

RESEARCH ARTICLE

Multimodal SAM-Adapter for Semantic Segmentation

IACOPO CURTI¹, PIERLUIGI ZAMA RAMIREZ¹, ALIOSCIA PETRELLI²,
AND LUIGI DI STEFANO¹, (Member, IEEE)

¹University of Bologna, 40136 Bologna, Italy

²SINA, 40132 Bologna, Italy

Corresponding author: Iacopo Curti (iacopo.curti2@unibo.it)

This work was supported in part by FSE+ 2021–2027 ai sensi dell’art. 24, comma 3, lett. a), della Legge 240/2010 e s.m.i. e del D.G.R. 693/2023 (RIF. PA: 2023-20090/RER-CUP: J19J23000730002), and in part by SINA.

ABSTRACT Semantic segmentation, a key task in computer vision with broad applications in autonomous driving, medical imaging, and robotics, has advanced substantially with deep learning. Nevertheless, current approaches remain vulnerable to challenging conditions such as poor lighting, occlusions, and adverse weather. To address these limitations, multimodal methods that integrate auxiliary sensor data (e.g., LiDAR, infrared) have recently emerged, providing complementary information that enhances robustness. In this work, we present MM SAM-adapter, a novel framework that extends the capabilities of the Segment Anything Model (SAM) for multimodal semantic segmentation. The proposed method employs an adapter network that injects fused multimodal features into SAM’s rich RGB features. This design enables the model to retain the strong generalization ability of RGB features while selectively incorporating auxiliary modalities only when they contribute additional cues. As a result, MM SAM-adapter achieves a balanced and efficient use of multimodal information. We evaluate our approach on three challenging benchmarks, DeLiVER, FMB, and MUSES, where MM SAM-adapter delivers *state-of-the-art* performance. To further analyze modality contributions, we partition DeLiVER and FMB into *RGB-easy* and *RGB-hard* subsets. Results consistently demonstrate that our framework outperforms competing methods in both favorable and adverse conditions, highlighting the effectiveness of multimodal adaptation for robust scene understanding.

INDEX TERMS Adapter, event cameras, thermal cameras depth, LiDAR, multimodal semantic segmentation, SAM.

I. INTRODUCTION

Semantic segmentation is a fundamental computer vision task that assigns a category label to each image pixel, with applications in fields such as autonomous driving, medical image analysis, and robotic navigation. The advent of deep learning has revolutionized semantic segmentation, leading to remarkable improvements in accuracy, efficiency, and generalization across environments [1]. Despite advances, semantic segmentation from RGB images fails in challenging situations such as poorly lit scenes, adverse weather conditions, and motion blur. To overcome these

limitations, additional sensors that provide complementary measurements, such as infrared or event cameras and LiDAR, may be deployed to enhance segmentation performance. The increasing adoption of auxiliary sensors has fostered the creation of numerous datasets and approaches [2] that use multiple input modalities to pursue semantic segmentation. Standard practice for state-of-the-art multimodal methods is to train networks on the target dataset starting from ImageNet [3] weights, i.e., building upon the foundational knowledge acquired by image classification.

However, the recent availability of large-scale annotated datasets and large GPU clusters has enabled the development of foundational segmentation models capable of accurately segmenting any RGB image. Among these, the Segment

The associate editor coordinating the review of this manuscript and approving it for publication was Yongjie Li.

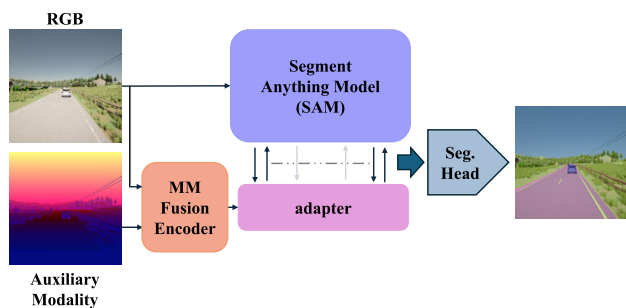


FIGURE 1. We propose adapting the Segment Anything Model’s rich and general knowledge for multimodal semantic segmentation using an external adapter network. Thanks to a Multimodal Fusion Encoder, our adapter leverages both the RGB and auxiliary modality (depth map in the Figure) information to achieve optimal performance on both challenging and easy scenarios.

Anything Model (SAM) [4], trained on 11 million images and 1 billion ground truth masks, stands out for its impressive generalization capabilities. The availability of such strong foundational models for segmentation made us question whether we could employ the segmentation-specific knowledge embedded in models like SAM to tackle multimodal semantic segmentation. Although SAM is designed to perform instance-level segmentation via point or box prompts on RGB images, we conjecture that the fine-grained features dealing with spatial and contextual relationships between objects and parts extracted by the SAM encoder may be used to address pixel-level semantic segmentation even when processing RGB images alongside other modalities. Thus, in this paper, we propose to deploy the rich and generalizable features learned by SAM to achieve multimodal semantic segmentation. In particular, our goal is to adapt SAM to multimodal semantic segmentation while retaining its previous general knowledge.

To achieve our goal, we draw inspiration from the strategy proposed by ViT-adapter [5], which pertains to tailoring a pre-trained Vision Transformer (ViT) [6] for dense tasks by an external adapter network. Based on the mechanism of cross-attention, this design can adapt the foundational ViT [6] features learned through image classification to tasks such as semantic segmentation, avoiding catastrophic forgetting. In our work, we employ a similar external adapter (the pink block in Figure 1), yet our goal is not only to adapt SAM features to a new task, namely from instance segmentation via prompts to semantic segmentation, but also to integrate features extracted from an auxiliary modality, such as, e.g., depth maps, LiDAR measurements, thermal or event data.

Our architecture combines SAM alongside an adapter network. This design follows our intuitions that the RGB features learned by SAM through massive pre-training on billions of images can serve as a strong foundation for multimodal semantic segmentation and that tailoring to the new task without forgetting reusable knowledge about instances and their parts may be achieved effectively through modern adaptation strategies such as [5]. To prioritize its general and rich RGB features, we combine SAM with a lighter adapter

network geared to incorporate the auxiliary modality. This asymmetric design materializes our intuition that, due to the outstanding segmentation capabilities of image-based models like SAM, multimodal semantic segmentation may best be tackled by relying mainly on RGB images and utilizing other modalities to handle challenging situations where the image turns out insufficiently informative. This starkly contrasts traditional multimodal architectures, which typically include networks of similar capacity to process the different modalities, thereby implicitly weighing each modality equally. Our experiments show that the SAM adaptation strategy outlined so far can handle challenging settings effectively. However, they also reveal that introducing an auxiliary modality can sometimes degrade performance when dealing with highly informative RGB content, e.g., images acquired in perfectly lit environments. To address this issue, we propose to adapt SAM using *fused* features computed from *both* the RGB and the auxiliary modality, as illustrated by the orange block in Figure 1. This design allows the model to learn to incorporate auxiliary information only when beneficial. For instance, in RGB-LiDAR segmentation, the model can learn to exploit LiDAR information in low-light settings while relying only on RGB features in perfectly lit environments. In summary, our design harnesses the strengths of foundational RGB models and leverages multimodal information to improve performance in challenging scenarios.

We evaluated our proposal on several multimodal benchmarks such as DeLiVER [7], FMB [8], and MUSES [9], achieving *state-of-the-art* performance in all the considered datasets. Furthermore, we observe that many existing multimodal benchmarks fail to clearly highlight the benefits of auxiliary sensing modalities, as most data samples can be effectively segmented using only RGB information. Hence, we also conducted an evaluation using manually crafted test set splits. Specifically, for both DeLiVER and FMB, we divide the test set into two splits: *RGB-easy* and *RGB-hard* (see Figure 3). The former includes samples in which the RGB image provides enough information for accurate segmentation. Conversely, in the samples assigned to the latter split, the RGB content alone is poorly informative, making it necessary to rely on the auxiliary modality. By evaluating multimodal approaches on these splits, we better assess their ability to integrate information from multiple sensing modalities synergistically. Notably, our approach consistently performs best on both the *RGB-easy* and *RGB-hard* subsets of the DeLiVER and FMB test sets. Furthermore, our multimodal adaptation method, built on the ViT-adapter architecture, outperforms alternative approaches such as LoRA [10], adapted for multimodal input, on the DeLiVER benchmark in the RGB-LiDAR scenario.

To summarize, our contributions are as follows:

- We propose to adapt SAM foundational features to pursue multimodal semantic segmentation. In particular, our multimodal adapter, outlined in Section III-A3, integrates *fused features*, obtained with a fusion module described in Section III-A2, with SAM’s RGB-only features.

- Our approach attains state-of-the-art performance on the DeLiVER [7], FMB [8], and MUSES [9] benchmarks.
- To highlight the effectiveness of multimodal methods in synergistically exploiting the available sensing modalities, we split the DeLiVER and FMB test sets into *RGB-hard* and *RGB-easy* samples and evaluate methods on these splits. Our method performs best in all settings.

The code is available at the following GitHub Repository.

II. RELATED WORK

Our work tackles multimodal semantic segmentation by proposing an adaptation strategy for SAM's foundational features. Hence, in this section, we first review the literature dealing with segmenting images either based on the RGB content or by also deploying auxiliary modalities. Then we pinpoint some recent approaches that attempt to deploy SAM alongside other modalities, as well as previous works aimed at adapting large pre-trained models effectively.

A. IMAGE SEGMENTATION

Image segmentation constitutes a fundamental discipline within computer vision, seeking to partition images into coherent and semantically meaningful regions to facilitate detailed analysis. Unlike conventional image classification, which assigns a single label to an entire image, segmentation operates at the pixel level, enabling a comprehensive understanding of visual content and spatial relationships. Segmentation tasks encompass several distinct paradigms such as semantic segmentation, instance segmentation, panoptic segmentation, and salient object detection. Semantic segmentation assigns a categorical label to each pixel in an image [11]. Instance segmentation [12] distinguishes individual instances of objects belonging to the same class and panoptic segmentation integrates the principles of both semantic and instance segmentation, yielding a unified pixel-wise representation wherein every pixel is attributed either to a specific object instance or a background class. Salient object detection (SOD) [13], [14] concentrates on detecting visually prominent objects within a scene, typically generating binary masks that delineate regions most likely to attract human attention. In this work, we focus on semantic segmentation, which involves predicting a category label for each pixel. Among segmentation During the last years, many semantic segmentation methods have been designed [5], [15], [16], [17], [18].

Among them, the most relevant to our work is the ViT-adapter [5], which concerns tailoring a Vision Transformer trained for ImageNet classification to handle dense tasks by leveraging a lightweight adapter network.

Thanks to the availability of large annotated RGB datasets, some foundational RGB segmentation models have been recently proposed [4], [19], [20]. In particular, the Segment Anything Model (SAM) [4] is designed for prompt-guided instance segmentation, i.e., it can segment objects and their parts in RGB images using prompts such as points, bounding boxes, and masks. Our work aims to rely on SAM's

foundational knowledge to pursue semantic segmentation from multiple modalities.

B. MULTIMODAL SEMANTIC SEGMENTATION

Multimodal semantic segmentation leverages multiple modalities to address situations where individual sensors may fail [2]. This section focuses on methods that include RGB images as one of the modalities. A key design choice is deciding when to fuse RGB with other modalities. Early methods [21] perform channel-wise concatenation of raw input modalities, which are then processed by a single modality-agnostic encoder, while later approaches typically fuse features at multiple levels within the encoder [7], [22], [23], [24], [25], [26], [27]. Recent state-of-the-art models [28], [29], [30] adopt a single late feature fusion stage downstream of the encoder before the decoder. Some methods process the input modalities with a shared modality-agnostic encoder [26], [29], [30], [31] while others with modality-specific encoders [7], [22], [25], [28], [32]. The machinery to combine features, namely the Fusion Module, has evolved over the years, from simple summation or concatenation to attention mechanisms [25], [26]. Our MM Fusion Encoder deploys modality-specific encoders and late feature fusion. However, peculiar to our design, the purpose of the extracted fused multimodal features is to adjust SAM's RGB features through an adapter network. The adapter enables the SAM backbone to incorporate auxiliary information at the encoder level, allowing it to leverage multimodal features for semantic segmentation.

