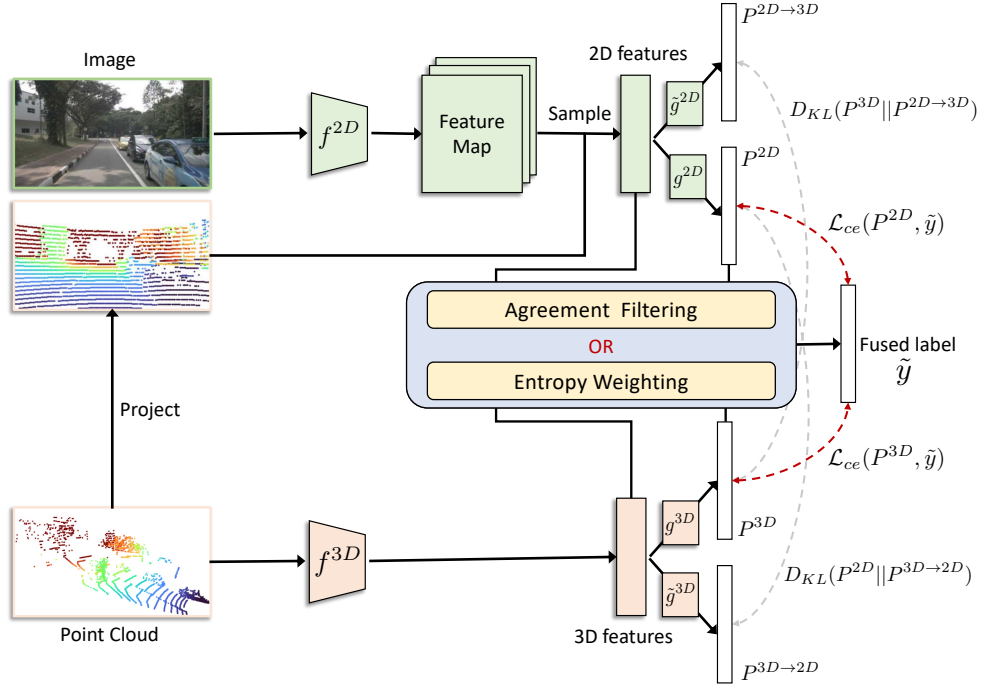


Figure 2. **Overview of framework.** Our framework consists of two streams corresponding to the 2D and 3D inputs respectively. The two modalities are processed by separate feature encoders  $f^{2D}$  and  $f^{3D}$ , with the 2D features being sampled by projecting the 3D points onto the corresponding RGB image. The four segmentation outputs consist of the main predictions  $P^{2D}$ ,  $P^{3D}$  and the modality translated predictions  $P^{2D \rightarrow 3D}$ ,  $P^{3D \rightarrow 2D}$  (using auxiliary heads  $\tilde{g}^{2D}$ ,  $\tilde{g}^{3D}$ ). The main predictions are used to compute pseudo-labels via a fusion strategy, agreement filtering or entropy weighting. The choice of approach is decided via a switching strategy. These pseudo-labels are used to train the framework using cross-entropy loss. On the other hand, we transfer knowledge across modalities using  $D_{KL}(P^{3D} || P^{2D \rightarrow 3D})$ , where the objective of the 2D translation head is to estimate the main 3D output and vice versa,  $D_{KL}(P^{2D} || P^{3D \rightarrow 2D})$ .

tween predictions from the different modalities as a score to combine and refine the pseudo-labels. On the other hand, entropy weighting employs an information-theoretic approach to fuse predictions across modalities, supplemented by a hypothesis testing over class-conditional feature space statistics to refine predictions. The optimal fusion approach is chosen *automatically* during adaptation via a switching strategy which utilizes meta-data derived from the source models.

Even with refined pseudo-labels, adaptation is still challenging due to the fact that semantic categories in the real-world are usually quite imbalanced. Thus, we present a method for estimating the class distribution, which in turn can facilitate the learning process. Using the estimated class distributions, the refined pseudo-labels are subsequently used to jointly adapt the source models to the target domain. An overview of our framework is shown in Figure 2.

### 3.3. Pseudo-label Fusion

The pseudo-labels obtained from the uni-modal source models are potentially noisy due to the domain shift between the source and the target domains. Using these directly to conduct the adaptation process can degrade the performance of the adapted models. Thus, we propose a switching

framework that automatically chooses between two complementary methods to fuse pseudo-labels across modalities in order to obtain a single pseudo-label to supervise adaptation.

**Agreement Filtering.** In agreement filtering (AF), we implement fusion by comparing the pseudo-labels across modalities and keeping only those that agree with each other. First, we calculate pseudo-labels for each modality, by taking the most likely prediction according to the unadapted source model. Next, we perform an initial refinement by filtering out any pseudo-label with a confidence score that is lower than the median confidence value for that class, *i.e.*

$$\tilde{y}_i^{2D} = \begin{cases} \arg \max_k \mathcal{M}_k^{2D}(x_i^{2D}), & \mathcal{M}_k^{2D}(x_i^{2D}) \geq m_k \\ \text{ignore}, & \text{otherwise} \end{cases} \quad (1)$$

$$m_k = \text{median}(\{\mathcal{M}_k^{2D}(x^{2D}) | x^{2D} \in \mathcal{D}_{\mathcal{T}}\}) \quad (2)$$

where  $k$  denotes the class index. We use the median here since it tends to be a better indicator of the central tendency given skewed data. Once both 2D and 3D pseudo-labels are filtered, we combine them via an agreement filter as follows,

$$\tilde{y}_i = \begin{cases} \tilde{y}_i^{2D}, & \tilde{y}_i^{2D} = \tilde{y}_i^{3D} \\ \text{ignore}, & \tilde{y}_i^{2D} \neq \tilde{y}_i^{3D} \end{cases} \quad (3)$$

