Learning Modality-agnostic Representation for Semantic Segmentation from Any Modalities

Xu Zheng¹[®], Yuanhuiyi Lyu¹[®], and Lin Wang^{1,2}[®]*

 ¹ Hong Kong University of Science and Technology, Guangzhou, China zhengxu128@gmail.com, yuanhuiyilv@hkust-gz.edu.cn
 ² Hong Kong University of Science and Technology, Hong Kong, China linwang@ust.hk

https://vlislab22.github.io/Any2Seg/

Abstract. Image modality is not perfect as it often fails in certain conditions, e.g., night and fast motion. This significantly limits the robustness and versatility of existing multi-modal (*i.e.*, Image+X) semantic segmentation methods when confronting modality absence or failure, as often occurred in real-world applications. Inspired by the open-world learning capability of multi-modal vision-language models (MVLMs), we explore a new direction in learning the modality-agnostic representation via knowledge distillation (KD) from MVLMs. Intuitively, we propose Any2Seg, a novel framework that can achieve robust segmentation from any combination of modalities in any visual conditions. Specifically, we first introduce a novel language-guided semantic correlation distillation (LSCD) module to transfer both inter-modal and intra-modal semantic knowledge in the embedding space from MVLMs, e.g., LanguageBind [82]. This enables us to minimize the modality gap and alleviate semantic ambiguity to combine any modalities in any visual conditions. Then, we introduce a modality-agnostic feature fusion (MFF) module that reweights the multi-modal features based on the inter-modal correlation and selects the fine-grained feature. This way, our Any2Seg finally yields an optimal modality-agnostic representation. Extensive experiments on two benchmarks with four modalities demonstrate that Any2Seg achieves the stateof-the-art under the multi-modal setting (+3.54 mIoU) and excels in the challenging modality-incomplete setting (+19.79 mIoU).

Keywords: Multi-modal Segmentation, MVLMs, KD

1 Introduction

The integration of multi-modal sensory input is pivotal for enabling intelligent agents to achieve robust and reliable scene understanding, e.g., semantic segmentation [1, 7, 8, 16, 20, 32, 34, 41, 44, 49, 66, 71–73, 75–77, 84]. As such, distinct visual modalities have been introduced into multi-modal systems, such as

^{*} Corresponding author



Fig. 1: (a) Our Any2Seg aims to learn a modality-agnostic representation for robust segmentation from any modalities guided by multi-modal VLMs; (b) Performance comparison between CMNeXt [60] and our Any2Seg under the multi-modal semantic segmentation (MSS) and modality-incomplete semantic segmentation (MISS) settings.

depth [5, 9, 47, 48, 78], LiDAR [46, 54, 86], and event [1, 61], to accompany RGB data to achieve robust scene understanding. Recently, endeavors have been made in scaling from dual modality to multiple modality fusion which is capable of processing a broader spectrum of input modalities [58, 60, 70, 74, 79, 83], *a.k.a.*, multi-modal fusion frameworks.

However, these methods predominantly rely on image modality, which is not always reliable as it often fails in certain conditions, *e.g.*, extreme lighting. This significantly limits the robustness and versatility of multi-modal (*i.e.*, RGB+X modality) frameworks [4, 27, 38, 58, 60] when confronting modality absence and failure, as often occurred in the wild. Humans can automatically combine sensory modalities when understanding the environment. Intuitively, for machines to emulate humans, the multi-modal segmentation model should also be able to dynamically combine various modalities across different environments. This is particularly beneficial for overcoming modality absence and failure challenges, *a.k.a.*, **system-level** and **sensor-level modality-incomplete** conditions [27].

The recent multi-modal vision-language models (MVLMs) have gained significant attention for their strong open-world learning capability and superiority in alleviating cross-modal heterogeneity by binding multiple visual modalities to a central modality (*e.g.*, language) in the representation space [14, 15, 28, 29, 63, 64, 80, 82]. Inspired by their success, we explore a **new** direction by learning modality-agnostic representation to tackle aforementioned challenges under the guidance of MVLMs. Intuitively, we present **Any2Seg**, a novel multi-modal segmentation framework that can achieve robust semantic segmentation from any combination of modalities in any visual conditions, as depicted in Fig. 1 (a).

To make this possible, we first propose a language-guided semantic correlation distillation (LSCD) module to minimize the modality gap and further alleviate semantic ambiguity to achieve the capability of combining any modalities in any visual conditions (Sec. 3.2). Specifically, to minimize the modality gap, LSCD first distills the inter-modal correlation with text descriptions of the targeted scene generated by LLaMA-adapter [64]. Then, to alleviate semantic ambiguity, LSCD extracts the intra-modal semantic knowledge from LanguageBind and transfers it to the segmentation model.

With the transferred inter- and intra-modal semantic knowledge, we then introduce a novel Modality-agnostic Feature Fusion (**MFF**) module to **learn an optimal modality-agnostic representation** from the multi-modal features (Sec. 3.3). Concretely, the inter-modal semantic correlation from the LSCD is utilized to reweight the multi-modal features. Next, we select the fine-grained feature across multi-modal features for better fusion from any combination of modalities. Finally, MFF aggregates the reweighted features and the fine-grained feature to learn the optimal modality-agnostic representation.

We conduct extensive experiments across dual, triple, and quadruple modality combinations on DLIVER [60] and MCubeS [24] datasets with two experimental settings: 1) multi-modal segmentation with all modalities as inputs, and 2) modality-incomplete segmentation which accepts any combination of modalities in any visual conditions during the inference. Experimental results demonstrate that Any2Seg significantly surpasses prior SoTA methods [3, 27, 59, 60], under the multi-modal segmentation setting (See Fig. 1 (b)). When extending to challenging modality-incomplete settings that accommodate inputs from four modalities, our Any2Seg significantly outperforms existing approaches, achieving a remarkable enhancement of +19.79 mIoU.

In summary, the main contributions of our paper are as follows: (I) We make the *first* attempt to explore the potential of MVLMs and propose Any2Seg that can achieve robust segmentation from any combination of modalities in any visual conditions by distilling knowledge from MVLMs; (II) We propose a novel LSCD module to achieve the capability of combining any modalities in any visual conditions; (III) We introduce a novel MFF module to learn the optimal modality-agnostic representation; (IV) We conduct extensive experiments on two multi-modal benchmarks with four modalities in two settings, and the results demonstrate the superiority and robustness of our Any2Seg.

2 Related Work

Multi-modal Semantic Segmentation Multi-modal semantic segmentation aims at improving the performance by introducing more visual modalities that consists of complementary scene information. Endeavors have been made in introducing distinct visual modalities, such as depth [5,9,12,19,21,38,42,47,48, 55,78], LiDAR [2,22,23,26,46,54,57,86], event [1,61], and thermal data [6,18,25, 33,37,51,53,62,65,81], to accompany the RGB data to achieve robust scene understanding. Recently, with the development of novel sensors, such as the event cameras [70], various approaches have been proposed to scale from dual modality fusion to multiple modality fusion for achieving robust scene understanding at all day time [3,9,31,46,50,58,60,62]. CMNeXt [60] is a representative work that

achieves multi-modal semantic segmentation with arbitrary-modal complements by taking RGB as fundamentals and other modalities as auxiliary inputs.