C. SAM BEYOND RGB

In 2024, several studies extended SAM to tasks involving input modalities beyond RGB [33], [34], [35]. Chen et al. [35] distill the SAM backbone into a student network to perform semantic segmentation from event data. Wang et al. [33] propose a salient object detection framework that uses a fusion module as a prompt generator for the SAM decoder. Xiao et al. [34] introduce a multimodal SAM tailored to prompt-based instance segmentation. Yao et al. [36] propose a SAM-adapter for the RGB-Event framework, based on cross-attention and gated blocks, where event data are processed through a domain-specific backbone. Liu et al. [37] leverage the full SAM architecture for RGB-Depth and RGB-Event frameworks in the context of mask segmentation. Differently, we employ the SAM encoder for multimodal semantic segmentation.

D. ADAPTERS

Adapters effectively transfer knowledge from powerful backbones to various tasks. Initially introduced in NLP [38] and later adapted for Computer Vision [39], they enhance a Transformer's general knowledge by integrating external learnable modules. Among these methods, LoRA [10] is one of the earliest approaches designed for parameter-efficient fine-tuning and has also been extended in the vision

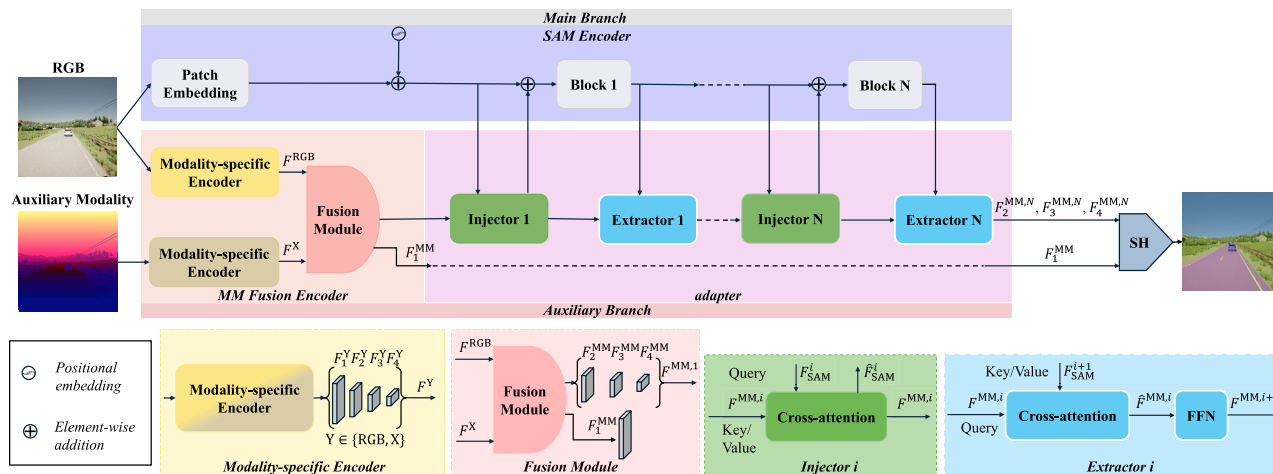


FIGURE 2. Overview of the MM SAM-adapter architecture. The top row shows the four main modules: SAM Encoder, MM Fusion Encoder, adapter and Segmentation Head (SH). The bottom row details the Modality-specific Encoders and the Fusion Module utilized by the MM Fusion Encoder as well as the Injector and Extractors modules utilized by the adapter.

domain. In particular, the ViT-adapter [5] aims to adapt Vision Transformers (ViT) pre-trained on image classification for dense prediction tasks. This architecture utilizes a convolutional encoder to extract spatial features, which are integrated with ViT features through a series of injectors and extractors that maintain continuous interaction with the main backbone. Unlike standard adapters, which often freeze the backbone during training, the ViT-adapter fine-tunes the main ViT encoder to achieve optimal performance. Recently, the adapter mechanism has been explored to customize SAM to various downstream tasks [40], [41], [42], [43]. However, all previous works focus on adapting SAM to perform tasks that require processing only RGB images. In contrast, we are the first to apply an adapter network to adapt SAM for processing multimodal inputs. Our adapter mitigates the risk of catastrophic forgetting and enables the integration of multimodal features into the SAM model.

III. MATERIALS AND METHOD

A. MM SAM-ADAPTER

Our method leverages the benefits of both fusing auxiliary modality features and adaptively integrating knowledge into the SAM backbone. Fusion allows the model to harness auxiliary modality information when the primary modality alone is insufficient for semantic segmentation, while the adaptation technique enables the SAM encoder to incorporate multimodal knowledge, mitigating the risk of catastrophic forgetting in the backbone. Notably, the overall encoder architecture is asymmetric, as the SAM backbone has a larger number of parameters compared to the combined adapter module and Multimodal Fusion Encoder. The scenario in which this parameter gap is reduced is analyzed in Section IV-C2. Thus, our approach prioritizes the foundational knowledge embedded in the SAM RGB backbone, while exploiting the *multimodal fused* knowledge captured by the Multimodal Fusion Encoder as well. These advantages

enable the model to achieve *state-of-the-art* performance across various benchmarks as discussed in Section IV and to perform best in *RGB-easy* and *RGB-hard* scenarios. As shown in Figure 2, our multimodal semantic segmentation framework, MM SAM-adapter, comprises a main branch with the SAM Image Encoder, a secondary branch featuring a Multimodal Fusion Encoder and an adapter module, and a Segmentation Head (SH). The following sections provide a detailed description of these components.

1) SAM ENCODER

Our goal is to leverage the segmentation-specific knowledge embedded in SAM’s weights. The overall SAM architecture consists of three main modules: an image encoder, a prompt encoder, and a decoder. As the prompt encoder and the decoder are tailored specifically for the prompt-guided instance segmentation task, we deploy only the image encoder in our multimodal semantic segmentation framework. The SAM image encoder is built upon a ViT architecture, which has been pre-trained in a self-supervised manner as a Masked AutoEncoder [44] and then fine-tuned on the SA-1B dataset to perform prompt-guided instance segmentation. We utilize the SAM encoder based on the ViT-Large (ViT-L) architecture, which includes $L = 24$ layers. As presented in Figure 2 and Section III-A3, akin to the architecture proposed in [5], we divide the encoder into $N = 4$ Blocks consisting of $\frac{L}{N}$ layers, each Block interacting with the adapter by a pair of Injector-Extractor modules.

2) MULTIMODAL FUSION ENCODER

As discussed in Section I, our method relies initially on a fusion stage, where RGB features are combined with auxiliary modality features, producing multimodal representations that are subsequently fed into the adapter module. To achieve this, we employ a Multimodal Fusion Encoder that processes the RGB image and the auxiliary measurements as inputs

to generate fused features. Thereby, the Fusion Encoder can learn to weigh the contribution of the different modalities within the signal provided to the adapter, which, in turn, can learn how to integrate it with SAM's features. Our Multimodal Fusion Encoder comprises three main modules: two modality-specific encoders and one fusion module, as described below.

a: MODALITY-SPECIFIC ENCODERS

The modality-specific encoders extract features independently from each of the two modalities to better handle their different nature: e.g., RGB images and LiDAR measurements are dense and sparse signals, respectively. Inspired by Vit-adapter [5], we employ convolutional networks as they yield spatial features at multiple resolutions that help transformers better capture local spatial information. In particular, we employ ConvNext Small [45] pre-trained on ImageNet-22k [3] for both modality-specific encoders. To ensure compatibility, when the auxiliary signal has a single channel, we replicated it three times before feeding it to ConvNext. Hence, given an RGB image $I^{\text{RGB}} \in \mathbb{R}^{H \times W \times 3}$ and a pixel-aligned auxiliary signal $I^X \in \mathbb{R}^{H \times W \times 3}$, both encoders produce four different spatial features, namely F_i^X and F_i^{RGB} with $i = 1, 2, 3, 4$ for the RGB and auxiliary modality encoder, respectively. The feature tensor size at resolution i is $\frac{H}{2^{i+1}} \times \frac{W}{2^{i+1}} \times D_i$, where D_i denotes the number of channels (e.g., in ConvNext Small 96, 128, 384, 768).

b: FUSION MODULE

The Fusion Module processes multi-scale modality-specific features (F_i^X, F_i^{RGB}) and generates fused multi-scale outputs, F_i^{MM} , for $i = 1, 2, 3, 4$. To align with SAM's feature dimensions, we apply four linear projection layers to match the channel dimension D . The fusion module must preserve information from both modalities to allow the adapter to dynamically select the relevant ones during inference, e.g., only the RGB information in optimal conditions. As shown in Section IV-C3, to achieve this goal, we can employ a simple feature concatenation of the RGB and auxiliary features, which preserves each modality's information. However, to obtain the best performance, we adopt the Road-Fusion module from Roadformer+ [29], which combines convolutional, self-attention, spatial-attention, and coordinate-attention layers for advanced heterogeneous feature fusion.

3) ADAPTER

The adapter module takes as input the multimodal features produced by the Multimodal Fusion Encoder and incorporates that knowledge within the SAM backbone. Akin to the ViT-adapter [5], our adapter is a side-tuning network – a parallel branch that processes information alongside the main backbone. This strategy enables the adapter to learn domain-specific knowledge, in this case, multimodal information, while mitigating the risk of catastrophic forgetting in the

main backbone. However, unlike the original proposal in [5], our adapter processes fused multimodal features rather than RGB-only features.

The adapter consists of a series of two modules: an injector, which introduces spatial multimodal knowledge to the SAM encoder, and an extractor, which retrieves hierarchical features from the SAM backbone.

a: INJECTOR

The $F_2^{\text{MM}}, F_3^{\text{MM}}$, and F_4^{MM} multi-scale multimodal fused features are flattened and stacked channel-wise to obtain $F^{\text{MM},1}$, a 2-D array of size $(\frac{HW}{8^2} + \frac{HW}{16^2} + \frac{HW}{32^2}) \times D$. $F^{\text{MM},1}$ will be the input to the first Injector module. Each injector, *Injector* i , computes a cross-attention between the SAM features, which are output of the block $i - 1$, F_{SAM}^i with dimensionality $\frac{HW}{16^2} \times D$, and $F^{\text{MM},i}$, the multiscale multimodal fused features after being processed by $i - 1$ injectors and extractors. In particular, F_{SAM}^i is taken as the query and $F^{\text{MM},i}$ is used as the key and value of a multi-scale deformable cross-attention [46]. Hence, in our adapter, the computation performed by an Injector can be expressed as:

$$\hat{F}_{\text{SAM}}^i = F_{\text{SAM}}^i + \gamma_i \text{Attn}(\Phi(F_{\text{SAM}}^i), \Phi(F^{\text{MM},i})) \quad (1)$$

where $\Phi(\cdot)$ denotes the layer normalization operation [47], $\text{Attn}(\cdot)$ represents the multi-scale deformable cross-attention, and γ_i is a D -dimensional learnable vector initialized to zero to mitigate the impact of the cross-attention at the beginning of the training process. The output from the i -th injector is then processed by the i -th SAM block, obtaining F_{SAM}^{i+1} with dimensionality $\frac{HW}{16^2} \times D$.

b: EXTRACTOR

This component operates on the output of SAM's blocks. The multi-scale deformable cross-attention [46] is computed between SAM's features, F_{SAM}^{i+1} , taken as key and values, and the multi-scale fused features $F^{\text{MM},i}$, taken as query. The features obtained by the cross-attention, $\hat{F}^{\text{MM},i}$, are fed to a feed-forward network (FFN), obtaining the final features $F^{\text{MM},i+1}$, which will be the input to the next injector module. Hence, the computation performed by an extractor of our adapter can be expressed as:

$$\hat{F}^{\text{MM},i} = F^{\text{MM},i} + \text{Attn}(\Phi(F^{\text{MM},i}), \Phi(F_{\text{SAM}}^{i+1})) \quad (2)$$

$$F^{\text{MM},i+1} = \hat{F}^{\text{MM},i} + \text{FFN}(\Phi(\hat{F}^{\text{MM},i})) \quad (3)$$

4) SEGMENTATION HEAD

We obtain the predictions by a Segmentation Head (SH) implemented as a Segformer [18] decoder. It processes the multimodal fused features, F_1^{MM} , with the adapter tokenized output, $F^{\text{MM},N}$, which is reshaped to obtain feature maps $F_i^{\text{MM},N}$, $i = 2, 3, 4$ at their original input resolutions.