Consequently, the new set of refined pseudo-labels contains only highly confident labels that agree across modalities. Certain types of bias will be specific to each modality. For example, a model trained on image data might associate green pixels with vegetation, but a model trained on point clouds does not sense color and cannot share this bias. Since we require every source model agree and each source model is trained on a different modality, we are left with a single common modality agnostic pseudo-label.

**Entropy Weighting.** Entropy weighting (EW) explores an alternative information-theoretic method for combining pseudo-labels. Inspired by [29], we perform an initial fusion by linearly combining the predictions from the source models by using uncertainty weights derived from the entropy of the output probabilities. Specifically,

$$p = w^{2D}\psi(\mathcal{M}^{2D}(x^{2D})) + w^{3D}\psi(\mathcal{M}^{3D}(x^{3D})) \quad (4)$$

$$w^{2D} = \frac{e^{-h(\mathcal{M}^{2D}(x^{2D}))}}{e^{-h(\mathcal{M}^{2D}(x^{2D}))} + e^{-h(\mathcal{M}^{3D}(x^{3D}))}} \quad (5)$$

$$w^{3D} = 1 - w^{2D} \quad (6)$$

where  $h(x) = -\sum_{k=1}^K \psi_k(x) \log(\psi_k(x))$  is the entropy over the output space,  $\psi$  is the softmax function, and  $p$  denotes the merged output probabilities. The merged probabilities are filtered using the class median of the merged probability, as defined previously for individual modalities in Eq. (1), giving a single fused pseudo-label.

**Entropy Weighting Refinement.** While this weighted combination is useful for resolving minor disagreements between modalities, it may also transfer noise across modalities leading to accurate pseudo-labels being rejected. In order to recover these pseudo-labels, we integrate target dataset statistics via hypothesis testing. The statistics we make use of are the class-wise mean and standard deviation calculated in the 2D and 3D feature spaces. We use the pseudo-labels previously accepted by Entropy Weighting as the class labels to calculate these statistics. The mean feature for class  $k$  in the 2D feature space is notated as  $\mu_k^{2D}$  and the corresponding standard deviation as  $\sigma_k^{2D}$  and similarly we define  $\mu_k^{3D}$  and  $\sigma_k^{3D}$  for the 3D feature space.

Given a sample consisting of paired observations  $x^{2D}$  and  $x^{3D}$  and the unadapted source models  $\mathcal{M}^{2D}$  and  $\mathcal{M}^{3D}$  we calculate a hypothesis pseudo-label using each modality individually as

$$k_{2D} = \arg \max_k \mathcal{M}^{2D}(x^{2D}) \quad (7)$$

$$k_{3D} = \arg \max_k \mathcal{M}^{3D}(x^{3D}). \quad (8)$$

Within the 2D feature space, we take  $k_{2D}$  as the null hypothesis and  $k_{3D}$  as the alternative hypothesis. We assume the likelihood function of a hypothesis to be a multivariate normal distribution parameterized by the class-wise mean and variance. We can then perform a likelihood ratio test:

$$\frac{\mathcal{N}(f^{2D}(x^{2D}); \mu_{k_{2D}}^{2D}, (\sigma_{k_{2D}}^{2D})^2)}{\mathcal{N}(f^{2D}(x^{2D}); \mu_{k_{3D}}^{2D}, (\sigma_{k_{3D}}^{2D})^2)} \leq \tau. \quad (9)$$

We set  $\tau = 1$  throughout our experiments and perform a sensitivity analysis of the threshold in the supplemental materials. We perform the same test in the 3D feature space, however in this case we take  $k_{3D}$  as the null hypothesis and  $k_{2D}$  as the alternative.

If both hypothesis tests agree on the correct hypothesis, e.g. the 2D feature space test rejects  $k_{2D}$  in favor of  $k_{3D}$  and the 3D feature space test keeps the  $k_{3D}$ , we take this as the correct pseudo-label. If the hypothesis tests do not agree on which hypothesis is correct, the pseudo-label is rejected. This allows us to recover previously rejected pseudo-labels and the entire set is used to supervise the adaptation process, as shown in Section 3.4.

**Automatic Switching Between Fusion Methods.** The performance of the two fusion techniques previously described varies depending on the particulars of the adaptation task. When the domain gap is small we expect EW to work better since it can resolve small disagreements across modalities, otherwise, AF is chosen since fusing discrete class labels has less potential to transfer noise. We cannot compute the domain gap directly, so we propose to measure it indirectly via the agreement between modalities. We justify the use of agreement as a proxy measure for domain gap via an empirical examination of agreement rates across various adaptation scenarios in Section 4.2.