The most significant challenge in multi-modal semantic segmentation is how to alleviate the cross-modal heterogeneity while efficiently extract comprehensive scene information. Existing works always design tailored components to overcome the modality gaps, such as the self-query hub in CMNeXt [60]. However, this kind of component only gives marginal performance gain and shows unsatisfactory performance when combining more modalities. Differently, inspired by the success of MVLMs, we pioneeringly unleash the potential of MVLMs and take them as teacher to minimize the modality gap and reduce semantic ambiguity for better multi-modal segmentation.

Multi-modal Vision-language Models (MVLMs). The integration of diverse modalities into a unified embedding space for cross-modal learning has emerged as a burgeoning research area. Initial efforts, exemplified by CLIP [35], have employed contrastive learning paradigms to align image and text pairs, aiming to achieve promising zero-shot generalization performance. Recent advancements have further expanded the scope of learning large-scale vision-language models to encompass multiple modalities, encompassing video [11, 40, 69], point cloud [17, 46, 63, 85], thermal [13, 59], event [70, 79], and other modalities.

ImageBind [13] leverages the binding property of image modality to bind the multi-modal embeddings with the image modality. Subsequent works [15,39,67], such as Point-Bind [14], extend ImageBind [15] to integrate point cloud data, resulting in remarkable zero-shot 3D capabilities. More recently, LanguageBind [82] employs language modality as the binding center across various modalities, aiming to address multi-modal tasks. In this paper, we make the **first** attempt to explore the potential of MVLMs and propose Any2Seg that can achieve robust segmentation from any combination of modalities in any visual conditions by distilling knowledge from MVLMs.

Modality-Incomplete Semantic Segmentation Research endeavors have been made in investigating robust frameworks that accommodate complete modality presence during training, despite the potential random absence of modalities during validation [10, 27, 30, 36, 45, 49, 68]. To address the challenge of missing modalities in multi-modal tasks, ShaSpec [49] has been designed to exploit all available inputs by learning shared and specific features, facilitated by a shared encoder alongside individual modality-specific encoders. Wang *et al.* [43] introduced a method to adaptively discern crucial modalities and distill knowledge from them, aiding in the cross-modal compensation for absent modalities.

More recently, Liu *et al.* [27] have broadened the scope by establishing the concept of modality-incomplete scene segmentation, addressing both systemlevel and sensor-level modality deficiencies. They introduced a missing-aware modal switch strategy that mitigates over-reliance on predominant modalities in multi-modal fusion and proactively regulates missing modalities during the training phase. Different from the ideas to achieve cross-modal compensation and design missing-aware strategies, we directly focus on learning modality-agnostic representation. Our Any2Seg prioritizes knowledge distillation from



Fig. 2: The overall framework of Any2Seg, incorporating the LSCD and MFF modules to learn an optimal modality-agnostic representation for robust segmentation.

MVLMs and fine-grained feature fusion, thereby ensuring robustness and efficiency in both multi-modal and modality-incomplete segmentation tasks.

3 The Proposed Any2Seg Framework

3.1 Overview

An overview of our Any2Seg is depicted in Fig. 2. It consists of two key modules: the Language-guided Semantic Correlation Distillation (**LSCD**) (Sec 3.2) and the Modality-agnostic Feature Fusion (**MFF**) (Sec 3.3) modules. The LSCD module is only for training, and the MFF module is utilized in both training and inference to aggregate multi-modal feature maps $\{f_r, f_d, f_e, f_l\}$ into modalityagnostic feature f_{ma} , as shown in the forward pass of Fig. 2.

Inputs: Our Any2Seg processes multi-modal visual data from four modalities, all within the same scene. We consider RGB images $\mathbf{R} \in \mathbb{R}^{h \times w \times 3}$, depth maps $\mathbf{D} \in \mathbb{D}^{h \times w \times C^{D}}$, LiDAR data $\mathbf{L} \in \mathbb{L}^{h \times w \times C^{P}}$, and event stack images $\mathbf{E} \in \mathbb{E}^{h \times w \times C^{E}}$ to illustrate our method, as depicted in Fig. 2. Here, we follow the data processing as [60], where the channel dimensions $C^{D} = C^{P} = C^{E} = 3$. Any2Seg also integrates the corresponding ground truth Y across K categories. For each training iteration, Any2Seg takes a mini-batch $\{r, d, e, l\}$ containing samples from all the input modalities.

MVLM and Backbone: Within Any2Seg, we integrate the cutting-edge LanguageBind [82], functioning as a 'teacher' for knowledge distillation. We incorporate the multi-modal encoders of [82] to extract corresponding embeddings

for all input modalities, as shown in the visual and the text encoders in Fig. 2. Concurrently, we employ SegFormer [52] as the segmentation backbone.

Outputs: During the inference, Any2Seg utilizes only the forward pass, as depicted in Fig. 2. Given a multi-modal data mini-batch $\{r, d, e, l\}$, the data is first fed into the encoder of the segmentation backbone. The high-level feature maps $\{f_r, f_d, f_l, f_e\}$ are then passed to our MFF module for extracting the modality-agnostic feature map f_{ma} . Finally, the segmentation head utilizes f_{ma} to generate the predictions. We now describe the two proposed modules in detail.

3.2 Language-guided Semantic Correlation Distillation (LCSD)

Building on the recent MVLMs have the strong open-world learning capability and superiority in alleviating cross-modal heterogeneity, we advocate for the knowledge distillation from these cutting-edge MVLM, *i.e.*, LanguageBind [82]. It takes language modality as the central modality to bind all the visual modalities at the high-level representation space.

Scene Description Generation. To obtain the text description of the targeted scene, we use the LLaMA-Adapter [64] to generate a detailed scene description with visual data and category names. This model adeptly crafts linguistic representation of scenes, exemplified by ['a road that goes along...'], from visual input batch $\{r, d, e, l\}$ and their corresponding category names, such as ['Road']. Following this, the category names, along with the generated scene description and the visual input batch $\{r, d, e, l\}$, are fed into the text and visual encoders of LanguageBind [82], respectively. This results in a series of high-level feature embeddings, including semantic embeddings E_C from category names, multiple visual embeddings $\{E_R, E_D, E_E, E_L\}$, and scene embedding E_S .³

Inter-modal Correlation Distillation: To minimize the modality gap, our Any2Seg employs the visual embeddings $\{E_r, E_d, E_e, E_l\}$ and scene embedding E_S to compute inter-modal correlation, as illustrated in Fig. 2. As discussed in LanguageBind [82], linguistic concepts possess an inherent modality-agnostic nature and are less susceptible to style variations in the visual modalities. Consequently, we utilize the scene embedding E_s as a central reference for calculating embedding distances relative to $\{E_r, E_d, E_e, E_l\}$, leading to the formation of the correlation matrix M_{corr}^e as follows:

$$M_{corr}^e = Concat(Cos(E_r, E_S), Cos(E_d, E_S), Cos(E_e, E_S), Cos(E_l, E_S)).$$
(1)

where *Cos* denotes the computation of vector similarities utilizing cosine similarity metrics. Furthermore, as shown in the *'Inter-modal Correlation'* of Fig. 2, the obtained modality-agnostic feature f_{ma} is employed as an anchor point for calculating the correlation with the multi-modal feature maps $\{f_r, f_d, f_e, f_l\}$:

 $M_{corr}^{f} = Concat(Cos(f_r, f_{ma}), Cos(f_d, f_{ma}), Cos(f_e, f_{ma}), Cos(f_l, f_{ma})).$ (2)

The loss function for mirroring similarity matrices to transfer the inter-modal correlation knowledge:

$$\mathcal{L}_{cr} = KL(M_{corr}^f, M_{corr}^e), \tag{3}$$

³ Exact operations can be found in the supplementary material.