B. DATASETS

We employ the following multimodal semantic segmentation datasets in our experiments:

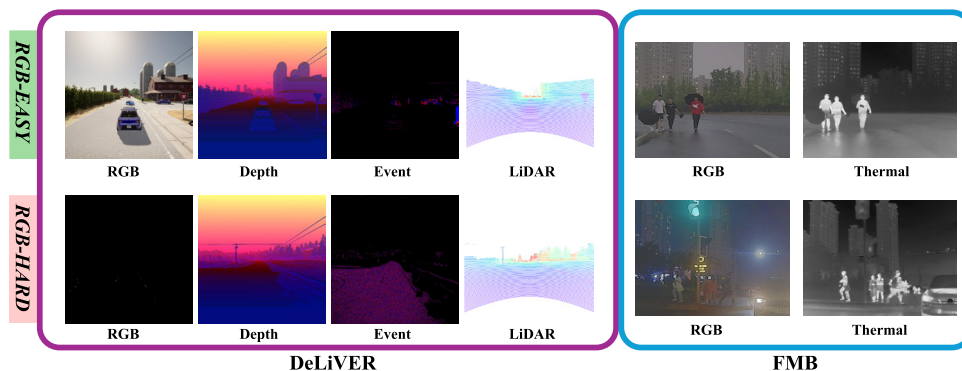


FIGURE 3. *RGB-easy* and *RGB-hard* examples for DeLiVER and FMB. The DeLiVER *RGB-hard* image has been stretched to enhance visibility (zoom in to better notice it).

DeLiVER [7] is a synthetic dataset comprising 3,983 training, 2,005 validation, and 1,897 test samples. It includes LiDAR, Depth, Event, and RGB data from urban environments with diverse lighting and weather conditions. The tensors containing the auxiliary modalities have the same spatial dimensions as the RGB images (1042×1042), with data from all modalities being pixel-wise aligned. The number of classes in DeLiVER is 25.

FMB [8] contains 1,500 pixel-aligned RGB-Thermal image pairs, 280 of which serve as the test set. The image pairs are acquired in urban driving scenes with different illumination and weather conditions and have spatial dimensions 800×600 . The number of classes is 14. To select the best checkpoint for all models, we create a validation set by randomly extracting 160 samples from the training set.

MUSES [9] consists of 1,500 training, 250 validation and 750 test samples collected under various daytime and weather conditions. It includes LiDAR, Event, Radar, and RGB data, though we excluded Radar due to its sparsity and insufficient information for multimodal segmentation. We align the data from all modalities at pixel level by using the official code [9]. MUSES employs the same 19 classes as the RGB-only Cityscapes dataset [48]. As the test annotations are withheld, test set evaluations are possible only through online submissions.

1) RGB-EASY AND RGB-HARD SPLITS

As highlighted in Section I, most samples in multimodal datasets are easily segmented by leveraging exclusively RGB information. For instance, from our experiments, approximately 97% of DeLiVER sample can be optimally segmented using solely the RGB image, while only the remaining 3% represents challenging scenarios in which the auxiliary information may be useful. To better assess the methods' capability to exploit information from multiple modalities synergistically, we divided the DeLiVER and FMB test sets into an *RGB-hard* and *RGB-easy* split. Purposely, we first train an RGB-only network on each multimodal dataset to obtain semantic predictions. Then, we visually compare the RGB-only predictions to the ground-truths while paying

attention to the information conveyed by the auxiliary modality. Thereby, we identify as *RGB-hard* those samples in which the RGB-only network fails to recognize elements that may be correctly detected by another modality, and as *RGB-easy* all other samples. Accordingly, the DeLiVER test set is divided into 1797 *RGB-easy* and 100 *RGB-hard* samples, whereas for FMB we get 183 and 97 samples for the *RGB-easy* and *RGB-hard* splits. Figure 3 shows samples for the *RGB-easy* and *RGB-hard* splits of DeLiVER and FMB. We point out that we adopt a similar procedure to build a balanced validation set of FMB, i.e., we randomly select 80 *RGB-easy* and 80 *RGB-hard* samples from the original training set. We did not perform a similar test set split on MUSES [9] as the ground truths are unavailable. Yet, we report results for all the official splits (day, night, different weather conditions). In particular, this disentangled evaluation allows us to assess the behaviour of the methods in the fog and nighttime scenarios – those that pose more challenges for the RGB modality, as highlighted in the original paper [9].

C. IMPLEMENTATION DETAILS

1) ARCHITECTURAL DETAILS

The architectural specifications are summarized in Table 1 and Table 2, which offers a detailed overview of our network design. A legend explaining the abbreviations used in Table 1 and Table 2 is provided in Table 3.

The **Fusion Module**, as detailed in Table 1 and employed throughout this paper, is based on the RoadFormer+ fusion block [29] (Road-fusion module). This module comprises several key components, including two Global Feature Enhancers (GFEs), two Local Feature Enhancers (LFEs), one Global Feature Recalibration Module (GFRM), one Local Feature Fusion Module (LFFM) and one Feature Enhancement and Integration Module (FEIM). The Road-fusion block initially extracts global and local features separately from RGB features and auxiliary modality features, both originating from the two Modality-Specific Encoders, by employing GFEs and LFEs. The former are a transformer-based modules, while the latter are convolution-based. Subsequently,

TABLE 1. Details of our architecture (SAM Encoder, Modality-Specific Encoders, Fusion Module). H and W represents the Height and Width of the input image, $i \in [1, 4]$, $N = 4$ and N_{classes} is the number of classes. SAM-L encoder has an embedding dimension of 1024.

Stage	Layer Type	Input	Output	Output Shape
SAM Encoder				
Input	Image	-	I^{RGB}	$H \times W \times 3$
Emb	Patch E., Pos. E.	I^{RGB}	F_1^{SAM}	$\frac{HW}{16^2} \times 1024$
Block i	-	\hat{F}_i^{SAM}	F_{i+1}^{SAM}	$\frac{HW}{16^2} \times 1024$
Modality-Specific Encoders				
Input	Image Aux. mod.	-	I^{RGB} I^{X}	$H \times W \times 3$ $H \times W \times 3$
RGB Enc.	ConvNeXt	I^{RGB}	F_1^{RGB} F_2^{RGB} F_3^{RGB} F_4^{RGB}	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
X Enc.	ConvNeXt	I^{X}	F_1^{X} F_2^{X} F_3^{X} F_4^{X}	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
Fusion module				
GFE-RGB	Patch E., MHSA, LN	F_1^{RGB} F_2^{RGB} F_3^{RGB} F_4^{RGB}	$F_1^{\text{RGB,G}}$ $F_2^{\text{RGB,G}}$ $F_3^{\text{RGB,G}}$ $F_4^{\text{RGB,G}}$	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
GFE-X	Patch E., MHSA, LN	F_1^{X} F_2^{X} F_3^{X} F_4^{X}	$F_1^{\text{X,G}}$ $F_2^{\text{X,G}}$ $F_3^{\text{X,G}}$ $F_4^{\text{X,G}}$	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
LFE-RGB	Conv $_{1 \times 1}$, RL, DWC $_{3 \times 3}$, RL, Conv $_{1 \times 1}$	F_1^{RGB} F_2^{RGB} F_3^{RGB} F_4^{RGB}	$F_1^{\text{RGB,L}}$ $F_2^{\text{RGB,L}}$ $F_3^{\text{RGB,L}}$ $F_4^{\text{RGB,L}}$	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
LFE-X	Conv $_{1 \times 1}$, RL, DWC $_{3 \times 3}$, RL, Conv $_{1 \times 1}$	F_1^{X} F_2^{X} F_3^{X} F_4^{X}	$F_1^{\text{X,L}}$ $F_2^{\text{X,L}}$ $F_3^{\text{X,L}}$ $F_4^{\text{X,L}}$	$\frac{H}{4} \times \frac{W}{4} \times 96$ $\frac{H}{8} \times \frac{W}{8} \times 192$ $\frac{H}{16} \times \frac{W}{16} \times 384$ $\frac{H}{32} \times \frac{W}{32} \times 768$
GFRM	CA RGB-X, CA X-RGB, Concat., LN, Av. Pool., Sigm.	$F_1^{\text{RGB,G}}, F_1^{\text{X,G}}$ $F_2^{\text{RGB,G}}, F_2^{\text{X,G}}$ $F_3^{\text{RGB,G}}, F_3^{\text{X,G}}$ $F_4^{\text{RGB,G}}, F_4^{\text{X,G}}$	$F_1^{\text{MM,G}}$ $F_2^{\text{MM,G}}$ $F_3^{\text{MM,G}}$ $F_4^{\text{MM,G}}$	$\frac{H}{4} \times \frac{W}{4} \times 192$ $\frac{H}{8} \times \frac{W}{8} \times 384$ $\frac{H}{16} \times \frac{W}{16} \times 768$ $\frac{H}{32} \times \frac{W}{32} \times 1536$
LFFM	Concat., Conv $_{1 \times 1}$ DWC $_{3 \times 3}$, GL, Conv $_{1 \times 1}$	$F_1^{\text{RGB,L}}, F_1^{\text{X,L}}$ $F_2^{\text{RGB,L}}, F_2^{\text{X,L}}$ $F_3^{\text{RGB,L}}, F_3^{\text{X,L}}$ $F_4^{\text{RGB,L}}, F_4^{\text{X,L}}$	$F_1^{\text{MM,L}}$ $F_2^{\text{MM,L}}$ $F_3^{\text{MM,L}}$ $F_4^{\text{MM,L}}$	$\frac{H}{4} \times \frac{W}{4} \times 192$ $\frac{H}{8} \times \frac{W}{8} \times 384$ $\frac{H}{16} \times \frac{W}{16} \times 768$ $\frac{H}{32} \times \frac{W}{32} \times 1536$
FEIM	W. Feat. Sum, Coord. A.	$F_1^{\text{MM,G}}, F_1^{\text{MM,L}}$ $F_2^{\text{MM,G}}, F_2^{\text{MM,L}}$ $F_3^{\text{MM,G}}, F_3^{\text{MM,L}}$ $F_4^{\text{MM,G}}, F_4^{\text{MM,L}}$	$F_1^{\text{MM,F}}$ $F_2^{\text{MM,F}}$ $F_3^{\text{MM,F}}$ $F_4^{\text{MM,F}}$	$\frac{H}{4} \times \frac{W}{4} \times 192$ $\frac{H}{8} \times \frac{W}{8} \times 384$ $\frac{H}{16} \times \frac{W}{16} \times 768$ $\frac{H}{32} \times \frac{W}{32} \times 1536$

TABLE 2. Details of our architecture (Projector, Injector, Extractor, Feature Refinement, Segmentation Head). H and W represent the Height and Width of the input image, $i \in [1, 4]$, $N = 4$ and N_{classes} is the number of classes. SAM-L encoder has an embedding dimension of 1024.