First, we estimate the expected rate of agreement between the unadapted source models. Models agree with each other either if they both predict the same correct class, or both predict the same incorrect class. Considering only the case where the models agree and are correct, we can lower bound the expected agreement rate using source meta-data, more specifically, the Top-1 accuracy of the model on the source data.<sup>2</sup> Considering the predictions from each model to be independent of each other, the expected agreement can be lower bounded by the product of the Top-1 accuracies, *i.e.*

$$P(\hat{y}^{2D} = \hat{y}^{3D}) \geq \mathcal{P}(\hat{y}^{2D} = y|x^{2D})\mathcal{P}(\hat{y}^{3D} = y|x^{3D}). \quad (10)$$

This is the likelihood that source models agree given the source data distribution, which is our null hypothesis.

<sup>2</sup>The Top-1 accuracy of a model is a statistic commonly reported for most models.

The actual agreement rate between the source models on the target dataset is calculated using AF, as described in Section 3.3, and is the proportion of accepted pseudo-labels. This gives us the likelihood that the models agree given the target dataset distribution, which we take as the alternative hypothesis.

We compare these two likelihoods using hypothesis ratio testing. If the ratio between them is close to unity, then neither data distribution is more likely based on the agreement rate. Thus, we can infer that the target data comes from a data distribution similar to the source, and EW is chosen as the fusion method. The more the ratio diverges from unity the less likely the source and target datasets have a similar distribution. Throughout our experiments we set threshold of 0.5 to automatically switch to AF. We examine the choice of threshold in Section 4.2

### 3.4. Training

Using pseudo-labels, refined using AF or EW as chosen by our switching method, we optimize the following loss to adapt the uni-modal models to the target domain,

$$\mathcal{L}_{pl} = \mathcal{L}_{ce}(P^{2D}, \tilde{y}) + \mathcal{L}_{ce}(P^{3D}, \tilde{y}). \quad (11)$$

Here,  $\mathcal{L}_{ce}$  denotes the weighted cross-entropy loss,  $P^{2D} = \psi(\mathcal{M}^{2D}(x^{2D}))$ , and  $P^{3D} = \psi(\mathcal{M}^{3D}(x^{3D}))$ . We use the weighted cross-entropy in order to address potential class imbalances during adaptation, using the estimates of the class distribution as the weights. Note that in the source-free setting we do not have access to labels either on the source or on the target domain. Thus, we estimate the class distribution by taking the mean of the output distribution of each modality across target instances as  $\tilde{p}^{2D} = \mathbb{E}_{\mathcal{D}_T} [P^{2D}]$ . We estimate  $\tilde{p}^{3D}$  by averaging over the output of  $\mathcal{M}^{3D}$  in a similar fashion. The final estimate of the class distribution is calculated by averaging  $\tilde{p}^{2D}$  and  $\tilde{p}^{3D}$  as  $\tilde{p} = \frac{\tilde{p}^{2D} + \tilde{p}^{3D}}{2}$ . Inspired from [16], we optimize an additional cross-modal cycle-consistency constraint to allow information transfer between modalities through mutual imitation, *i.e.*

$$\mathcal{L}_{xM} = D_{KL}(P^{2D} || P^{3D \rightarrow 2D}) + D_{KL}(P^{3D} || P^{2D \rightarrow 3D}). \quad (12)$$

$P^{3D \rightarrow 2D}$  is an output head of  $\mathcal{M}^{3D}$  which attempts to predict  $P^{2D}$  and likewise  $P^{2D \rightarrow 3D}$  is a prediction of  $P^{3D}$  by  $\mathcal{M}^{2D}$  (see Figure 2). Minimizing the KL divergence  $D_{KL}$  allows for information flow between the two models. We combine these to get,  $\mathcal{L}_{tot} = \mathcal{L}_{pl} + \lambda \mathcal{L}_{xM}$ , where  $\lambda$  is a hyper-parameter to balance the two loss functions and is typically set to 0.1. More training details and pseudo-code can be found in the supplementary document.

## 4. Experiments

In this section, we demonstrate the ability of our proposed framework to perform source-free adaptation of uni-modal

models to multi-modal data. Experiments across multiple adaptation scenarios demonstrate that the pseudo-label refinement across modalities, in addition to the cross-modal consistency criterion, provides a consistent improvement in both 2D and 3D target domain performance over baselines. Even in scenarios where the 2D and 3D source data are derived from distinct domains, our framework is able to adapt the disconnected models to multi-modal datasets effectively. Visualizations of our segmentation results can be found in the supplementary materials.

### 4.1. Experimental Setup

**Datasets and Experimental Scenarios.** We use three autonomous vehicle datasets in our experiments, which are organized into three different source-free cross-modal UDA adaptation scenarios following the experimental setup defined in [16]. The critical difference is that our experiments have no access to the source data, and the source models are trained separately. Two of the adaptation experiments are derived from the *nuScenes* dataset [5]. The *nuScenes* dataset gathers data in multiple cities, throughout different times of day with multiple cameras and LIDAR. We extract approximately 34K images and the corresponding points clouds, and define two adaptation scenarios. The *first scenario* takes all image and point clouds that were gathered in the USA as the source domain and those gathered in Singapore as the target domain. The *second scenario* treats samples gathered during the day as the source domain and those gathered during the night as the target domain.