Fig. 3: We use the Modality-agnostic Feature Fusion Module to obtain modalityagnostic features for multi-modal segmentation.

where KL denotes the Kullback-Leibler Divergence. After minimizing modality gaps among visual modalities, we strive to reduce the semantic ambiguity by transferring the intra-modal semantic knowledge.

Intra-modal Semantic Distillation: We employ semantic embeddings E_c which are extracted by the text encoder of LanguageBind [82] from category names, such as ['Road'], as supervision signals for segmentation model, as shown in Fig. 2. Subsequently, we use the ground truth $y \in Y$ and the modality-agnostic feature to compute high-level class-wise representation, as shown in the *Class-wise Rep* in Fig. 2, with Mask Average Pooling (MAP):

$$\{e_0, \dots, e_K\} = MAP(f_{ma}, y). \tag{4}$$

The detailed MAP process is as follows: 1) we first up-sample the modalityagnostic feature f_{ma} to fit the size of ground truth y; 2) we then use y as classwise masks to derive per-class features; 3) lastly, we perform average pooling on these per-class features to get the final class-wise semantic representation $\{e_0, \ldots, e_K\}$. Building upon this, LSCD initially transfers intra-modal semantic knowledge with implicit relationship transfer. Specifically, we emulate the selfsimilarity matrix of E_c and $\{e_0, \ldots, e_K\}$ to achieve the intra-modal semantic knowledge distillation:

$$\mathcal{L}_{se} = KL(Cos(\{e_0, \dots, e_K\}, \{e_0, \dots, e_K\}^T), Cos(E_C, E_C^T)).$$
(5)

The overall knowledge distillation loss of LSCD can be formulated as:

$$\mathcal{L}_{kd} = \mathcal{L}_{cr} + \mathcal{L}_{se}.$$
 (6)

3.3 Modality-agnostic Feature Fusion (MFF) Module

Building upon LSCD module, we then present the MFF module, a tailored architecture designed to *learn an optimal modality-agnostic representation across modalities*, thus improving the robustness of multi-modal semantic segmentation in any visual conditions. The input of MFF module is the

modality-agnostic feature f_{ma}^{t-1} from the iteration at t-1, and the visual features $\{f_r, f_d, f_e, f_l\}$. Note that the f_{ma}^0 is initialized by the average of multi-modal features $\{f_r, f_d, f_e, f_l\}$. As shown in Fig. 3, we first compute the cross-modal correlation M_{corr}^e by calculating the cosine similarity between the f_{ma}^{t-1} and $\{f_r, f_d, f_e, f_l\}$. Then M_{corr}^e is multiplied (\otimes) with visual features $\{f_r, f_d, f_e, f_l\}$ to obtain the weighted feature maps. Then, the weighted feature maps are averaged to obtain the modality-balanced feature f_{mb} .

Afterwards, MFF module computes the similarity between $\{f_r, f_d, f_e, f_l\}$ and f_{mb} to obtain cross-modal similarity maps C_S . After that, we perform max similarity index selection for each element location across four similarity maps (denoted as gray maps in Fig. 3) to get the max index map M_I . MFF module uses the max index map M_I to select elements from $\{f_r, f_d, f_e, f_l\}$ to obtain the modality-selected feature f_{ms} . Finally, MFF module aggregates (\oplus) modality-balanced feature f_{mb} and modality-selected feature f_{ms} to get the modality-agnostic feature f_{ma}^t for iteration at t, which is subsequently fed into the segmentation head, resulting in the prediction P_{ma} , from which the supervised loss is computed as:

$$\mathcal{L}_{sup} = \operatorname{CE}(P_{ma}, Y),\tag{7}$$

where CE represents the Cross-Entropy loss function. The modality-agnostic feature f_{ma}^t for iteration at t is then used as input for iteration at t + 1.

3.4 Optimization, Training, and Validation

The total training objective for learning the multi-modal semantic segmentation model containing three losses (Eq.3, Eq.5, Eq.7) is defined as:

$$\mathcal{L} = \mathcal{L}_{se} + \mathcal{L}_{cr} + \mathcal{L}_{sup}.$$
(8)

Different from most multi-modal semantic segmentation frameworks [58,60] process multi-modal inputs with parallel modality-specific encoders, our method takes multi-modal data batch as input and simultaneously processes them with shared-weights encoders across modalities, the loss functions for each modality are integrated together before backward. During training, the multi-modal data flow is depicted in Fig. 2, and both the LSCD and MFF modules are included. For the validation phase with multi-modal data, Any2Seg performs solely the segmentation forward pass with MFF module.

Multi-modal Semantic Segmentation (**MSS**) validation means that we pass all the modalities in training for validation. **System-level** Modality-incomplete Semantic Segmentation (**MISS**) validation denotes that we evaluate all the possible combinations of input modalities and average all results to obtain final mean performance. **Sensor-level MISS** means that we evaluate the segmentation model with specific condition data from all modalities, such as cloudy weather.

4 Experiments

Datasets. DELIVER [60] is a comprehensive multi-modal dataset featuring RGB, depth, LiDAR, and event data across 25 semantic categories, under var-

Method	Modal	Backbone	mIoU	Method	Modal	Backbone	mIoU
HRFuser [56]		HRFuser-T	47.95	HRFuser [56]		HRFuser-T	43.13
SegFormer [52]	R	MiT-B0	52.10	TokenFusion [46]	B+L		53.01
SegFormer [52]		MiT-B2	57.20	CMX [58]	n+L	MiT-B2	56.37
HRFuser [56]		HRFuser-T	51.88	CMNeXt [60]			58.04
FPT [27]		MultiMAE	57.38	Ours			57.92
TokenFusion [46]	$ ^{R+D}$		60.25	HRFuser [56]		HRFuser-T	51.83
CMX [58]		MiT-B2	62.67	CMNeXt [60]	R+D+E	MiT-B2	64.44
CMNeXt [60]			63.58	Ours			66.75 + 2.31
Ours			67.12 + 3.54	HRFuser [56]		HRFuser-T	52.72
HRFuser [56]		HRFuser-T	42.22	CMNeXt [60]	R+D+L	MiT-B2	65.50
TokenFusion [46]	B+E		45.63	Ours			67.20 + 1.70
CMX [58]		MiT-B2	56.62	HRFuser [56]		HRFuser-T	52.97
CMNeXt [60]			57.48	CMNeXt [60]	$ ^{R+D+E+L}$	MiT-B2	66.30
Ours			57.83 +0.35	Ours			68.25 + 1.95

Table 1: Quantitative results in MSS setting across dual, triple, and quadruple modalities on DELIVER. R: RGB, D: Depth, E: Event, L: LiDAR.

ious environmental conditions and sensor failures. We follow the official processing and split guidelines from [60]. MCubeS [24] focuses on material segmentation, offering 20 categories with RGB, Near-Infrared (NIR), Degree of Linear Polarization (DoLP), and Angle of Linear Polarization (AoLP) images. Implementation Details.