Stage	Layer Type	Input	Output	Output Shape
Projector				
S-to-V	Conv $_{1 \times 1}$, Flatten	$F_1^{\text{MM},F}$ $F_2^{\text{MM},F}$ $F_3^{\text{MM},F}$ $F_4^{\text{MM},F}$	F_1^{MM} F_2^{MM} F_3^{MM} F_4^{MM}	$\frac{HW}{4^2} \times 1024$ $\frac{HW}{8^2} \times 1024$ $\frac{HW}{16^2} \times 1024$ $\frac{HW}{32^2} \times 1024$
Feat. Stack.	Concat.	F_2^{MM} F_3^{MM} F_4^{MM}	$F^{\text{MM},1}$	$(\frac{HW}{8^2} + \frac{HW}{16^2} + \frac{HW}{32^2}) \times 1024$
Injector i				
Q-K Inter	MSDA W. Feat. Sum	$F_i^{\text{SAM}}, F^{\text{MM},i}$	\hat{F}_i^{SAM}	$\frac{HW}{16^2} \times 1024$
Extractor i				
Q-K Inter	MSDA	$F^{\text{MM},i}, F_{i+1}^{\text{SAM}}$	$\hat{F}^{\text{MM},i}$	$(\frac{HW}{8^2} + \frac{HW}{16^2} + \frac{HW}{32^2}) \times 1024$
FFN	Lin., DWC $_{3 \times 3}$ GL, Lin.	$\hat{F}^{\text{MM},i}$	$F^{\text{MM},i+1}$	$(\frac{HW}{8^2} + \frac{HW}{16^2} + \frac{HW}{32^2}) \times 1024$
Feature Refinement				
Slice	Reshape	$F^{\text{MM},N}$	$F_2^{\text{MM},N}$ $F_3^{\text{MM},N}$ $F_4^{\text{MM},N}$	$\frac{H}{8} \times \frac{W}{8} \times 1024$ $\frac{H}{16} \times \frac{W}{16} \times 1024$ $\frac{H}{32} \times \frac{W}{32} \times 1024$
Segmentation Head				
Prep.	Tran. Conv. & Sum, Feature Sum	$F_2^{\text{MM},N}, F_1^{\text{MM}}$ $F_1^{\text{MM},N}, F_N^{\text{SAM}}$ $F_2^{\text{MM},N}, F_N^{\text{SAM}}$ $F_3^{\text{MM},N}, F_N^{\text{SAM}}$ $F_4^{\text{MM},N}, F_N^{\text{SAM}}$	$F_1^{\text{MM},N}$ $F_{1,\text{mix}}^{\text{MM},N}$ $F_{2,\text{mix}}^{\text{MM},N}$ $F_{3,\text{mix}}^{\text{MM},N}$ $F_{4,\text{mix}}^{\text{MM},N}$	$\frac{H}{4} \times \frac{W}{4} \times 1024$ $\frac{H}{4} \times \frac{W}{4} \times 1024$ $\frac{H}{8} \times \frac{W}{8} \times 1024$ $\frac{H}{16} \times \frac{W}{16} \times 1024$ $\frac{H}{32} \times \frac{W}{32} \times 1024$
up	Conv $_{1 \times 1}$, BN, RL, Interp., Concat.	$F_{1,\text{mix}}^{\text{MM},N}$ $F_{2,\text{mix}}^{\text{MM},N}$ $F_{3,\text{mix}}^{\text{MM},N}$ $F_{4,\text{mix}}^{\text{MM},N}$	$F_{\text{up}}^{\text{MM}}$	$\frac{H}{4} \times \frac{W}{4} \times 512$
Pred.	Conv $_{1 \times 1}$, BN, RL, Conv $_{1 \times 1}$, Interp.	$F_{\text{up}}^{\text{MM}}$	$F_{\text{Seg}}^{\text{MM}}$	$H \times W \times N_{\text{classes}}$

the global RGB features and auxiliary modality features are fused using the GFRM. This module performs cross-attention operations between the two modalities: one operation utilizes RGB modality features as queries and auxiliary modality features as keys/values, whereas the other operation reverses these roles. The resulting global feature representations are then concatenated, and the obtained multimodal representation is further refined through an average pooling

layer followed by a sigmoid activation function. In parallel with the global features extraction, the local features from both modalities are fused using the LFFM. This module first performs a channel-wise concatenation, after which the resulting feature map is processed through a series of convolutional layers employing one GeLU activation function. Subsequently, local and global fused features are combined through the GFRM. This module employs

TABLE 3. Legend of abbreviations used in the Table 1 and Table 2.

Abbreviation	Meaning
Emb	Embedding
Patch E.	Patch Embedding
Pos. E.	Positional Embedding
Aux mod	Auxiliary Modality
RGB Enc.	RGB Encoder
X Enc.	X Encoder
GFE	Global Feature Enhancer
MHSA	Multi-Head Self-Attention
LN	Layer normalization
Concat.	Concatenation
LFE	Local Feature Enhancer
DWC _{3×3}	Depth-wise Separable Convolution
Conv _{1×1}	1 × 1 Convolution
GFRM	Global Feature Recalibration Module
Sigm.	Sigmoid
Av. Pool.	Average Pooling
CA	Cross-Attention
LFFM	Local Feature Fusion Module
FEIM	Feature Enhancement and Integration Module
W. Feat. Sum	Weighted Feature Sum
Coord. A.	Coordinate Attention
MSDA	Multi-Scale Deformable Attention
BN	Batch Normalization
Prep.	Preprocessing
Up	Upsampling
Pred.	Prediction
Tran. Conv.	Transpose Convolution
Q-K Inter	Query-Key Interaction
S-to-V	Spatial-to-Vector
Feat. Stack.	Feature Stacking
Interp.	Interpolation
Lin.	Linear
RL	ReLU
GL	GeLU

coordinate attention to effectively encode spatial relationships within the features, enhancing their representational capacity. The resulting attention-enhanced representation is further refined through a series of convolutional layers. The output of the fusion module is subsequently passed through a series of projection layers to align the channel dimensions of the fused features with those of SAM.

The **Segmentation Head** (SH) will take as input F_1^{MM} and the adapter output, $F^{MM,N}$ reshaped into $F_i^{MM,N}$, $i = 2, 3, 4$. The F_1^{MM} is summed with $F_2^{MM,N}$, which has been upsampled using transpose convolution [49] to match the spatial resolution, resulting in $F_1^{MM,N}$. Each spatial feature is subsequently enriched by summing it with the SAM tokens from the final layer, which are first interpolated to match the spatial resolution. This results in the mixed feature representation $F_{i,mix}^{MM,N}$, where $i = 2, 3, 4$. Subsequently, the Segformer head [18] processes each spatial feature $F_{i,mix}^{MM,N}$, using distinct MLPs, composed of 1×1 convolutional layers, batch normalization layers and ReLU activation functions. A channel-wise concatenation operation with bilinear interpolation is then applied, yielding F_{up}^{MM} with

dimensionality $\frac{H}{4} \times \frac{W}{4} \times 512$. Finally, in the last stage, another MLP is used to generate the segmentation prediction, with dimensionality $H \times W \times N_{classes}$, where $N_{classes}$ denotes the number of classes.

2) TRAINING DETAILS

Our method was trained for 100 epochs using the OHEM cross-entropy loss on two NVIDIA RTX 3090 GPUs with batch size 8 – we employ gradient accumulation for memory constraints –, a base learning rate of $2e^{-4}$, which follows a polynomial strategy with a power of 0.9, and exponential warm-up with base 0.1. Therefore, the learning rate evolves according to the following formula:

$$\eta(p) = \begin{cases} \eta_{base} \cdot (wr)^{1 - \frac{p}{N_w}}, & \text{if } p \leq N_w, \\ (\eta_{base} - \eta_{min}) \cdot \left(1 - \frac{p}{P_{max}}\right)^\alpha + \eta_{min}, & \text{if } p > N_w. \end{cases} \quad (4)$$

In the learning rate schedule described above, $\eta(p)$ denotes the learning rate at epoch or iteration p . The parameter η_{base} represents the base learning rate that the schedule aims to reach at the end of the warm-up phase ($2e^{-4}$), while η_{min} defines the minimum learning rate that serves as a lower bound throughout training ($\eta_{min} = 0$). During the warm-up phase, which lasts for N_w epochs ($N_w = 10$), the learning rate increases exponentially from an initial value determined by scaling η_{base} by the factor wr , which represent the warm-up ratio ($wr = 0.1$). Specifically, at the beginning of training, the learning rate is set to $\eta_{base} \times wr$, and it grows smoothly and exponentially toward η_{base} as training progresses through the warm-up steps. After the warm-up period, the learning rate follows a polynomial decay controlled by the exponent α ($\alpha = 0.9$), which determines how sharply the learning rate decreases. The variable P_{max} specifies the total number of epochs in the training process, ensuring that the schedule spans the entire duration of learning without abrupt transitions. The optimizer used during training is AdamW [50] with a weight decay of $1e^{-2}$. We use a layer-wise learning rate decay of 0.9, as in ViT-adapter architecture [5]. In this approach, the learning rate assigned to each layer decreases exponentially as one moves from the higher layers toward the lower layers of the backbone. The learning rate for the parameters in layer ℓ is defined as:

$$\eta_\ell = \eta_{curr} \times \gamma^{L-\ell-1}, \quad (5)$$

where η_{curr} denotes the learning rate assigned to the highest layer, γ is the decay rate ($\gamma = 0.9$), L is the total number of layers of the transformer backbone, and ℓ is the index of the current layer, starting from zero.

Parameters associated with embeddings, special tokens, and normalization layers are assigned to layer zero and are typically exempted from weight decay. Conversely, the newly introduced layers in the architecture, such as adapter modules and decoding heads, are trained with higher learning

TABLE 4. DeLiVER test set results in the RGB-Depth (RGB-D), RGB-LiDAR (RGB-L), and RGB-Event (RGB-E) setups.

Method	All			RGB-easy			RGB-hard		
	RGB-D	RGB-L	RGB-E	RGB-D	RGB-L	RGB-E	RGB-D	RGB-L	RGB-E
CMNeXt [7]	53.87	51.32	50.81	54.07	51.97	51.56	49.66	39.29	37.08
GeminiFusion [26]	54.98	50.57	51.71	55.07	51.07	52.32	53.00	40.93	40.27
RoadFormer+ [29]	55.95	54.56	54.10	56.12	55.41	54.80	52.39	40.26	41.98
MM SAM-adapter	57.35	57.14	55.70	57.62	57.75	56.29	53.35	45.46	44.67

TABLE 5. FMB test set results in the RGB-Thermal (RGB-T) setup across different scenarios.

Method	All	RGB-easy	RGB-hard
	RGB-T	RGB-T	RGB-T
CMNeXt [7]	61.66	64.81	56.85
GeminiFusion [26]	64.75	68.05	61.03
RoadFormer+ [29]	64.57	66.31	61.79
MM SAM-adapter	66.10	68.45	62.59

rates to encourage rapid adaptation to the downstream task. This parameter grouping ensures that lower layers of the backbone, which encode generic visual features learned during pre-training, are updated more conservatively, thereby preserving their valuable representations. In order to improve robustness of our model, we perform online data augmentations based on random resize with a ratio from 0.5 to 2.0, random horizontal flipping, photometric distortion, random Gaussian blur with probability $p = 0.2$, and random crop to 1024×1024 resolution for DeLiVER and MUSES, 800×800 for FMB. Our method on FMB has been trained for 200 epochs.

3) COMPETITOR DETAILS

For competitors evaluation, we use official weights when available; otherwise, we train them following the official paper guidelines. In the **DeLiVER** benchmark, CMNeXt and GeminiFusion were evaluated using the official weights, which are based on the most effective backbone, MiT-B2 [18]. In contrast, we trained RoadFormer+ from scratch, as no pre-trained weights were available, using ConvNeXt-Large, the best-performing configuration reported in the paper. For the **FMB** [8] and **MUSES** [9] datasets, we had to train all competitors due to the unavailability of pre-trained weights. Specifically, for **FMB**, we followed the official guidelines to ensure optimal performance, training CMNeXt with the MiT-B5 [18] backbone, GeminiFusion with Swin Transformer Large [51], and RoadFormer+ with ConvNeXt-Large [45]. Regarding **MUSES**, since it was recently published with no available training guidelines, we adopted the same configuration as **DeLiVER**, given the similarity between the datasets in terms of modalities and high-resolution images. Consequently, we trained CMNeXt and GeminiFusion with the MiT-B2 backbone and RoadFormer+ with the ConvNeXt-Large backbone. CAFuser [30], according to its authors, is inherently a multimodal model.