The *third scenario* considers two unrelated datasets as the source and target domain. The source domain is the *A2D2* dataset [11], which consists of approximately 40K images and point clouds. The target domain is the *SemanticKITTI* dataset [4], containing approximately 18K samples. The classes provided are dissimilar, thus, they are merged into larger super classes which are kept consistent between the domains; the details of this process can be found in the supplemental materials. This scenario is particularly challenging due to the different sensors used in each dataset, in addition to the temporal and geographic shifts.

While these experiments make no explicit use of pairing information when training the source models, there is the possibility of implicitly learning some cross-modal correlations since 2D and 3D samples are drawn from the same split of the dataset. Thus, we propose additional experiments, termed *crossover experiments*, with no corresponding data points shared between the 2D and 3D datasets. We take the source/target splits defined on the *nuScenes* dataset and use each to train a single modality, *i.e.* we use images from the USA split to train the 2D model and the point clouds from the Singapore split to train the 3D model. The *SemanticKITTI* dataset is used as the target dataset for all experiments. To increase the overlap of class labels between *nuScenes* and

Table 1. **Results on primary adaptation scenarios.** Our framework outperforms both the uni-modal and multi-modal source-free adaptation baselines by a large margin across all scenarios. Note that xMUDA needs source data, and the results below are provided assuming that the source data is available during adaptation. Despite xMUDA requiring access to source data and models that are jointly trained on multi-modal data, our method has comparable performance and even outperforms it in the A2D2/SemanticKITTI scenario.

| Method        | Multimodal | Source-free | USA/Singapore     |                   |                   | Day/Night         |                   |                   | A2D2/SemKITTI     |                   |                   |
|---------------|------------|-------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
|               |            |             | 2D                | 3D                | 2D+3D             | 2D                | 3D                | 2D+3D             | 2D                | 3D                | 2D+3D             |
| No adaptation | -          | -           | 49.67             | 47.39             | 55.96             | 41.85             | 42.32             | 49.98             | 32.85             | 34.01             | 40.52             |
| xMUDA         | ✓          | ✗           | 61.10             | 54.10             | 63.20             | 47.10             | 46.70             | 50.80             | 43.70             | 48.50             | 49.10             |
| xMUDA SF      | ✓          | ✓           | 51.88±0.12        | 51.04±0.27        | 56.59±1.46        | 36.43±1.05        | 42.23±1.62        | 44.46±1.49        | 33.74±0.15        | 36.79±0.23        | 42.68±0.19        |
| SHOT          | ✗          | ✓           | 52.21±0.00        | 48.20±0.01        | 58.18±0.00        | 34.70±0.03        | 40.61±0.00        | 40.83±0.01        | 33.26±0.23        | 36.92±0.06        | 42.09±0.18        |
| SUMMIT (Ours) | ✓          | ✓           | <b>57.92±1.43</b> | <b>52.95±0.72</b> | <b>61.10±1.97</b> | <b>44.73±0.01</b> | <b>44.53±0.00</b> | <b>50.73±0.01</b> | <b>44.68±0.30</b> | <b>48.56±0.34</b> | <b>49.93±0.13</b> |
| Oracle        | -          | -           | 64.67             | 57.19             | 71.29             | 46.11             | 40.31             | 45.60             | 58.84             | 71.75             | 74.57             |

SemanticKITTI, we include additional labels for background objects from the *nuScenes-lidarseg* dataset [9].

**Baselines.** We compare our method to (i) xMUDA [16], which is a cross-modal adaption method, (ii) a version of xMUDA trained with only the data available in our setting called xMUDA-SF, and (iii) SHOT [18] which is one of the seminal works in source-free UDA. Note that xMUDA requires source data, and the results are provided assuming that the source data is available during adaptation. In xMUDA-SF we remove any loss functions that work on the source data and begin the adaptation process with independently trained source models. The performance of xMUDA-SF highlights the reliance of xMUDA on paired source data. To apply SHOT, we treat each modality as its own uni-modal UDA problem. We use this comparison to show the potential of modeling and exploiting cross-correlations. Each of the baseline methods uses the unrefined uni-modal pseudo-labels. In addition, we compare with the unadapted source models as a lower bound on performance (**no adaptation**), and models adapted using oracle ground truth labels as an upper bound (**Oracle**). We keep the same baselines for the crossover experiments, but we do not compare them to the original xMUDA method since no source pairing information exists, which xMUDA needs. We do not compare against multi-source source-free adaptation methods such as [2], because they assume that each source comes from the same modality.

**Experimental Objectives.** The first three adaptation scenarios described above demonstrate that our method does not need access to paired source data, compared to state-of-the-art methods like [16], and that it can work with uni-modal trained models as the source. Furthermore, by comparing to existing source-free UDA methods, we will highlight the need to explicitly model cross-modal correlations.

Through the crossover experiments, we further show the reliance of [16] on *paired* data. Critically, we demonstrate our method works with truly independent source models, opening up the possibility of adapting pretrained uni-modal models together to work on multi-modal data.

**Implementation Details.** Architectures for the 2D and 3D

models follow from [16]. Specifically, the 2D network is a UNet architecture [26] with a ResNet34 [13] backbone. To associate points across modalities, the point cloud is projected onto the image plane, and the corresponding features are sampled. These are then passed through a final classification layer to give us predictions. The 3D model uses the official PyTorch implementation of the Sparse Convolutional Networks UNet [12]. The source models are initially trained on each modality independently of each other for 100K iterations. Once the source model is trained, we save this model and use it as the same starting point for all adaptation methods to ensure a fair comparison.