Our Any2Seg is trained on 8 NVIDIA

A800 GPUs, starting with a learning rate of 6 e^{-5} , adjusted by a poly strategy (power 0.9) across 200 epochs, including an initial 10-epoch warm-up at 0.1 times the learning rate. We use the AdamW optimizer with a batch size of 1 per GPU. Data augmentation includes random resizing (0.5-2.0 ratio), horizontal flipping, color jitter, gaussian blur, and cropping to 1024 × 1024 on [60] and 512 × 512 on [24]. **Multi-modal Semantic Segmen-**

tation (MSS) Tabs. 1 and 2 present a detailed quantitative evaluation,

Table 2: MSS results on MCubeS [24].

Method	Modal	mIoU
MCubeSNet [24]	I	33.70
MCubeSNet [24]	I-A	39.10
MCubeSNet [24]	I-A-D	42.00
MCubeSNet [24]	I-A-D	42.86
CMNeXt [60] (MiT-B2)	I-A	48.82
CMNeXt [60] (MiT-B2)	I-A-D	49.48
CMNeXt [60] (MiT-B2)	I-A-D-N	51.54
Ours (MiT-B2)	I-A	48.95
Ours (MiT-B2)	I-A-D	50.67
Ours (MiT-B2)	I-A-D-N	51.91

contrasting our proposed framework with previous SoTA methods. Any2Seg outperforms other methods in almost all multi-modal scenarios. Specifically, for dual-modality inputs on the DELIVER dataset, Any2Seg significantly exceeds the prior SoTA, CMNeXt [60], with mIoU improvements of +3.54 and

Table 3: Results of system-level MISS evaluation. All methods are trained with four modalities, and the metric is mIoU for all numbers.





Fig. 4: System-level Modality-Incomplete Semantic Segmentation (MISS) validation results on DEVLIER Dataset (D: depth; L: LiDAR; and E: Event).

+0.35 for the R+D and R+E configurations, respectively. In scenarios using triplet modality inputs, our Any2Seg continues to lead, surpassing the SoTA with mIoU enhancements of +2.31 and +1.70 for the R+D+E and R+D+L modalities. Most notably, by incorporating all four modalities, our Any2Seg not only maintains its leading edge but also reaches an unprecedented performance benchmark with a 68.25 mIoU, achieving a significant improvement of +1.95 mIoU over the CMNeXt [60]. Additionally, Any2Seg consistently delivers the highest performance on the MCubeS across various methods with only the proposed MFFM, as shown in Tab. 2.

Modality-Incomplete Semantic Segmentation (MISS)

System-level MISS: As shown in Tab. 3, Any2Seg significantly outperforms the CMNeXt framework [60] in nearly all configurations, with a notable mean improvement of +19.79 mIoU, especially in RGB and Depth combinations. While performance dips in scenarios with sparse modalities like Event and Li-DAR, Any2Seg's design to harness multi-modal information enables it to excel in dense modality scenarios, significantly outdoing the benchmarks with a remarkable 60.15 mIoU in the DL (depth and LiDAR) modality compared to CMNeXt's 3.83 mIoU. This highlights the efficacy of Any2Seg in integrating multi-modal information to improve segmentation, as shown in Fig. 4.

Sensor-level MISS: Our evaluation also encompasses sensor-level failures, leveraging the diverse adverse conditions present in the DELIVER dataset's training

Table 4: Snesor-level MISS results on DELIVER. MB: Motion Blur; OE: Over-Exposure; UE: Under-Exposure; LJ: LiDAR-Jitter; and EL: Event Low-resolution. The **bold** and <u>underline</u> denote the best and the second-best performance, respectively.

Method	Modality	Cloudy	Foggy	Night	Rainy	Sunny	MB	OE	UE	LJ	\mathbf{EL}	Mean
HRFuser [56] SegFormer [52]	R R	$49.26 \\ 59.99$	$48.64 \\ 57.30$	$42.57 \\ 50.45$	$50.61 \\ 58.69$	$50.47 \\ 60.21$	48.33 57.28	$35.13 \\ 56.64$	$26.86 \\ 37.44$	$49.06 \\ 57.17$	49.88 59.12	47.95 57.20
TokenFusion [46] CMX [58] HRFuser [56]	R-D R-D R-D	50.92 63.70 54.80	52.02 62.77 51.48	$\begin{array}{r} 43.37 \\ 60.74 \\ 49.51 \end{array}$	50.70 62.37 51.55	$52.21 \\ 63.14 \\ 52.12$	49.22 59.50 50.92	$\begin{array}{c} 46.22 \\ 60.14 \\ 41.51 \end{array}$	36.39 55.84 44.00	49.58 62.65 54.10	49.17 63.26 52.52	49.86 62.66 51.88
HRFuser [56] HRFuser [56]	R-D-E R-D-E-L	$54.04 \\ 56.20$	50.83 52.39	50.88 49.85	$51.13 \\ 52.53$	$\begin{array}{c} 52.61 \\ 54.02 \end{array}$	49.32 49.44	$\begin{array}{c} 41.75\\ 46.31 \end{array}$	$\begin{array}{c} 47.89\\ 46.92 \end{array}$	$54.65 \\ 53.94$	52.33 52.72	51.83 52.97
[60] w/ S-2 [60] w/ S-2 [60] w/ S-2	R-D-E R-D-L R-D-E-L	60.02 68.28 <u>68.70</u>	$56.16 \\ 63.28 \\ 65.67$	54.03 62.64 62.46	54.82 63.01 67.50	$58.29 \\ 66.06 \\ 66.57$	$53.70 \\ 62.58 \\ 62.91$	$57.04 \\ 64.44 \\ 64.59$	$51.98 \\ 58.73 \\ 60.00$	$58.54 \\ 65.37 \\ 65.92$	55.87 65.80 65.48	$56.05 \\ 64.02 \\ 64.98$
Ours w/ S-2 Ours w/ S-2 Ours w/ S-2	R-D-E R-D-L R-D-E-L	68.35 68.38 69.19	66.03 <u>67.18</u> 67.64	64.60 66.68 67.25	$66.61 \\ 65.06 \\ 67.00$	67.31 <u>67.73</u> 68.25	63.62 63.34 66.23	64.61 66.11 <u>64.89</u>	62.34 64.49 65.69	65.55 <u>67.54</u> 69.05	64.20 66.40 66.73	65.32 66.29 67.19

and validation sets, such as cloudiness and over-exposure. These conditions serve as effective data augmentation, against which Any2Seg consistently excels, surpassing prior state-of-the-art methods under challenging conditions, for instance, achieving 69.19 mIoU in cloudy weather, as depicted in Tab. 4. This performance attests to our LSCD and MFF modules' proficiency in minimizing modality gaps, fusing multi-modal data and extracting durable, modality-agnostic features for reliable segmentation across various conditions and sensor failure cases.