Moreover, its training requires an additional text prompt describing the condition of the scene, which is not provided in **FMB**, but can be implicitly inferred from **DeLiVER** and **MUSES**. Thus, the results are reported following the original paper's settings.

We highlight that, differently from most competitors that employ different backbones depending on the dataset to achieve the best performance, such as CMNeXt [7] and GeminiFusion [26], our framework remains consistent across datasets yet achieves superior performance.

IV. RESULTS

This section compares MM SAM-adapter with *state-of-the-art* multimodal semantic segmentation models such as CMNext [7], GeminiFusion [26], RoadFormer+ [29], and CAFuser [30]. We use Mean Intersection Over Union (mIoU) as the semantic segmentation evaluation metric. In addition, we analyze the design choices underlying the proposed multimodal adapter.

A. MAIN RESULTS

Table 4 reports the results on the DeLiVER test set in three multimodal setups: RGB-Depth (RGB-D), RGB-LiDAR (RGB-L), and RGB-Events (RGB-E). By analyzing the performance in the *RGB-hard* case, we note that methods are more effective when employing the Depth auxiliary modality rather than Event or LiDAR data. We ascribe it to the perfect synthetic Depth maps provided with DeLiVER, which convey extremely rich auxiliary information. In this RGB-D scenario, the improvements over competitors are marginal (e.g., 53.35 vs. 53.00 mIoU of our method vs. GeminiFusion). However, our method shines when the auxiliary modality is more noisy and realistic, such as in the RGB-L and RGB-E setups. In these cases, we achieve large performance improvements over competitors (e.g., 45.46 vs. 40.93 mIoU of our method vs. GeminiFusion in the RGB-L scenario). Finally, by examining our method's results in the *RGB-easy* scenario across the RGB-D, RGB-L, and RGB-E setups, we note that it achieves very similar performance independently of the auxiliary modality (i.e., 57.62, 57.75, 56.29 mIoU in the RGB-D, RGB-L, and RGB-E setups, respectively). Conversely, other methods generally perform less consistently (e.g., CMNeXt obtains 54.07, 51.97, and 51.56 in the RGB-D, RGB-L, and RGB-E setups, respectively). In general, these results highlight how our method more effectively utilizes the synergies between

TABLE 6. MUSES test results in the RGB-LiDAR (RGB-L) and RGB-Event (RGB-E) setups for different weather conditions.

Method	All		Day		Night		Clear		Fog		Rain		Snow	
	RGB-L	RGB-E	RGB-L	RGB-E	RGB-L	RGB-E	RGB-L	RGB-E	RGB-L	RGB-E	RGB-L	RGB-E	RGB-L	RGB-E
CMNeXt [7]	72.36	70.49	75.42	74.02	64.27	60.72	72.81	71.24	59.98	60.91	72.48	68.78	71.92	70.54
GeminiFusion [26]	74.22	68.62	75.90	72.04	67.67	58.19	74.79	70.40	62.95	60.53	73.30	65.27	74.16	68.87
RoadFormer+ [29]	80.38	77.70	82.56	80.62	74.51	69.02	79.39	78.44	69.64	69.29	80.69	76.40	79.49	77.70
MM SAM-adapter	81.07	79.92	83.34	83.39	74.97	72.38	80.82	81.52	74.12	68.97	80.00	78.92	80.85	78.72

TABLE 7. Comparison with all-modalities competitors on the DELIVER test (Test mIoU) and validation (Val mIoU) sets.

Method	Modalities	Val mIoU	Test mIoU
CMNeXt [7]	RGB, Depth, LiDAR, Event	66.30	53.00
GeminiFusion [26]	RGB, Depth, LiDAR, Event	66.90	54.46
CAFuser [30]	RGB, Depth, LiDAR, Event	67.80	55.60
MM SAM-adapter	RGB, LiDAR	61.89	57.14
MM SAM-adapter	RGB, Event	60.74	55.70
MM SAM-adapter	RGB, Depth	69.60	57.35

TABLE 8. Comparison with all-modalities competitors on the MUSES test set.

Method	Modalities	mIoU
CMNeXt	RGB, LiDAR, Event, Radar	72.40
GeminiFusion	RGB, LiDAR, Event, Radar	75.30
CAFuser	RGB, LiDAR, Event, Radar	78.18
MM SAM-adapter	RGB, Event	79.92
MM SAM-adapter	RGB, LiDAR	81.07

modalities, e.g., in the case of good RGB images, the auxiliary modality contribute less to final results. Our method excels at fusing multimodal features and seamlessly adapting the SAM RGB backbone to incorporate this information. As can be seen from these results, our method establishes a new *state-of-the-art* on the DeLiVER test set in the RGB-D, RGB-L, and RGB-E setups, and achieves remarkable performance in both *RGB-easy* and *RGB-hard* scenarios.

Table 5 shows the results on the FMB test set, which features RGB and Thermal images (RGB-T). Remarkably, our method achieves *state-of-the-art* performance also in this real dataset. We note a significant performance drop between *RGB-easy* and *RGB-hard* samples, which validates our manually defined splits. As evidenced by the results, in the *RGB-hard* scenario, our method demonstrates superior exploitation of auxiliary modality information, yielding a measurable performance gain (e.g., 62.59 vs. 61.79 mIoU of our method vs RoadFormer+). In the *RGB-easy* scenario, our approach likewise surpasses competing methods, though with a narrower margin (e.g., 68.45 vs. 68.05 mIoU of our method vs. GeminiFusion).

Table 6 presents the result on the MUSES dataset [9] for the RGB-LiDAR (RGB-L) and RGB-Event (RGB-E) setups. We highlight that the results of each model were submitted to the MUSES online benchmark to construct this table. Notably, we achieve *state-of-the-art* performance overall and in most daytime and weather conditions. In particular, we highlight how, in the foggy scenario, one of the most challenging for the RGB modality alone, our method largely surpasses the best competitor performance in the

RGB-L setup (i.e., 74.12 vs 69.64 mIoU of our method vs RoadFormer+) and provides by far the most effective solution to tackle this difficult setting. In the nighttime scenario—also a challenging condition—our method consistently outperforms prior approaches, achieving higher mIoU scores (e.g., 72.48 vs. 69.02 mIoU in the RGB-E setting and 74.97 vs. 74.51 mIoU in the RGB-L setting, of our method vs. RoadFormer+). Remarkably, our method demonstrates strong performance under daytime conditions, a scenario where RGB images are particularly informative (e.g., 83.39 vs. 80.62 mIoU in the RGB-E Day setting, and 83.34 vs. 82.56 mIoU in the RGB-L Day setting, of our method vs RoadFormer+).

Qualitative results of our method and the competing models are shown in Figure 4, in Figure 5. A closer inspection of the qualitative results in Figure 4 reveals that our method demonstrates a superior ability to capture fine-grained scene details, including road features (e.g., zebra crossings and lane markings), sidewalk elements (e.g., streetlights), and subtle distinctions between contiguous surfaces (e.g., terrain and concrete). Notably, our method produces more accurate qualitative predictions than competing approaches in both the *RGB-hard* and *RGB-easy* settings. In Figure 5, which illustrates examples from both FMB and MUSES under the *RGB-hard* condition, our model exhibits a stronger understanding of the overall environment and background elements. In particular, the building partially occluded by trees in the FMB *RGB-hard* scenario is segmented with noticeably higher precision compared to the competing methods. These qualitative results highlight the robustness of our method in various scenarios and datasets, aligning with the findings from the quantitative results. Additional examples are provided in Section S.

B. ALL-MODALITIES COMPETITORS

This section compares MM SAM-adapter with those methods that are capable of processing more than two modalities simultaneously.

Table 7 shows the results on the DeLiVER validation and test sets. We employ the official weights of the competitors trained on all available modalities (RGB, LiDAR, Depth, and Event). We note that, even though trained on *only two modalities*, our MM SAM-adapter achieves *state-of-the-art* performance on the DeLiVER [7] test set regardless of the adopted auxiliary modality, and the best results also on the validation set when processing RGB and Depth. Table 8 reports results on the MUSES dataset. Again,

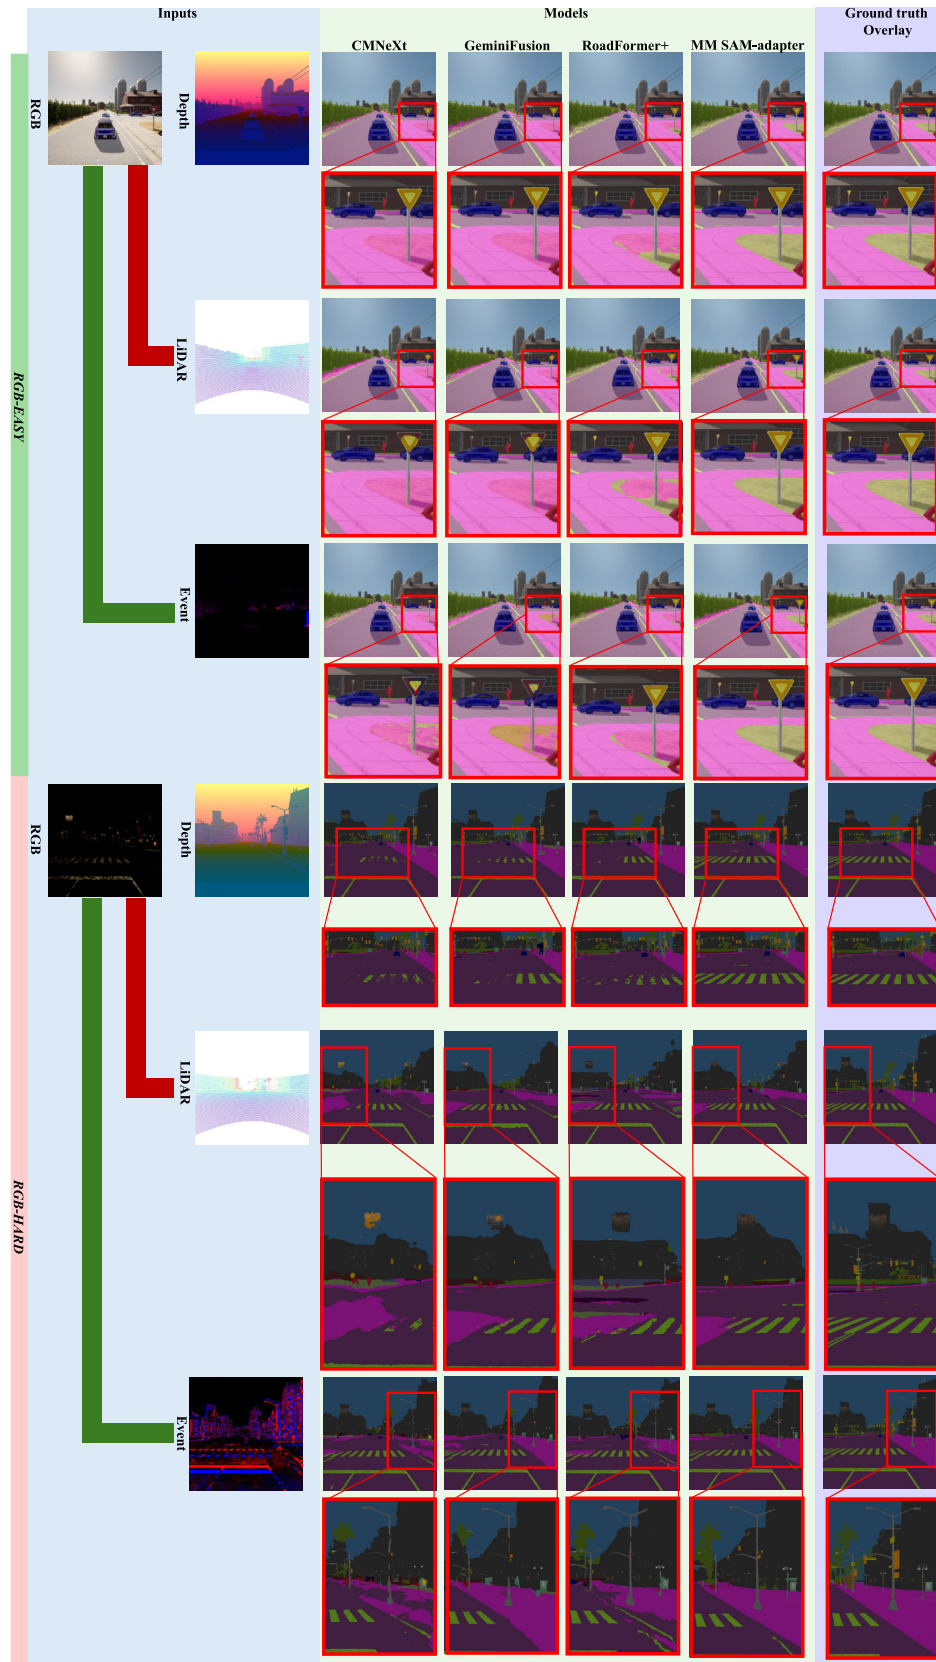


FIGURE 4. DeLiVER [7] test set predictions in RGB-Depth, RGB-LiDAR, RGB-Event framework. *RGB-easy* (top) and *RGB-hard* (bottom) samples. Notably, the input RGB image in *RGB-hard* case has been stretched with an exponential operator.