$\mathcal{L}_{xM}$  requires each model to provide a secondary output that predicts the output of the other modality. We initialize this second output head with a copy of the weights from the original output head. We set the  $\lambda$  for  $\mathcal{L}_{xM}$  according to the values used in [16], and the initial learning rate to  $1e - 6$ . Pseudo-labels are calculated offline before adaptation and are kept fixed throughout the adaptation process. Each method adapts the model for an additional 100K iterations. More details can be found in the supplementary document.

## 4.2. Results

We report the results on the three primary source-free adaptation scenarios in Table 1. We report the mIoU for each modality, 2D and 3D, along with the average across the softmax output of both modalities, as 2D+3D. Unsurprisingly, the highest performance is achieved by xMUDA, *however, this utilizes the source data during adaptation*. Its source-free counterpart, xMUDA-SF, obtains a much more modest improvement over the unadapted source, and in many cases leads to worse performance compared to no adaptation. This clearly highlights the reliance of xMUDA on the pairing information being present when training the source model. The improvement of SHOT is more consistent than xMUDA-SF, however since it is a uni-modal method, its improvement is limited when compared to our method.

Our framework achieves high-performance improvements over the source models across the board, with an improvement as high as 12% over the other source-free methods. We see this in the A2D2/SemanticKITTI experiment, with an

Table 2. **Results on crossover adaptation scenarios.** Differences in sensor configuration and label distribution result in all methods experiencing a drop in performance. While our method is not an exception to this, we obtain consistent improvements of up to 12% over the unadapted source model and other adaptation approaches.

| Method        | Multimodal | Source-free | (USA-2D,Singapore-3D)/SemKITTI |                   |                   | (Singapore-2D,USA-3D)/SemKITTI |                   |                   | (Day-2D,Night-3D)/ SemKITTI |                   |                   | (Night-2D,Day-3D)/SemKITTI |                   |                   |
|---------------|------------|-------------|--------------------------------|-------------------|-------------------|--------------------------------|-------------------|-------------------|-----------------------------|-------------------|-------------------|----------------------------|-------------------|-------------------|
|               |            |             | 2D                             | 3D                | 2D+3D             | 2D                             | 3D                | 2D+3D             | 2D                          | 3D                | 2D+3D             | 2D                         | 3D                | 2D+3D             |
| No adaptation | -          | -           | 18.98                          | 28.94             | 26.77             | 20.22                          | 28.45             | 25.51             | 24.21                       | 28.54             | 29.96             | 16.14                      | 26.60             | 20.52             |
| xMUDA SF      | ✓          | ✓           | 20.02±0.52                     | 32.68±1.01        | 28.10±0.65        | 20.05±0.28                     | 33.64±1.10        | 28.21±0.03        | 24.64±0.14                  | 34.01±0.10        | 29.54±0.17        | 16.13±0.03                 | 27.97±0.03        | 25.88±0.03        |
| SHOT          | ✗          | ✓           | 19.48±0.04                     | 34.07±0.06        | 27.39±0.03        | 20.40±0.04                     | 32.39±0.02        | 28.14±0.05        | 24.55±0.07                  | 34.08±0.03        | 29.13±0.07        | 16.11±0.02                 | 28.10±0.02        | 25.83±0.03        |
| SUMMIT (Ours) | ✓          | ✓           | <b>28.84±0.72</b>              | <b>33.42±0.73</b> | <b>31.53±0.63</b> | <b>28.74±0.70</b>              | <b>34.07±1.01</b> | <b>31.41±0.47</b> | <b>32.19±0.24</b>           | <b>38.71±0.27</b> | <b>35.51±0.16</b> | <b>28.51±0.16</b>          | <b>32.83±0.14</b> | <b>30.78±0.13</b> |
| Oracle        | -          | -           | 55.36                          | 68.73             | 69.82             | 56.37                          | 67.69             | 70.60             | 58.20                       | 67.82             | 70.82             | 56.41                      | 68.69             | 69.90             |

Table 3. **Estimated source and target agreement.** The ratio between source and target agreement is used to estimate the domain gap. In the first two cases, which are partitions of the same data, the ratio is nearly one, so EW is selected. A lower ratio, as in the last case, indicates a high domain gap, so AF is selected.

|                  | USA/Singapore | Day/Night | A2D2/SemKITTI |
|------------------|---------------|-----------|---------------|
| Source Agreement | 94.25         | 95.33     | 84.75         |
| Target Agreement | 88.50         | 87.40     | 26.00         |
| Ratio            | 0.94          | 0.92      | 0.31          |

Table 4. **Pseudo-label accuracy.** EW consistently admits a higher proportion of correct pseudo-labels than AF. However, in the case of a large domain gap, such as A2D2/SemanticKITTI, EW accepts a large amount of incorrect labels, which AF filters out.

| Method | USA/Singapore |           |        | Day/Night |           |        | A2D2/SemKITTI |           |        |
|--------|---------------|-----------|--------|-----------|-----------|--------|---------------|-----------|--------|
|        | Correct       | Incorrect | Ignore | Correct   | Incorrect | Ignore | Correct       | Incorrect | Ignore |
| AF     | 87.29         | 0.13      | 12.58  | 88.45     | 0.10      | 11.46  | 23.76         | 1.73      | 74.51  |
| EW     | 93.40         | 0.65      | 5.94   | 94.51     | 0.56      | 4.93   | 55.42         | 24.72     | 19.85  |

improvement of 12% on individual modalities compared to xMUDA-SF and SHOT which provides only a few points of improvement. On the 2D+3D we have a smaller, but still significant, improvement of 7% over SHOT. We also note that in several cases we are within 1% of the original xMUDA score, which assumes access to the source dataset.