5 Ablation Study

5.1 Effectiveness of Loss Functions

We first conduct ablation studies to evaluate the proposed losses. The results in Tab. 5 show that all losses play positive roles. For instance, adding \mathcal{L}_{sup} to the baseline leads to gains of +1.04 mIoU, +1.16 F1, and +2.32 Acc. Obviously, using all three losses results in the highest performance, *i.e.*, 63.74 mIoU, 73.33 F1, and 73.64 Acc, demonstrating their effectiveness.

5.2 Ablation of Components in MFF module

To assess the effectiveness of MFFM's key components—feature re-weighting, fine-grained feature selection, and residual connections—we undertake a thorough evaluation across MSS and MISS settings within two datasets, covering four modalities as shown in Tab. 6.

Fine-grained feature selection. Initially, modality-balanced re-weighting alone yields a mIoU of 39.63 and 56.24 in MSS validation on both datasets. The incorporation of our fine-grained feature selection further augments the performance

\mathcal{L}_{base}	ſ	ſ	ſ	RGB-Depth										
	∠ se	~ cr	\sim sup	mIoU	Δ	F1	Δ	Acc	Δ					
\checkmark				61.93	-	71.53	-	70.35	-					
\checkmark	\checkmark			62.89	+0.96	72.52	+0.99	72.14	+1.79					
\checkmark		\checkmark		62.86	+0.93	72.36	+0.83	72.81	+2.46					
			\checkmark	62.97	+1.04	72.69	+1.16	72.67	+2.32					
	\checkmark		\checkmark	63.41	+1.48	73.15	+1.62	73.19	+2.84					
		\checkmark	\checkmark	63.14	+1.21	72.97	+1.44	73.35	+3.00					
	\checkmark	\checkmark	\checkmark	63.74	+1.81	73.33	+1.80	73.64	+3.29					

Table 5: Ablation of different losses on DELIVER with SegFormer-B0 in MSS setting.

Table 6: Ablation of MFF with SegFormer-B0 in both MSS and MISS settings.

	MCubeS Dataset with all four modalities																
Re-	F	usior	ı	MSS							MISS						
weight	Selectio	on Re	sidu	alµnIoU	Δ	Acc	Δ	F1	Δ	mIoU	Δ	Acc	Δ	F1	Δ		
\checkmark				39.63	-	51.06	-	47.97	-	27.88	-	35.10	-	35.68	-		
\checkmark	🗸			42.49	+2.86	54.90	+3.84	52.09	+4.12	2 29.10	+1.22	38.24	+3.14	36.56	+0.88		
\checkmark			\checkmark	42.86	+3.23	54.91	+3.85	52.62	+4.65	5 25.10	-2.78	33.25	-1.85	32.95	-2.73		
\checkmark	✓		\checkmark	44.20	+4.57	56.60	+5.54	54.12	+6.13	5 29.14	+1.26	38.81	+3.71	40.75	+5.07		
				Ι	DELIV	ER D	ataset	with a	ull four	r moda	lities						
Re-	F	usior	1			Μ	ISS			MISS							
weight	Selectio	on Re	sidu	al mIoU	Δ	Acc	Δ	F1	Δ	mIoU	Δ	Acc	Δ	F1	Δ		
\checkmark				56.24	-	67.33	-	73.08	-	35.47	-	44.36	-	53.96	-		
\checkmark	🗸			61.61	+5.37	70.98	+3.65	71.61	-1.47	44.49	+9.02	52.79	+8.43	54.28	+0.32		
\checkmark			\checkmark	60.96	+4.72	71.35	+4.02	74.12	+1.04	36.22	+0.75	45.06	+0.70	52.51	-1.45		
\checkmark	🗸		\checkmark	62.97	+6.73	76.03	+8.70	72.67	-0.41	47.79	+12.32	257.12	+12.76	3 56.69	+2.79		

significantly, enhancing the mIoU by +2.86 and +5.37, Acc by +3.84 and +3.65 on both datasets. Concurrently, MISS validation also exhibits improvements, with increases of +1.22 and +9.02 mIoU, +3.14 and +8.43 Acc. Obviously, the performance gains brought by our fine-grained feature selection are larger in modality-incomplete segmentation settings, which indicates that our proposed selection algorithm is efficient in selecting the fine-grained feature-agnostic features from the multi-modal inputs.

Residual connections. The residual connections which aggregates the balanced feature with the selected feature, leads to significant gains of +3.23 and +4.72 in mIoU, +3.85 and +4.02 in Acc, and +4.65 and +1.04 in F1 during MSS on both datasets, as depicted in Tab. 6. However, using the residual connection alone, without the fine-grained feature selection, results in mixed outcomes across all three segmentation metrics: -2.78 and +0.75 in mIoU, -1.85 and +0.70 in Acc,



Learning Modality-agnostic Representation for Semantic Segmentation

13

(b) Night with under-exposure conditions (d) w/ Semantic KD

Depth

RGB

Fig. 5: Visualization of multi-modal features under different conditions on DEVLIER. (a) sun with over-exposure, and (b) night with under-exposure.

Modality-agnostic

and -2.73 and -1.45 in F1. These findings highlight that while the residual connection mechanism enhances MSS, the fine-grained feature selection mechanism plays a pivotal role in improving MISS performance.

Effectiveness of the Overall MFF module. Integrating fine-grained feature selection with residual connections substantially enhances performance: in multimodal validation, mIoU increases by +4.57 and +6.73 and Acc by +5.54 and +8.70; modality-incomplete segmentation sees mIoU gains of +1.26 and +12.32 and Acc of +3.71 and +12.76 across datasets. Meanwhile, using the selection mechanism alone results in a 61.61 mIoU in multi-modal segmentation on DE-LIVER, which rises to 62.97 with the residual connection, as shown in Tab. 6. In modality-incomplete segmentation, combining all mechanisms increases performance from 36.22 to 47.79 mIoU. Meanwhile the qualitative results in Fig. 5 illustrates that modality-agnostic features merge the advantages of different modalities, *i.e.*, RGB's textual information and depth cameras' lighting in-variance. These results highlight MFF's crucial impact on extracting modality-agnostic features and improving segmentation accuracy in varied scenarios.

5.3 Ablation of KD in Two Embedding Spaces

Effectiveness of Intra-modal Semantic KD. We conduct ablation studies on DELIVER, employing both single and dual modality inputs while contrasting the effects of solely using the basic loss \mathcal{L}_{base} against the inclusion of our proposed semantic KD loss \mathcal{L}_{se} . The outcomes are shown in Tab. 7, reveal that with a

Table 7: Ablation study of \mathcal{L}_{se} .

Table 8: Ablation study of \mathcal{L}_{cr} .