TABLE 9. Analysis of our contributions on the DeLiVER test set. SH = Segmentation Head, SPM = Spatial Prior Module. SAM* = SAM backbone frozen.

Method	Auxiliary Encoders	adapter Modalities	All mIoU	RGB-easy mIoU	RGB-hard mIoU
SAM + SH	None	None	53.32	54.28	35.57
SAM + SH + adapter	1 SPM	RGB	55.90	56.93	37.13
SAM* + SH + LoRA	1 ConvNext-S	RGB	53.18	54.28	35.05
SAM + SH + adapter	1 ConvNext-S	RGB	55.98	57.07	37.93
SAM* + SH + LoRA	1 ConvNext-S	LiDAR	49.77	50.81	32.59
SAM + SH + adapter	1 ConvNext-S	LiDAR	54.90	55.49	44.06
MM SAM*-LoRA	2 ConvNext-S	RGB & LiDAR	53.97	54.74	40.22
MM SAM-adapter	2 ConvNext-S	RGB & LiDAR	57.14	57.75	45.46

MM SAM-adapter achieves *state-of-the-art* performance independently of the adopted auxiliary modality, i.e. with both RGB + Event as well as RGB + LiDAR, highlighting the potential of our approach for broad applicability to real contexts. *These results are publicly available in the MUSES online benchmark.*¹

Therefore, the experiments in this section confirm that our approach—based on adapting the SAM backbone with *multimodal fused features*—is capable of effectively leveraging auxiliary knowledge while preserving the strong priors derived from the RGB-pretrained encoder.

C. ABLATION STUDIES

1) MAIN CONTRIBUTIONS

We conduct experiments on DeLiVER to highlight the impact of the key contributions in our proposed multimodal adapter architecture, reporting results in Table 9 and Figure 6.

The first row shows the results of the SAM image encoder fine-tuned for semantic segmentation by simply appending a SegFormer [18] Segmentation Head. The second and fourth rows, instead, report the results achieved by fine-tuning the SAM encoder via adapter-based approaches that utilize two different backbones: a Spatial Prior Module (SPM), as proposed in ViT-adapter [5], and a standard ConvNext-Small (ConvNext-S). In these two experiments, the adapter branch takes only RGB images as input, akin to the original ViT-adapter. We observe that the side-tuning SAM-adapter outperforms standard fine-tuning, likely because it better preserves SAM’s prior knowledge, reducing catastrophic forgetting. The third row depicts the result of SAM image encoder frozen, adapted using a different adaptation strategy based on LoRA [10]. To ensure a fair comparison with the SAM-adapter, we employ a ConvNext-S before the first LoRA layer as explained in Appendix V-A. When comparing the second, third, and fourth rows, we observe that the side-tuning SAM-adapter surpasses the LoRA adaptation, likely due to its ability to smoothly inject new domain-specific spatial information. However, RGB-only adapters still struggle in the *RGB-hard* scenario, with a performance gap of about 20% mIoU vs. *RGB-easy*. In the fifth row, SAM with LoRA adapter performs poorly when the LiDAR modality is introduced at the first layer and unexpectedly struggles even in the *RGB-hard* scenario

¹“MUSES online benchmark,” May 2025: <https://muses.vision.ee.ethz.ch/benchmarks#semanticSegmentation>

(from 50.81 in *RGB-easy* to 32.59 mIoU in *RGB-hard*). This is likely because a single LoRA layer is insufficient to integrate a modality like LiDAR into the RGB-based SAM backbone, given the inherent differences between RGB and LiDAR data. In the sixth row, we demonstrate that replacing the RGB input with the LiDAR modality in the side-tuning adapter branch significantly improves performance in the *RGB-hard* scenario (e.g., 44.06 vs. 37.93 mIoU for LiDAR vs. RGB adapter with the same ConvNext-S backbone). However, this comes at the cost of a performance drop in the *RGB-easy* scenario (from 57.07 to 55.49 mIoU), as noisier LiDAR measurements may corrupt highly informative RGB features. The seventh row shows the results of adapting using LoRA, yet by employing multimodal data as input. Multimodal LoRA utilizes our MM Fusion Encoder, which comprises two modality-specific encoders and a Fusion module. The fused encoder features are fed into the first LoRA layer. The results align with our main idea that we need to inject multimodal information to achieve effective adaptation. Indeed, multimodal LoRA obtains better performance than its single-modality counterpart in both *RGB-easy* and *RGB-hard* scenarios. However, we highlight that the best results are achieved by our proposed MM SAM-adapter strategy, as shown in the last row. Notably, our approach achieves the best performance across all scenarios. Figure 6 shows the predictions obtained by some of the methods in Table 9. Our MM SAM adapter yields way more accurate predictions than the RGB-only counterparts in the *RGB-hard* scenario. In the *RGB-easy* scenario, our model provides very similar segmentations compared to the RGB-only adapter, while performing much better than the model that uses the LiDAR-only auxiliary branch. These qualitative results vouch for our framework’s ability to learn which modality to prioritize during inference, namely to discard or deploy LiDAR features depending on whether the RGB content is sufficiently informative or not. This capability is achieved through the combination of the Multimodal Fusion Encoder with the adapter module, which smoothly contaminates SAM foundational knowledge with multimodal one.

2) ASYMMETRIC VERSUS SYMMETRIC ARCHITECTURE

To validate our asymmetric network design, we replace the two ConvNext-Small (ConvNext-S) encoders in the auxiliary branch with two ConvNext-Base (ConvNext-B) networks, so as to realize a symmetric architecture where the auxiliary branch (Multimodal fusion encoder + adapter) and the SAM encoder have a similar number of parameters. In Table 10, we present results on the DeLiVER test set for the RGB-LiDAR setup. Interestingly, the symmetric MM SAM-adapter performs worse, which underscores the importance of maintaining an asymmetric structure to prioritize SAM knowledge over the auxiliary modality one. Additionally, our asymmetric design enables to use lighter networks for the auxiliary branch, reducing computational complexity without sacrificing performance. This

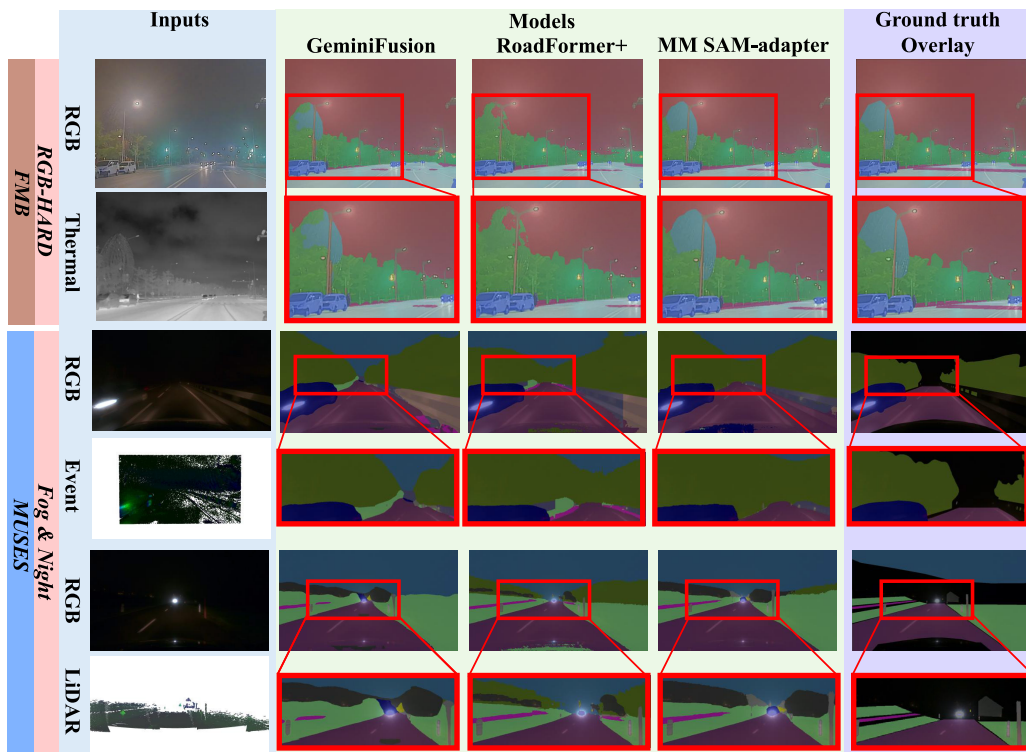


FIGURE 5. FMB test set and MUSES validation set predictions of the methods in Table 5 and in Table 6. Black pixels in the last column denote missing ground truth labels.

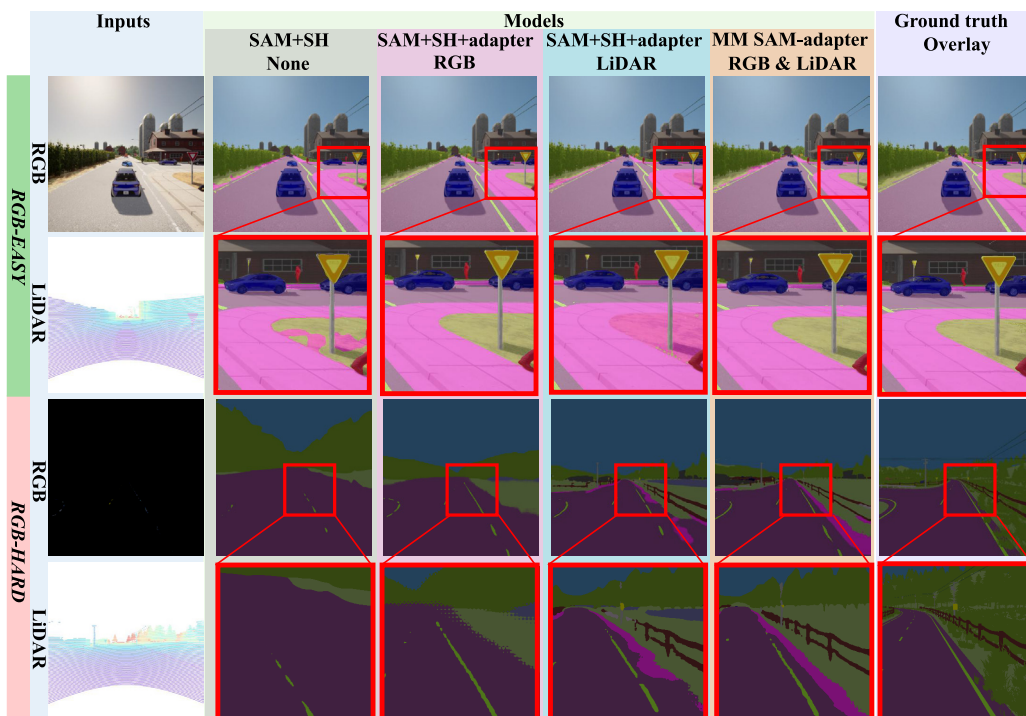


FIGURE 6. DeLIVER test set predictions of some of the methods in Table 9 in RGB-LiDAR framework. RGB-easy (top) and RGB-hard (bottom) samples.

experiment aligns with our intuition that RGB is typically the primary source of information, while auxiliary modalities

are needed mainly when RGB data is lacking. Therefore, adopting a symmetric architecture could undermine this

TABLE 10. Symmetric vs asymmetric architecture. Results on DeLiVER test set in the RGB-LiDAR setup.