In Table 2, we report the performance on the *crossover experiments*. Note that SHOT outperforms xMUDA-SF in several cases, indicating that in the crossover experiments, xMUDA-SF tends to transfer more noise between the models. This may be due to lower performing source models since this adaptation scenario goes between two different datasets as opposed to the previous experiments where two of the experiments transfer between partitions of a single dataset. Our framework outperforms existing baselines, with improvements ranging from 4%-10% over the unadapted source. The largest improvement is on Night 2D & Day 3D, despite the poor performance of the initial 2D source model. In some cases the performance of SHOT on the 3D modality approaches our performance; however, our method achieves a stronger improvement on the 2D and 2D+3D modality, showing that our method transfers useful information from the 3D to the 2D modality without transferring noise.

**Analysis of Automatic Switching.** To gain further insight into the performance of our switching method, we analyze

Table 5. We examine the interplay between AF, EW,  $\mathcal{L}_{xM}$ . We note that both AF and EW on their own provide an improvement across modalities and work well with the cross-modal loss. The cross-modal loss provides the most improvement in the 2D+3D performance.

| Method                 | USA/Singapore |       |       | Day/Night |       |       | A2D2/SemKITTI |       |       |
|------------------------|---------------|-------|-------|-----------|-------|-------|---------------|-------|-------|
|                        | 2D            | 3D    | 2D+3D | 2D        | 3D    | 2D+3D | 2D            | 3D    | 2D+3D |
| No Adaptation          | 49.67         | 47.39 | 55.96 | 41.85     | 42.32 | 49.98 | 32.85         | 34.01 | 40.52 |
| AF                     | 55.73         | 53.25 | 58.48 | 38.68     | 41.38 | 39.41 | 44.92         | 48.54 | 50.05 |
| AF+ $\mathcal{L}_{xM}$ | 56.78         | 53.29 | 58.83 | 37.78     | 41.36 | 39.10 | 45.01         | 48.34 | 49.60 |
| EW                     | 57.32         | 51.97 | 62.28 | 44.23     | 44.61 | 50.72 | 35.93         | 43.16 | 41.23 |
| EW+ $\mathcal{L}_{xM}$ | 57.47         | 52.12 | 62.32 | 44.73     | 44.53 | 50.72 | 36.82         | 43.26 | 41.24 |
| Oracle                 | 64.67         | 57.19 | 71.29 | 46.11     | 40.31 | 45.60 | 58.84         | 71.75 | 74.57 |

Table 6. We examine the interplay between agreement filtering (AF), entropy weighting (EW), and the cross-modal loss  $\mathcal{L}_{xM}$  for the *crossover* experiments. We note that in the *crossover* experiments AF provides substantial improvements, while EW does not. Also worth noting is that most of the performance gain comes from just using AF, with  $\mathcal{L}_{xM}$  providing modest improvements.

| Method                 | USA-2D,Singapore-3D |       |       | Singapore-2D,USA-3D |       |       | Day-2D,Night-3D |       |       | Night-2D,Day-3D |       |       |
|------------------------|---------------------|-------|-------|---------------------|-------|-------|-----------------|-------|-------|-----------------|-------|-------|
|                        | 2D                  | 3D    | 2D+3D | 2D                  | 3D    | 2D+3D | 2D              | 3D    | 2D+3D | 2D              | 3D    | 2D+3D |
| No adaptation          | 18.98               | 28.94 | 26.77 | 20.22               | 28.45 | 25.51 | 24.21           | 28.54 | 29.96 | 16.14           | 26.60 | 20.52 |
| AF                     | 29.34               | 33.00 | 31.86 | 28.50               | 34.59 | 31.32 | 31.76           | 38.19 | 35.07 | 28.22           | 32.61 | 30.55 |
| AF+ $\mathcal{L}_{xM}$ | 28.02               | 34.27 | 30.81 | 29.55               | 32.91 | 31.95 | 32.42           | 38.54 | 35.67 | 28.63           | 32.98 | 30.87 |
| EW                     | 16.39               | 20.38 | 18.41 | 16.20               | 17.87 | 17.22 | 19.61           | 18.93 | 20.60 | 7.27            | 7.84  | 8.23  |
| EW+ $\mathcal{L}_{xM}$ | 15.64               | 18.57 | 16.88 | 16.26               | 19.49 | 18.08 | 19.64           | 19.33 | 20.68 | 7.71            | 7.97  | 8.60  |
| Oracle                 | 55.36               | 68.73 | 69.82 | 56.37               | 67.69 | 70.60 | 58.20           | 67.82 | 70.82 | 56.41           | 68.69 | 69.90 |

when particular pseudo-label fusion approaches are chosen, and the resulting quality of pseudo-labels that are being selected. In Table 3 we show the source and target agreement calculated according to the approach described in Section 3.3, and the ratio between them on the first three experiments. For the USA/Singapore and Day/Night experiments, which are both derived by partitioning a single dataset, we obtain a ratio close to one and EW is selected. For the A2D2/SemanticKITTI experiment, which adapts across datasets, the ratio drops to 0.31 which is less than our threshold of 0.5 leading to AF being selected.