	\mathcal{L}_{base}	$\mathcal{L}_{s\epsilon}$	mAcc	Δ	1	mIoU	Δ	F1	Δ		\mathcal{L}_{base}	\mathcal{L}_{cr}	mAcc	Δ	mIoU	Δ	F1	Δ
R	√ √	~	$59.79 \\ 61.15$	- +1.	.36	51.39 53.03 -	- +1.6	61.59 463.63 -	- +2.04	R+D	√ ✓	~	61.93 62.86	- + 0.9 3	71.53 72.36 -	- +0.8	70.35 3 72.81	- + 2.46
	\mathcal{L}_{base}	\mathcal{L}_{se}	mAcc	Δ	7	mIoU	Δ	F1	Δ		\mathcal{L}_{base}	\mathcal{L}_{cr}	mAcc	Δ	mIoU	Δ	F1	Δ
R+F		~	60.36 60.20	-0.	16	52.21 53.11 -	- +0.9	62.64 063.24-	- 	R+E	\checkmark	~	52.21 53.26	+1.05	62.74 63.55 -	- +0.8	60.36 1 61.02	- +0.66

singular modality input, specifically RGB data, the integration of our semantic space KD results in notable improvements: a +1.36 increase in mAcc, a +1.64 enhancement in mIoU, and a +2.04 uplift in F1 score. In scenarios involving dual modalities, the application of our KD loss \mathcal{L}_{se} yields further gains: +1.05 in mAcc, +0.81 in mIoU, and +0.66 in F1 score. Fig. 5 (c) and (d) illustrate that our semantic KD facilitates enhanced differentiation of semantics in the high-level space. These findings robustly affirm the efficacy of implementing KD within the semantic embedding space.

Effectiveness of Inter-modal Correlation KD. We conduct ablation studies on DELIVER dataset with RGB-Event and RGB-Depth modalities while contrasting the effects of solely using the basic \mathcal{L}_{base} against the inclusion of our proposed KD loss \mathcal{L}_{cr} , as shown in Tab. 8. Experimental results shows that in RGB-Depth setting, the integration of our inter-modal correlation KD leads to notable improvements: +0.93 mIoU, +0.83 F1, and +2.46 Acc score. Consistently, in RGB-Event setting, our \mathcal{L}_{se} also achieves +0.90 mIoU improvement.

6 Conclusion

In this paper, we made the first attempt to explore the potential of current MVLMs and introduced **Any2Seg** framework that can achieve robust segmentation from any combination of modalities in any visual conditions with the guidance from MVLMs. We proposed a Language-guided Semantic Correlation Distillation (**LSCD**) module for distilling semantic and cross-modal knowledge from MVLMs and a modality-agnostic feature fusion (**MFF**) module to transform multi-modal representation into modality-agnostic forms, ensuring robust segmentation. Our Any2Seg significantly outperformed previous multi-modal methods in both multi-modal and modality-incomplete settings with two public benchmarks, achieving new state-of-the-art segmentation performance.

Limitation and Future Work. Despite the contributions, there are still limitations of Any2Seg: (a) current MVLMs are pretrained at image-level rather than pixel-level to binding multi-modal data, restricting the dense pixel-level multimodal segmentation performance; and (b) MSS and MISS performance remains sub-optimal in scenarios with sparse data inputs. Future research will concentrate on developing pixel-wise MVLMs and improving the modality-agnostic module to achieve uniform and enhanced performance across all scenarios.

7 Acknowledgement

This paper is supported by the National Natural Science Foundation of China (NSF) under Grant No. NSFC22FYT45, the Guangzhou City, University and Enterprise Joint Fund under Grant No.SL2022A03J01278, and Guangzhou Fundamental and Applied Basic Research (Grant Number: 2024A04J4072).

References

- Alonso, I., Murillo, A.C.: Ev-segnet: Semantic segmentation for event-based cameras. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Borse, S., Klingner, M., Kumar, V.R., Cai, H., Almuzairee, A., Yogamani, S., Porikli, F.: X-align: Cross-modal cross-view alignment for bird's-eye-view segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3287–3297 (2023)
- Broedermann, T., Sakaridis, C., Dai, D., Van Gool, L.: Hrfuser: A multi-resolution sensor fusion architecture for 2d object detection. arXiv preprint arXiv:2206.15157 (2022)
- Cao, J., Zheng, X., Lyu, Y., Wang, J., Xu, R., Wang, L.: Chasing day and night: Towards robust and efficient all-day object detection guided by an event camera. arXiv preprint arXiv:2309.09297 (2023)
- Cao, J., Leng, H., Lischinski, D., Cohen-Or, D., Tu, C., Li, Y.: Shapeconv: Shapeaware convolutional layer for indoor rgb-d semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7088–7097 (2021)
- Chen, G., Shao, F., Chai, X., Chen, H., Jiang, Q., Meng, X., Ho, Y.S.: Modalityinduced transfer-fusion network for rgb-d and rgb-t salient object detection. IEEE Transactions on Circuits and Systems for Video Technology 33(4), 1787–1801 (2022)
- Chen, J., Deguchi, D., Zhang, C., Zheng, X., Murase, H.: Clip is also a good teacher: A new learning framework for inductive zero-shot semantic segmentation. arXiv preprint arXiv:2310.02296 (2023)
- Chen, J., Deguchi, D., Zhang, C., Zheng, X., Murase, H.: Frozen is better than learning: A new design of prototype-based classifier for semantic segmentation. Pattern Recognition 152, 110431 (2024)
- Chen, L.Z., Lin, Z., Wang, Z., Yang, Y.L., Cheng, M.M.: Spatial information guided convolution for real-time rgbd semantic segmentation. IEEE Transactions on Image Processing 30, 2313–2324 (2021)
- Chen, M., Yao, J., Xing, L., Wang, Y., Zhang, Y., Wang, Y.: Redundancy-adaptive multimodal learning for imperfect data. arXiv preprint arXiv:2310.14496 (2023)
- Cheng, Y., Wei, F., Bao, J., Chen, D., Zhang, W.: Cico: Domain-aware sign language retrieval via cross-lingual contrastive learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19016– 19026 (2023)
- Cong, R., Lin, Q., Zhang, C., Li, C., Cao, X., Huang, Q., Zhao, Y.: Cir-net: Crossmodality interaction and refinement for rgb-d salient object detection. IEEE Transactions on Image Processing **31**, 6800–6815 (2022)