Method	Architecture	All mIoU	RGB-easy mIoU	RGB-hard mIoU
MM SAM-adapter	Symmetric	55.27	55.93	42.22
MM SAM-adapter	Asymmetric	57.14	57.75	45.46

TABLE 11. Comparison of fusion modules. Results on the DeLiVER test in the RGB-LiDAR setup. Addition or Concatenation sums or concatenates the two modality features, respectively. Road-fusion is the RoadFormer+ [29] fusion module.

Method	Fusion Module	All mIoU	RGB-easy mIoU	RGB-hard mIoU
MM SAM-adapter	Addition	56.09	56.66	45.53
MM SAM-adapter	Concatenation	56.87	57.51	44.78
MM SAM-adapter	Road-Fusion	57.14	57.75	45.46

logic by assigning excessive emphasis to the auxiliary modality.

3) ALTERNATIVE FUSION MODULES

Our architecture is flexible, as it can incorporate different fusion modules. Even the simpler fusion strategies enable the model to achieve remarkable performance. In Table 11, we explore alternative architectures for the Fusion Module. In particular, we replace Road-Fusion [29] with simpler fusion techniques such as addition and concatenation. First of all, we point out that our intuition of adapting SAM using fused multimodal features proves effective with all the considered fusion techniques: even simple ones, like addition and concatenation, can yield *state-of-the-art* results on the RGB-LiDAR setup of DeLiVER. We also notice that in the case of addition-based fusion, the results for *RGB-hard* samples are comparable to those of Road-Fusion. In contrast, a drop in performance is observed in the *RGB-easy* case. This aligns with our expectations, as addition-based fusion makes it harder for the network to determine what information to use based on the scenario due to LiDAR noise being directly injected in SAM features. Conversely, with concatenation-based fusion, the results are close to those of Road-Fusion, especially in the *RGB-easy* scenario. Indeed, concatenating the features from the two modalities allows the adapter to easily sift out only the required information at test time. Nevertheless, the Road-Fusion strategy generates better fused features, enabling the adapter to inject and extract knowledge effectively, thereby achieving superior performance in both *RGB-easy* and *RGB-hard* scenarios.

4) MODALITY SPECIFIC VERSUS MODALITY AGNOSTIC ENCODERS

To assess whether using two modality-specific encoders in the Multimodal Fusion Encoder is more effective than a shared, modality-agnostic encoder, we carry out an ablation

TABLE 12. Single modality-agnostic vs two modality-specific encoders. Results on the DeLiVER test set in the RGB-LiDAR setup. MA = Modality-Agnostic. MS = Modality-Specific.

Method	Auxiliary Encoders	All mIoU	RGB-easy mIoU	RGB-hard mIoU
MM SAM-adapter	MA	56.21	56.81	44.79
MM SAM-adapter	MS	57.14	57.75	45.46

TABLE 13. Comparison of our method vs SAM-frozen version. Results on DeLiVER test set in the RGB-LiDAR setup.

Method	SAM frozen	All mIoU	RGB-easy mIoU	RGB-hard mIoU
MM SAM-adapter	✓	55.35	56.00	43.30
MM SAM-adapter	✗	57.14	57.75	45.46

TABLE 14. Comparison of our method vs scratch-trained ViT-L backbone version. Results on DeLiVER test set in the RGB-LiDAR setup. ST = scratch-training of ViT-L with no pretrained weights, FT = fine-tuning of SAM from pretrained weights.

Method	SAM training	All mIoU	RGB-easy mIoU	RGB-hard mIoU
MM SAM-adapter	ST	53.73	54.46	40.97
MM SAM-adapter	FT	57.14	57.75	45.46

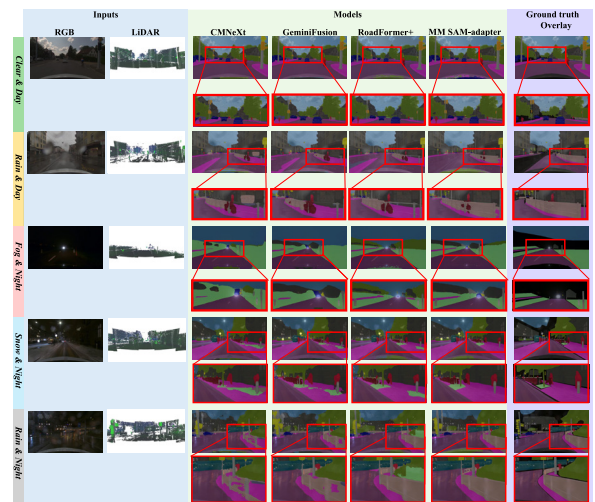


FIGURE 7. MUSES [9] validation set predictions in RGB-LiDAR framework. Different conditions have been illustrated.

study on the DeLiVER [7] test set in the RGB-LiDAR setup. The findings, shown in Table 12, indicate that employing two modality-specific encoders performs better than a single modality-agnostic encoder for the RGB-LiDAR modality pair. This can be attributed to the fundamental differences between RGB and LiDAR data: while RGB images provide dense visual information, LiDAR data is inherently sparse.

5) FROZEN VERSUS FINETUNED SAM

In Table 13, we compare our method with a variant where the SAM image backbone remains frozen during

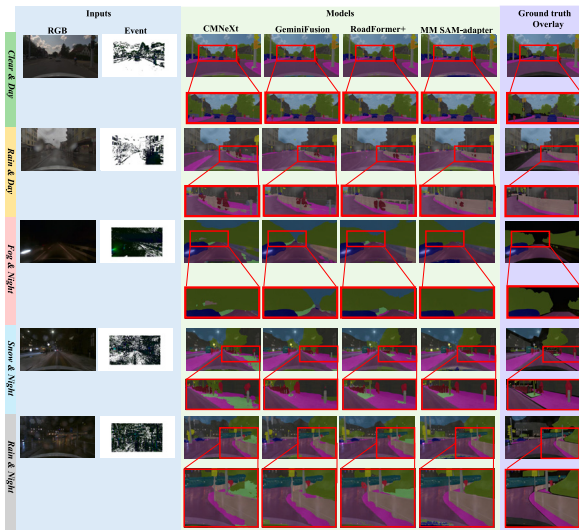


FIGURE 8. MUSES [9] validation set predictions in RGB-Event framework. Different conditions have been illustrate.




FIGURE 9. DeLiVER [7] test set predictions in RGB-LiDAR framework. RGB-easy (top) and RGB-hard (bottom) samples. Notably, the input RGB image in RGB-hard case has been stretched with an exponential operator.

training. As noted by ViT-adapter [5], this adaptation strategy benefits significantly from fine-tuning the SAM

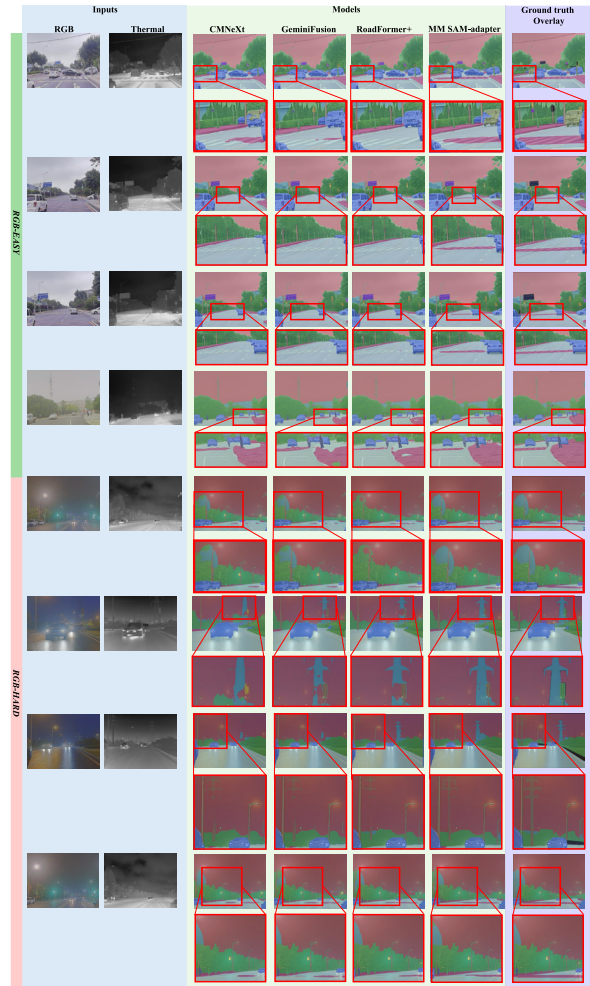


FIGURE 10. FMB [8] test set predictions. RGB-easy (top) and RGB-hard (bottom) samples.

backbone. Consequently, fine-tuning SAM achieves the best performance. Remarkably, the results obtained with the SAM-frozen backbone are nearly the same as CAFuser (i.e., 55.35 vs. 55.60 mIoU – see Table 7). This demonstrates that, even with the backbone entirely frozen, our multimodal adapter and fusion modules effectively guide the SAM backbone by integrating multimodal information in a beneficial way.

6) IMPORTANCE OF SAM PRE-TRAINED WEIGHTS

In Table 14, we present a reference baseline by evaluating our proposed architecture without using the pre-trained weights of the SAM image encoder. As expected, leveraging SAM’s pre-trained weights is essential for achieving *state-of-the-art* performance.

V. CONCLUSION

We propose MM SAM-adapter, a novel multimodal semantic segmentation framework that leverages an adapter strategy to harness SAM’s rich foundational knowledge. By enabling the adapter to utilize fused features from the RGB and one auxiliary modality, MM SAM-adapter ensures robust

performance in challenging conditions while maintaining optimal accuracy in simpler scenarios, resulting in state-of-the-art results across all the considered benchmarks.

One limitation of the MM SAM-adapter is that it currently supports only two input modalities due to constraints imposed by the road-fusion module. While our approach already outperforms existing all-modalities methods, an exciting avenue for future research is to extend this framework to accommodate more complex scenarios. This will require designing an innovative and efficient fusion module capable of effectively integrating more than two modalities. Additionally, exploring the potential of this framework in other tasks, such as panoptic segmentation, presents another promising avenue for further research.

APPENDIX A ADDITIONAL IMPLEMENTATION DETAILS A. LORA MULTIMODAL IMPLEMENTATION

The LoRA [10] multimodal network, showed in Table 9 has been implemented to assess the choice of the adapter method, presented in [5]. LoRA layers have been incorporated into every layer of the SAM Vision Transformer. However, unlike the standard LoRA implementation, we introduce modality-specific encoders and a fusion module at the first LoRA layer, to make it compatible with multimodal input. In the case of RGB-only data, a ConvNeXt network extracts features from the RGB image, which are then projected to a dimensionality compatible with SAM tokens. The standard LoRA layers subsequently process these projected features. For LiDAR input, the ConvNeXt network directly processes the LiDAR data, and the resulting features follow the same processing as in the RGB-only case. In the RGB-LiDAR scenario, the fusion module is identical to that used in the MM SAM-adapter, with two modality-specific encoders dedicated to processing RGB and LiDAR inputs, respectively.

B. COMPETITOR DETAILS

Each competitor has been trained and tested using the official repository provided by their authors. We used CMNeXt [7] GitHub repository, RoadFormer+ [29] GitHub repository, GeminiFusion [26] GitHub repository.