The efficacy of the switching is shown by the pseudo-label accuracy in Table 4. In the first two cases, EW gives correct pseudo-labels for over 90% of samples, while allowing less than 1% of incorrect labels. In the third experiment AF is selected and while far fewer correctly labeled samples are accepted the amount of incorrect labels is also much lower (< 2%), i.e., AF selects fewer, but correct, samples. Throughout the *crossover experiments*, AF is selected, a full analysis of which is included in the supplemental materials.






Figure 3. Ground truth data for comparison with qualitative results. Each color represents a different segmentation class.

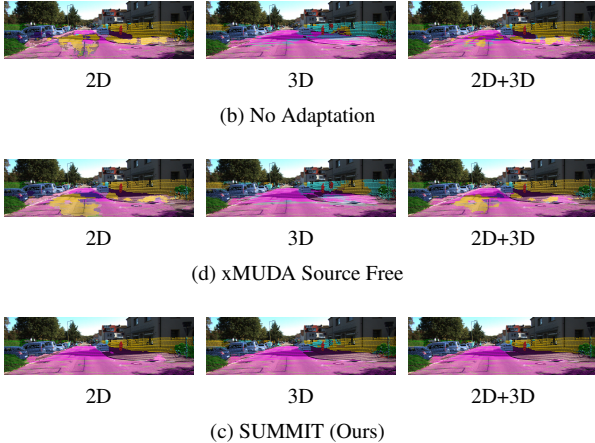


Figure 4. Qualitative results on the A2D2/SemanticKITTI experiment, comparing our method to the unadapted source model and xMUDA Source Free. Note the patches of misclassified road, which xMUDA SF exacerbates, but our method fixes. Images are best viewed digitally.

**Ablation Study.** We perform an ablation study on the adaptation and the *crossover experiments* where we examine the performance of each of our complementary pseudo-label fusion methods both with and without the cross-modal loss. In Table 5 we can clearly see the majority of improvements coming from the pseudo label fusion, with some additional improvements from the cross-modal loss. We note in the USA/Singapore and Day/Night experiments entropy weighting outperforms the agreement filtering method, due to the smaller domain gap. However, in this case, agreement filtering still gives strong performance. In the A2D2/SemanticKITTI agreement filtering performs better, since there is a larger domain gap. We also note that our switching method consistently picks the best method for pseudo-label fusion.

In Table 6 we show the results on the *crossover experiments*. We see that AF far outperforms EW in this setting. Since the *crossover experiments* are all across datasets, the domain gap is larger and AF is consistently selected. Most of the performance increase still comes solely from our pseudo-label fusion. We still see minor improvements from the cross-modal loss, but in some cases, this comes at the expense of one of the modalities.

**Qualitative Results.** We visualize the output of our adaptation methods as well as no adaptation and xMUDA SF baselines in Fig. 4, for the adaptation experiment, and in Fig.

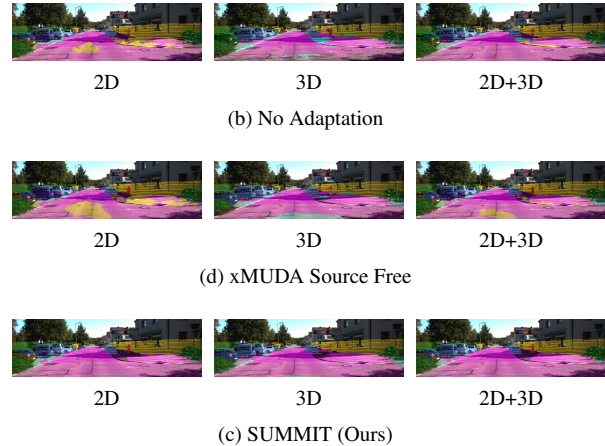


Figure 5. Qualitative results on the Day/Night to SemanticKITTI *crossover* experiments. Once again we see xMUDA SF exacerbating misclassifications. Images are best viewed digitally.

5, for the *crossover* experiment. They can both be compared to the ground truth segmentation in Fig. 3. In both figures, we see that in the unadapted source, large portions of the road and sidewalk are misclassified in the 2D output. We see similar errors in the 3D and 2D+3D output. When the source-free xMUDA adaptation is applied, we see that the results can get worse, with even larger portions of the road and the sidewalk being misclassified. When we apply our method we see that most of the misclassifications are fixed, critically the road and sidewalk are now correctly classified. Some small portions of the building are misclassified, but they are relatively small regions. The combined output does fix some of the smaller misclassifications and the combination does not introduce any new errors. The images are best viewed digitally, where they can be zoomed in. A short video is provided in the supplementary materials.

## 5. Conclusion

In this paper we introduced a new multi-modal UDA setting, Source-free Adaptation of Uni-modal Models to Multi-Modal Targets (SUMMIT), where source models are trained independently on each modality and source data are not needed during adaptation. We addressed this new setting by fusing information across modalities to improve pseudo-labeling. We proposed a data driven switching method that chooses between two complementary methods for fusing pseudo-labels across modalities, which provides an improvement of up to 12% over competing baselines.