- 16 X. Zheng et al.
- Girdhar, R., El-Nouby, A., Liu, Z., Singh, M., Alwala, K.V., Joulin, A., Misra, I.: Imagebind: One embedding space to bind them all. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15180– 15190 (2023)
- 14. Guo, Z., Zhang, R., Zhu, X., Tang, Y., Ma, X., Han, J., Chen, K., Gao, P., Li, X., Li, H., et al.: Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. arXiv preprint arXiv:2309.00615 (2023)
- Han, J., Zhang, R., Shao, W., Gao, P., Xu, P., Xiao, H., Zhang, K., Liu, C., Wen, S., Guo, Z., et al.: Imagebind-llm: Multi-modality instruction tuning. arXiv preprint arXiv:2309.03905 (2023)
- Huang, K., Shi, B., Li, X., Li, X., Huang, S., Li, Y.: Multi-modal sensor fusion for auto driving perception: A survey. arXiv preprint arXiv:2202.02703 (2022)
- Huang, T., Dong, B., Yang, Y., Huang, X., Lau, R.W., Ouyang, W., Zuo, W.: Clip2point: Transfer clip to point cloud classification with image-depth pretraining. arXiv preprint arXiv:2210.01055 (2022)
- Hui, T., Xun, Z., Peng, F., Huang, J., Wei, X., Wei, X., Dai, J., Han, J., Liu, S.: Bridging search region interaction with template for rgb-t tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13630–13639 (2023)
- Ji, W., Yan, G., Li, J., Piao, Y., Yao, S., Zhang, M., Cheng, L., Lu, H.: Dmra: Depth-induced multi-scale recurrent attention network for rgb-d saliency detection. IEEE Transactions on Image Processing **31**, 2321–2336 (2022)
- Jia, Z., You, K., He, W., Tian, Y., Feng, Y., Wang, Y., Jia, X., Lou, Y., Zhang, J., Li, G., et al.: Event-based semantic segmentation with posterior attention. IEEE Transactions on Image Processing **32**, 1829–1842 (2023)
- Lee, M., Park, C., Cho, S., Lee, S.: Spsn: Superpixel prototype sampling network for rgb-d salient object detection. In: European Conference on Computer Vision. pp. 630–647. Springer (2022)
- Li, J., Dai, H., Han, H., Ding, Y.: Mseg3d: Multi-modal 3d semantic segmentation for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21694–21704 (2023)
- Li, Y., Yu, A.W., Meng, T., Caine, B., Ngiam, J., Peng, D., Shen, J., Lu, Y., Zhou, D., Le, Q.V., et al.: Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17182–17191 (2022)
- Liang, Y., Wakaki, R., Nobuhara, S., Nishino, K.: Multimodal material segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19800–19808 (2022)
- Liao, G., Gao, W., Li, G., Wang, J., Kwong, S.: Cross-collaborative fusion-encoder network for robust rgb-thermal salient object detection. IEEE Transactions on Circuits and Systems for Video Technology 32(11), 7646–7661 (2022)
- Liu, H., Lu, T., Xu, Y., Liu, J., Li, W., Chen, L.: Camliflow: Bidirectional cameralidar fusion for joint optical flow and scene flow estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5791– 5801 (2022)
- Liu, R., Zhang, J., Peng, K., Chen, Y., Cao, K., Zheng, J., Sarfraz, M.S., Yang, K., Stiefelhagen, R.: Fourier prompt tuning for modality-incomplete scene segmentation. arXiv preprint arXiv:2401.16923 (2024)
- Lyu, Y., Zheng, X., Kim, D., Wang, L.: Omnibind: Teach to build unequal-scale modality interaction for omni-bind of all. arXiv preprint arXiv:2405.16108 (2024)

Learning Modality-agnostic Representation for Semantic Segmentation

17

- Lyu, Y., Zheng, X., Zhou, J., Wang, L.: Unibind: Llm-augmented unified and balanced representation space to bind them all. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26752–26762 (2024)
- Maheshwari, H., Liu, Y.C., Kira, Z.: Missing modality robustness in semisupervised multi-modal semantic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1020–1030 (2024)
- Man, Y., Gui, L.Y., Wang, Y.X.: Bev-guided multi-modality fusion for driving perception. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21960–21969 (2023)
- Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate lidar semantic segmentation. In: 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 4213–4220. IEEE (2019)
- Pang, Y., Zhao, X., Zhang, L., Lu, H.: Caver: Cross-modal view-mixed transformer for bi-modal salient object detection. IEEE Transactions on Image Processing 32, 892–904 (2023)
- Park, S.J., Hong, K.S., Lee, S.: Rdfnet: Rgb-d multi-level residual feature fusion for indoor semantic segmentation. In: Proceedings of the IEEE international conference on computer vision. pp. 4980–4989 (2017)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- 36. Reza, M.K., Prater-Bennette, A., Asif, M.S.: Robust multimodal learning with missing modalities via parameter-efficient adaptation. arXiv preprint arXiv:2310.03986 (2023)
- Shivakumar, S.S., Rodrigues, N., Zhou, A., Miller, I.D., Kumar, V., Taylor, C.J.: Pst900: Rgb-thermal calibration, dataset and segmentation network. In: 2020 IEEE international conference on robotics and automation (ICRA). pp. 9441–9447. IEEE (2020)
- Song, M., Song, W., Yang, G., Chen, C.: Improving rgb-d salient object detection via modality-aware decoder. IEEE Transactions on Image Processing 31, 6124– 6138 (2022)
- Su, Y., Lan, T., Li, H., Xu, J., Wang, Y., Cai, D.: Pandagpt: One model to instruction-follow them all. arXiv preprint arXiv:2305.16355 (2023)
- 40. Sun, W., Zhang, J., Wang, J., Liu, Z., Zhong, Y., Feng, T., Guo, Y., Zhang, Y., Barnes, N.: Learning audio-visual source localization via false negative aware contrastive learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6420–6429 (2023)
- Sun, Y., Zuo, W., Liu, M.: Rtfnet: Rgb-thermal fusion network for semantic segmentation of urban scenes. IEEE Robotics and Automation Letters 4(3), 2576– 2583 (2019)
- Wang, F., Pan, J., Xu, S., Tang, J.: Learning discriminative cross-modality features for rgb-d saliency detection. IEEE Transactions on Image Processing **31**, 1285–1297 (2022)
- Wang, H., Ma, C., Zhang, J., Zhang, Y., Avery, J., Hull, L., Carneiro, G.: Learnable cross-modal knowledge distillation for multi-modal learning with missing modality. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 216–226. Springer (2023)
- 44. Wang, J., Wang, Z., Tao, D., See, S., Wang, G.: Learning common and specific features for rgb-d semantic segmentation with deconvolutional networks. In: Com-

puter Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14. pp. 664–679. Springer (2016)