APPENDIX B ADDITIONAL QUALITATIVE RESULTS

We present qualitative results of our proposed method and competitors in Figure 7, Figure 8, Figure 9, Figure 10.

REFERENCES

- [1] H. Thisanke, C. Deshan, K. Chamith, S. Seneviratne, R. Vidanaarachchi, and D. Herath, "Semantic segmentation using vision transformers: A survey," *Eng. Appl. Artif. Intell.*, vol. 126, Nov. 2023, Art. no. 106669.
- [2] Y. Zhang, D. Sidibé, O. Morel, and F. Mériaudeau, "Deep multimodal fusion for semantic image segmentation: A survey," *Image Vis. Comput.*, vol. 105, Jan. 2021, Art. no. 104042.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [4] A. M. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W. Lo, P. Dollár, and R. Girshick, "Segment anything," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Jun. 2023.
- [5] Z. Chen, Y. Duan, W. Wang, J. He, T. Lu, J. Dai, and Y. Qiao, "Vision transformer adapter for dense predictions," in *Proc. 11th Int. Conf. Learn. Represent.*, 2022.
- [6] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16×16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent.*, 2020. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [7] J. Zhang, R. Liu, H. Shi, K. Yang, S. Reiß, K. Peng, H. Fu, K. Wang, and R. Stiefelhofen, "Delivering arbitrary-modal semantic segmentation," in *Proc. CVPR*, 2023.
- [8] J. Liu, Z. Liu, G. Wu, L. Ma, R. Liu, W. Zhong, Z. Luo, and X. Fan, "Multi-interactive feature learning and a full-time multi-modality benchmark for image fusion and segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 8081–8090.
- [9] T. Broedermann, D. Brüggemann, C. Sakaridis, K. Ta, O. Liagouris, J. Corkill, and L. V. Gool, "MUSES: The multi-sensor semantic perception dataset for driving under uncertainty," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2024, pp. 21–38.
- [10] J. E. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. ICLR*, 2021, p. 3.
- [11] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 7, pp. 3523–3542, Jul. 2022.
- [12] A. M. Hafiz and G. M. Bhat, "A survey on instance segmentation: State of the art," *Int. J. Multimedia Inf. Retr.*, vol. 9, no. 3, pp. 171–189, Sep. 2020.
- [13] H. Khan, M. T. Usman, I. Rida, and J. Koo, "Attention enhanced machine instinctive vision with human-inspired saliency detection," *Image Vis. Comput.*, vol. 152, Dec. 2024, Art. no. 105308. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S026288562400413X>
- [14] J. Chen, H. Zhang, M. Gong, and Z. Gao, "Collaborative compensative transformer network for salient object detection," *Pattern Recognit.*, vol. 154, Oct. 2024, Art. no. 110600. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0031320324003510>
- [15] B. Cheng, A. G. Schwing, and A. Kirillov, "Per-pixel classification is not all you need for semantic segmentation," in *Proc. NeurIPS*, Apr. 2021.
- [16] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 1280–1289.
- [17] J. Jain, J. Li, M. Chiu, A. Hassani, N. Orlov, and H. Shi, "OneFormer: One transformer to rule universal image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 2989–2998.
- [18] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Álvarez, and P. Luo, "SegFormer: Simple and efficient design for semantic segmentation with transformers," in *Proc. Neural Inf. Process. Syst. (NeurIPS)*, 2021.
- [19] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. K. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer, "SAM 2: Segment anything in images and videos," in *Proc. 13th Int. Conf. Learn. Represent.*, 2024. [Online]. Available: <https://openreview.net/forum?id=Ha6RTeWmd0>
- [20] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Gao, and Y. J. Lee, "Segment everything everywhere all at once," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023.
- [21] C. Couprie, C. Farabet, L. Najman, and Y. LeCun, "Indoor semantic segmentation using depth information," 2013, *arXiv:1301.3572*.
- [22] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "FuseNet: Incorporating depth into semantic segmentation via fusion-based CNN architecture," in *Proc. Comput. Vis. (ACCV)*, 2017, pp. 213–228.
- [23] Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, "MFNet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2017, pp. 5108–5115.
- [24] Y. Sun, W. Zuo, and M. Liu, "RTFNet: RGB-thermal fusion network for semantic segmentation of urban scenes," *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2576–2583, Jul. 2019.

- [25] J. Zhang, H. Liu, K. Yang, X. Hu, R. Liu, and R. Stiefelhagen, "CMX: Cross-modal fusion for RGB-X semantic segmentation with transformers," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 12, pp. 14679–14694, Dec. 2023.
- [26] D. Jia, J. Guo, K. Han, H. Wu, C. Zhang, C. Xu, and X. Chen, "GeminiFusion: Efficient pixel-wise multimodal fusion for vision transformer," in *Proc. 41st Int. Conf. Mach. Learn.*, vol. 235, R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berenkamp, Eds., pp. 21753–21767. [Online]. Available: <https://proceedings.mlr.press/v235/jia24b.html>
- [27] Y. Wang, X. Chen, L. Cao, W. Huang, F. Sun, and Y. Wang, "Multimodal token fusion for vision transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 12176–12185.
- [28] J. Li, Y. Zhang, P. Yun, G. Zhou, Q. Chen, and R. Fan, "RoadFormer: Duplex transformer for RGB-normal semantic road scene parsing," *IEEE Trans. Intell. Vehicles*, vol. 9, no. 7, pp. 5163–5172, Jul. 2024, doi: [10.1109/TIV.2024.3388726](https://doi.org/10.1109/TIV.2024.3388726).
- [29] J. Huang, J. Li, N. Jia, Y. Sun, C. Liu, Q. Chen, and R. Fan, "RoadFormer+: Delivering RGB-X scene parsing through scale-aware information decoupling and advanced heterogeneous feature fusion," *IEEE Trans. Intell. Vehicles*, vol. 10, no. 5, pp. 3156–3165, May 2025, doi: [10.1109/TIV.2024.3448251](https://doi.org/10.1109/TIV.2024.3448251).
- [30] T. Broedermann, C. Sakaridis, Y. Fu, and L. V. Gool, "CAFuser: Condition-aware multimodal fusion for robust semantic perception of driving scenes," *IEEE Robot. Autom. Lett.*, vol. 10, no. 4, pp. 3134–3141, Apr. 2025.
- [31] B. Li, D. Zhang, Z. Zhao, J. Gao, and X. Li, "StitchFusion: Weaving any visual modalities to enhance multimodal semantic segmentation," 2024, [arXiv:2408.01343](https://arxiv.org/abs/2408.01343).
- [32] S. Dong, Y. Feng, Q. Yang, Y. Huang, D. Liu, and H. Fan, "Efficient multimodal semantic segmentation via dual-prompt learning," 2023, [arXiv:2312.00360](https://arxiv.org/abs/2312.00360).
- [33] K. Wang, D. Lin, C. Li, Z. Tu, and B. Luo, "Adapting segment anything model to multi-modal salient object detection with semantic feature fusion guidance," 2024, [arXiv:2408.15063](https://arxiv.org/abs/2408.15063).
- [34] A. Xiao, W. Xuan, H. Qi, Y. Xing, N. Yokoya, and S. Lu, "Segment anything with multiple modalities," 2024, [arXiv:2408.09085](https://arxiv.org/abs/2408.09085).
- [35] Z. Chen, Z. Zhu, Y. Zhang, J. Hou, G. Shi, and J. Wu, "Segment any event streams via weighted adaptation of pivotal tokens," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 3890–3900.
- [36] B. Yao, Y. Deng, Y. Liu, H. Chen, Y. Li, and Z. Yang, "SAM-event-adapter: Adapting segment anything model for event-RGB semantic segmentation," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2024, pp. 9093–9100.
- [37] P. Liu, J. Deng, L. Duan, W. Li, and F. Lv, "Segmenting anything in the dark via depth perception," *IEEE Trans. Multimedia*, vol. 27, pp. 2975–2986, 2025.
- [38] N. Hounsby, A. Giurgiu, S. Jastrzębski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, "Parameter-efficient transfer learning for NLP," in *Proc. 36th Int. Conf. Mach. Learn.*, Jun. 2019, pp. 2790–2799. [Online]. Available: <https://proceedings.mlr.press/v97/hounsby19a.html>
- [39] Y. Xin, J. Yang, S. Luo, Y. Du, Q. Qin, K. Cen, Y. He, B. Fu, X. Yang, G. Zhai, M.-H. Yang, and X. Liu, "Parameter-efficient finetuning for pre-trained vision models: A survey and benchmark," 2024, [arXiv:2402.02242](https://arxiv.org/abs/2402.02242).
- [40] T. Chen, L. Zhu, C. Ding, R. Cao, Y. Wang, S. Zhang, Z. Li, L. Sun, Y. Zang, and P. Mao, "SAM-adapter: Adapting segment anything in underperformed scenes," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2023, pp. 3367–3375.
- [41] Y. Song, Q. Zhou, X. Lu, Z. Shao, and L. Ma, "SU-SAM: A simple unified framework for adapting segment anything model in underperformed scenes," 2024, [arXiv:2401.17803](https://arxiv.org/abs/2401.17803).
- [42] T. Chen, A. Lu, L. Zhu, C. Ding, C. Yu, D. Ji, Z. Li, L. Sun, P. Mao, and Y. Zang, "SAM2-adapter: Evaluating & adapting segment anything 2 in downstream tasks: Camouflage, shadow, medical image segmentation, and more," 2024, [arXiv:2408.04579](https://arxiv.org/abs/2408.04579).
- [43] J. Wu, W. Ji, Y. Liu, H. Fu, M. Xu, Y. Xu, and Y. Jin, "Medical SAM adapter: Adapting segment anything model for medical image segmentation," 2023, [arXiv:2304.12620](https://arxiv.org/abs/2304.12620).
- [44] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 15979–15988.
- [45] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 11966–11976.
- [46] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable DETR: Deformable transformers for end-to-end object detection," in *Proc. Int. Conf. Learn. Represent.*, 2021. [Online]. Available: <https://openreview.net/forum?id=gZ9hCDWe6ke>
- [47] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, [arXiv:1607.06450](https://arxiv.org/abs/1607.06450).
- [48] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223. [Online]. Available: <https://api.semanticscholar.org/CorpusID:502946>
- [49] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, "Deconvolutional networks," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2010, pp. 2528–2535, doi: [10.1109/CVPR.2010.5539957](https://doi.org/10.1109/CVPR.2010.5539957).
- [50] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. Int. Conf. Learn. Represent.*, 2017. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>
- [51] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9992–10002.



IACOPO CURTI received the master's degree in automation engineering from the University of Bologna, in 2023, where he is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering. His research interests include image processing and deep learning frameworks, in particular multimodal architectures for dense prediction tasks.



PIERLUIGI ZAMA RAMIREZ received the Ph.D. degree in computer science and engineering, in 2021. He is currently an Assistant Professor at the University of Bologna. He co-authored more than 30 publications on computer vision research topics, such as semantic segmentation, depth estimation, anomaly detection, domain adaptation, neural fields, and 3D computer vision.



ALIOSCIA PETRELLI received the master's degree in computer science engineering and the Ph.D. degree in computer science from the University of Bologna, in 2005 and 2016, respectively. He spent four years as a Research Fellow at the Computer Vision Laboratory (CVLab) of the Department of Computer Science and Engineering (DISI), Bologna. His research interests focus on computer vision, including 3D surface matching, visual search, and machine learning.



LUIGI DI STEFANO (Member, IEEE) received the Ph.D. degree in electronic engineering and computer science from the University of Bologna, in 1994. He is currently a Full Professor at the Department of Computer Science and Engineering, University of Bologna, where he founded and leads the Computer Vision Laboratory (CVLab). He has been a Scientific Consultant for major computer vision and machine learning companies. He is the author of more than 150 articles and several patents. His research interests include image processing, computer vision, and machine/deep learning. He is a member of the IEEE Computer Society and IAPR-IC.