**Acknowledgments** This work was supported in part by the US Department of Defense Laboratory University Collaboration Initiative program, National Science Foundation (Grant No. 172434 and 1901379), National Institute for Food and Agriculture (Award No. 2021-67022-33453) and the UC Multi-campus Research Programs Initiative.

## References

- [1] Sk Miraj Ahmed, Suhas Lohit, Kuan-Chuan Peng, Michael J Jones, and Amit K Roy-Chowdhury. Cross-modal knowledge transfer without task-relevant source data. In *ECCV*, 2022. 3
- [2] Sk Miraj Ahmed, Dripta S Raychaudhuri, Sujoy Paul, Samet Oymak, and Amit K Roy-Chowdhury. Unsupervised multi-source domain adaptation without access to source data. In *CVPR*, 2021. 2, 3, 7
- [3] Nikita Araslanov and Stefan Roth. Self-supervised augmentation consistency for adapting semantic segmentation. In *CVPR*, 2021. 3
- [4] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *ICCV*, 2019. 1, 6
- [5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020. 1, 6
- [6] Wei-Lun Chang, Hui-Po Wang, Wen-Hsiao Peng, and Wei-Chen Chiu. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *CVPR*, 2019. 3
- [7] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *ICCV*, 2019. 3
- [8] Jiahua Dong, Yang Cong, Gan Sun, Yuyang Liu, and Xiaowei Xu. Cslc: Critical semantic-consistent learning for unsupervised domain adaptation. In *ECCV*, 2020. 1, 3
- [9] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *arXiv preprint arXiv:2109.03805*, 2021. 7
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. In *JMLR*, 2016. 3
- [11] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, et al. A2d2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320*, 2020. 1, 6
- [12] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 2018. 1, 7
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7
- [14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 1, 3
- [15] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *WACV*, 2020. 3
- [16] Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Emilie Wirbel, and Patrick Pérez. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *CVPR*, 2020. 1, 2, 3, 6, 7
- [17] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *CVPR*, 2019. 3
- [18] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020. 2, 3, 7
- [19] Yuang Liu, Wei Zhang, and Jun Wang. Source-free domain adaptation for semantic segmentation. In *CVPR*, 2021. 3
- [20] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 3
- [21] Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *ECCV*, 2020. 1, 3
- [22] Luigi Musto and Andrea Zinelli. Semantically adaptive image-to-image translation for domain adaptation of semantic segmentation. In *BMVC*, 2020. 3
- [23] Sujoy Paul, Yi-Hsuan Tsai, Samuel Schuster, Amit K Roy-Chowdhury, and Manmohan Chandraker. Domain adaptive semantic segmentation using weak labels. In *ECCV*, 2020. 3
- [24] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 1
- [25] Dripta S Raychaudhuri, Sujoy Paul, Jeroen Vanbaaar, and Amit K Roy-Chowdhury. Cross-domain imitation from observations. In *International Conference on Machine Learning*, pages 8902–8912. PMLR, 2021. 3
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 7
- [27] Prabhu Teja S and Francois Fleuret. Uncertainty reduction for model adaptation in semantic segmentation. In *CVPR*, 2021. 3
- [28] Fatemeh Sadat Saleh, Mohammad Sadegh Aliakbarian, Mathieu Salzmann, Lars Petersson, and Jose M. Alvarez. Effective use of synthetic data for urban scene semantic segmentation. In *ECCV*, 2018. 3
- [29] Maohao Shen, Yuheng Bu, and Gregory Wornell. On the benefits of selectivity in pseudo-labeling for unsupervised multi-source-free domain adaptation. *arXiv preprint arXiv:2202.00796*, 2022. 5
- [30] M Naseer Subhani and Mohsen Ali. Learning from scale-invariant examples for domain adaptation in semantic segmentation. In *ECCV*, 2020. 3
- [31] Hanzhe Teng, Dimitrios Chatziparaschis, Xinyue Kan, Amit K Roy-Chowdhury, and Konstantinos Karydis. Centroid distance keypoint detector for colored point clouds. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1196–1205, 2023. 1
- [32] Luan Tran, Kihyuk Sohn, Xiang Yu, Xiaoming Liu, and Manmohan Chandraker. Gotta adapt 'em all: Joint pixel and feature-level domain adaptation for recognition in the wild. In *CVPR*, 2019. 3

- [33] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018. 3
- [34] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 3
- [35] Yuxi Wang, Jian Liang, Jun Xiao, Yuran Yang, Shuqi Mei, and Zhaoxiang Zhang. Domain adaptive semantic segmentation without source data: Align, teach and propagate. *arXiv preprint arXiv:2106.11653*, 2021. 3
- [36] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. In *ACM Transactions of Graphics*, 2019. 1
- [37] Haifeng Xia, Handong Zhao, and Zhengming Ding. Adaptive adversarial network for source-free domain adaptation. In *ICCV*, 2021. 2, 3
- [38] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *CVPR*, 2020. 3
- [39] Hao-Wei Yeh, Baoyao Yang, Pong C Yuen, and Tatsuya Harada. Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In *WACV*, 2021. 2, 3
- [40] Yang Zou, Zhiding Yu, B. V. K. Vijaya Kumar, and Jinsong Wang. Domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018. 3