- Wang, S., Caesar, H., Nan, L., Kooij, J.F.: Unibev: Multi-modal 3d object detection with uniform bev encoders for robustness against missing sensor modalities. arXiv preprint arXiv:2309.14516 (2023)
- Wang, Y., Chen, X., Cao, L., Huang, W., Sun, F., Wang, Y.: Multimodal token fusion for vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12186–12195 (2022)
- 47. Wang, Y., Huang, W., Sun, F., Xu, T., Rong, Y., Huang, J.: Deep multimodal fusion by channel exchanging. Advances in neural information processing systems 33, 4835–4845 (2020)
- Wang, Y., Sun, F., Lu, M., Yao, A.: Learning deep multimodal feature representation with asymmetric multi-layer fusion. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 3902–3910 (2020)
- Wang, Y., Mao, Q., Zhu, H., Deng, J., Zhang, Y., Ji, J., Li, H., Zhang, Y.: Multimodal 3d object detection in autonomous driving: a survey. International Journal of Computer Vision pp. 1–31 (2023)
- Wei, S., Luo, C., Luo, Y.: Mmanet: Margin-aware distillation and modalityaware regularization for incomplete multimodal learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20039– 20049 (2023)
- 51. Wu, W., Chu, T., Liu, Q.: Complementarity-aware cross-modal feature fusion network for rgb-t semantic segmentation. Pattern Recognition **131**, 108881 (2022)
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems 34, 12077–12090 (2021)
- 53. Xie, Z., Shao, F., Chen, G., Chen, H., Jiang, Q., Meng, X., Ho, Y.S.: Cross-modality double bidirectional interaction and fusion network for rgb-t salient object detection. IEEE Transactions on Circuits and Systems for Video Technology (2023)
- Yan, X., Gao, J., Zheng, C., Zheng, C., Zhang, R., Cui, S., Li, Z.: 2dpass: 2d priors assisted semantic segmentation on lidar point clouds. In: European Conference on Computer Vision. pp. 677–695. Springer (2022)
- Ying, X., Chuah, M.C.: Uctnet: Uncertainty-aware cross-modal transformer network for indoor rgb-d semantic segmentation. In: European Conference on Computer Vision. pp. 20–37. Springer (2022)
- Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution vision transformer for dense predict. Advances in Neural Information Processing Systems 34, 7281–7293 (2021)
- Zhang, B., Wang, Z., Ling, Y., Guan, Y., Zhang, S., Li, W.: Mx2m: Masked crossmodality modeling in domain adaptation for 3d semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3401–3409 (2023)
- Zhang, J., Liu, H., Yang, K., Hu, X., Liu, R., Stiefelhagen, R.: Cmx: Crossmodal fusion for rgb-x semantic segmentation with transformers. arXiv preprint arXiv:2203.04838 (2022)
- Zhang, J., Liu, H., Yang, K., Hu, X., Liu, R., Stiefelhagen, R.: Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. IEEE Transactions on Intelligent Transportation Systems (2023)
- 60. Zhang, J., Liu, R., Shi, H., Yang, K., Reiß, S., Peng, K., Fu, H., Wang, K., Stiefelhagen, R.: Delivering arbitrary-modal semantic segmentation. In: Proceedings of the

Learning Modality-agnostic Representation for Semantic Segmentation

IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1136–1147 (2023)

- Zhang, J., Yang, K., Stiefelhagen, R.: Issafe: Improving semantic segmentation in accidents by fusing event-based data. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1132–1139. IEEE (2021)
- 62. Zhang, Q., Zhao, S., Luo, Y., Zhang, D., Huang, N., Han, J.: Abmdrnet: Adaptiveweighted bi-directional modality difference reduction network for rgb-t semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2633–2642 (2021)
- 63. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8552–8562 (2022)
- 64. Zhang, R., Han, J., Zhou, A., Hu, X., Yan, S., Lu, P., Li, H., Gao, P., Qiao, Y.: Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199 (2023)
- Zhang, T., Guo, H., Jiao, Q., Zhang, Q., Han, J.: Efficient rgb-t tracking via crossmodality distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5404–5413 (2023)
- 66. Zhang, Y., Zhou, Z., David, P., Yue, X., Xi, Z., Gong, B., Foroosh, H.: Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9601–9610 (2020)
- Zhang, Y., Gong, K., Zhang, K., Li, H., Qiao, Y., Ouyang, W., Yue, X.: Meta-transformer: A unified framework for multimodal learning. arXiv preprint arXiv:2307.10802 (2023)
- 68. Zhao, Z., Palani, H., Liu, T., Evans, L., Toner, R.: Multi-modality guidance network for missing modality inference. arXiv preprint arXiv:2309.03452 (2023)
- 69. Zheng, J., Wang, Y., Tan, C., Li, S., Wang, G., Xia, J., Chen, Y., Li, S.Z.: Cvt-slr: Contrastive visual-textual transformation for sign language recognition with variational alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23141–23150 (2023)
- Zheng, X., Liu, Y., Lu, Y., Hua, T., Pan, T., Zhang, W., Tao, D., Wang, L.: Deep learning for event-based vision: A comprehensive survey and benchmarks. arXiv preprint arXiv:2302.08890 (2023)
- Zheng, X., Luo, Y., Wang, H., Fu, C., Wang, L.: Transformer-cnn cohort: Semisupervised semantic segmentation by the best of both students. arXiv preprint arXiv:2209.02178 (2022)
- 72. Zheng, X., Luo, Y., Zhou, P., Wang, L.: Distilling efficient vision transformers from cnns for semantic segmentation. arXiv preprint arXiv:2310.07265 (2023)
- Zheng, X., Pan, T., Luo, Y., Wang, L.: Look at the neighbor: Distortion-aware unsupervised domain adaptation for panoramic semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18687–18698 (2023)
- Zheng, X., Wang, L.: Eventdance: Unsupervised source-free cross-modal adaptation for event-based object recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17448–17458 (2024)
- Zheng, X., Zhou, P., Vasilakos, A.V., Wang, L.: 360sfuda++: Towards source-free uda for panoramic segmentation by learning reliable category prototypes. arXiv preprint arXiv:2404.16501 (2024)

- 20 X. Zheng et al.
- Zheng, X., Zhou, P., Vasilakos, A.V., Wang, L.: Semantics distortion and style matter: Towards source-free uda for panoramic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 27885– 27895 (2024)
- 77. Zheng, X., Zhu, J., Liu, Y., Cao, Z., Fu, C., Wang, L.: Both style and distortion matter: Dual-path unsupervised domain adaptation for panoramic semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1285–1295 (2023)
- Zhou, H., Qi, L., Wan, Z., Huang, H., Yang, X.: Rgb-d co-attention network for semantic segmentation. In: Proceedings of the Asian conference on computer vision (2020)
- 79. Zhou, J., Zheng, X., Lyu, Y., Wang, L.: E-clip: Towards label-efficient event-based open-world understanding by clip. arXiv preprint arXiv:2308.03135 (2023)
- Zhou, J., Zheng, X., Lyu, Y., Wang, L.: Exact: Language-guided conceptual reasoning and uncertainty estimation for event-based action recognition and more. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18633–18643 (2024)
- Zhou, W., Zhang, H., Yan, W., Lin, W.: Mmsmcnet: Modal memory sharing and morphological complementary networks for rgb-t urban scene semantic segmentation. IEEE Transactions on Circuits and Systems for Video Technology (2023)
- 82. Zhu, B., Lin, B., Ning, M., Yan, Y., Cui, J., Wang, H., Pang, Y., Jiang, W., Zhang, J., Li, Z., et al.: Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. arXiv preprint arXiv:2310.01852 (2023)
- Zhu, J., Lai, S., Chen, X., Wang, D., Lu, H.: Visual prompt multi-modal tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9516–9526 (2023)
- Zhu, J., Luo, Y., Zheng, X., Wang, H., Wang, L.: A good student is cooperative and reliable: Cnn-transformer collaborative learning for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11720–11730 (2023)
- Zhu, X., Zhang, R., He, B., Zeng, Z., Zhang, S., Gao, P.: Pointclip v2: Adapting clip for powerful 3d open-world learning. arXiv preprint arXiv:2211.11682 (2022)
- Zhuang, Z., Li, R., Jia, K., Wang, Q., Li, Y., Tan, M.: Perception-aware multisensor fusion for 3d lidar semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16280–16290 (2021)