

Complementary Random Masking for RGB-Thermal Semantic Segmentation

Ukcheol Shin¹, Kyunghyun Lee², In So Kweon², Jean Oh¹

Abstract—RGB-thermal semantic segmentation is one potential solution to achieve reliable semantic scene understanding in adverse weather and lighting conditions. However, the previous studies mostly focus on designing a multi-modal fusion module without consideration of the nature of multi-modality inputs. Therefore, the networks easily become over-reliant on a single modality, making it difficult to learn complementary and meaningful representations for each modality. This paper proposes 1) a complementary random masking strategy of RGB-T images and 2) self-distillation loss between clean and masked input modalities. The proposed masking strategy prevents over-reliance on a single modality. It also improves the accuracy and robustness of the neural network by forcing the network to segment and classify objects even when one modality is partially available. Also, the proposed self-distillation loss encourages the network to extract complementary and meaningful representations from a single modality or complementary masked modalities. We achieve state-of-the-art performance over three RGB-T semantic segmentation benchmarks. Our source code is available at https://github.com/UkcheolShin/CRM_RGBTseg.

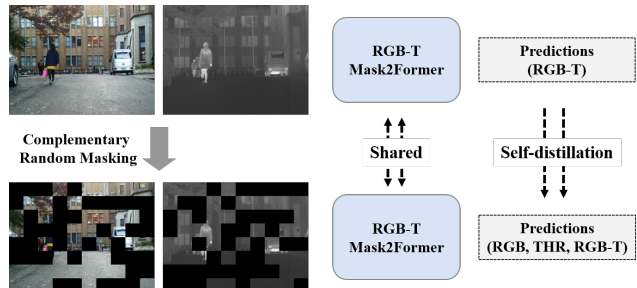
I. INTRODUCTION

Robust and reliable semantic scene understanding is a crucial fundamental ability for autonomous driving to ensure the safe and reliable operation of autonomous vehicles. RGB-thermal semantic segmentation is one potential solution to achieve reliable semantic scene understanding in adverse weather and lighting conditions. For example, in foggy or low-light conditions, the RGB camera may struggle to capture objects in the scene due to reduced visibility. In contrast, the thermal camera can still detect the heat signatures of objects. Combining the information from both modalities enables reliable and accurate semantic segmentation in adverse scenarios. Therefore, it naturally led to recent active studies on semantic segmentation of RGB-thermal images [11], [36], [35], [37], [48], [17], [21], [47].

The primary research direction is designing a multi-modal fusion module [44], [46], [8], [21] that can effectively combine the information from both RGB and thermal modalities to improve the accuracy of semantic segmentation. However, without consideration of the nature of multi-modal inputs, the networks easily fall into a sub-optimal solution, where the network becomes over-reliant on a single modality, as shown in Fig. 2 and Tab. I. In addition, this implies that the networks are susceptible to a wide range of fault cases,

¹U. Shin and J. Oh are with Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, 15217, United States {ushin, hyaejino}@andrew.cmu.edu

²K. Lee and I. S. Kweon are with the School of Electrical Engineering, KAIST, Daejeon, 34141, Republic of Korea. {kyunghyun.lee, iskweon77}@kaist.ac.kr



(a) Complementary random masking for RGB-T segmentation



(b) RGB image (c) Thermal image (d) Ground Truth



(e) RTFNet [36] (f) CMXNet [21] (g) Ours

Fig. 1: Complementary random masking for RGB-thermal semantic segmentation. Our proposed method aims to learn meaningful and complementary representations from RGB and thermal images by using complementary masking of RGB-T inputs and ensuring consistency between augmented and original inputs. The proposed method leads to robust and reliable segmentation results in day-light, low-light, and modality-dropped scenarios.

such as sensor disconnection, lens occlusion, and other input quality degeneration.

In this paper, we focus on learning complementary and meaningful representations from both RGB and thermal modalities to prevent the over-reliance problem on a single modality and eventually improve the accuracy of the segmentation model. For this purpose, our intuitive ideas are as follows and shown in Fig. 1: 1) We augment input RGB-T images with random masking to prevent the network from over-reliantly utilizing one modality for the semantic segmentation task. 2) We enforce consistency between the prediction results of augmented and original images to encourage the network to extract meaningful representations even from partially occluded modalities or a single modality.




Fig. 2: **Input modality dependency comparison of RGB-T semantic segmentation networks.** Common multi-modal fusion approaches often result in a sub-optimal solution, where the neural network becomes over-reliant on a single modality, as shown in (e) and (f). On the other hand, our proposed method prevents the over-reliance issue (*i.e.*, (h) and (i)).

Our contributions can be summarized as follows:

- We propose a complementary random masking strategy that randomly masks one modality and masks the other modality in a complementary manner to improve model robustness and accuracy.
- We propose a self-distillation loss between the prediction result from clean input modalities and the multiple prediction results from masked input modalities to learn complementary and non-local representations.
- Ours proposed method achieves state-of-the-art results over three RGB-T benchmark datasets (*i.e.*, MFNet [11], PST900 [35], and KP [14], [17] datasets).

II. RELATED WORK

A. RGB-Thermal Semantic Segmentation

Recently, thermal images have been widely adopted in various applications, such as detection [19], [45], tracking [20], [15], feature matching [18], [26], depth estimation [32], [31], [33], and SLAM [34], [16], to achieve high-level robustness against adverse weather and lighting conditions.

RGB-thermal semantic segmentation networks have also been proposed to overcome the limitation of RGB semantic segmentation networks [4], [13], [40], [6], [5] that are often vulnerable to extreme conditions, such as low-light, rainy, snowy, and sandy conditions. Most previous RGB-T fusion networks focused on designing a multi-modal fusion module that can effectively combine the information from both modalities to improve the accuracy of semantic segmentation. Specifically, they proposed various types of feature fusion modules, such as the naïve feature-level fusion [11], [36], [35], multi-scale feature fusion [36], [46], [37], [21], and attention-weighted fusion [44], [46], [8], [21].

TABLE I: **Quantitative comparisons of RGB-T segmentation on MF dataset [11] by input modality.** Previous RGB-T segmentation networks [36], [21] show a critical vulnerability to modality drop and over-reliance on a single modality, which can hinder the learning of complementary representations from multiple modalities.

Methods	RGB-T	RGB drop		THR drop	
	mIoU \uparrow	mIoU \uparrow	Diff \downarrow	mIoU \uparrow	Diff \downarrow
RTFNet	53.2	45.6	-7.6	10.5	-42.7
CMXNet	58.0	44.7	-13.3	39.2	-18.8
Ours	61.2	53.1	-8.1	52.7	-8.5

However, if the nature of multi-modal inputs is not considered in the network training, the networks easily become over-reliant on a single modality. This can hinder the network from learning complementary and meaningful representations for each modality, which is necessary to accurately and robustly segment objects.

B. RGB-Thermal Knowledge Distillation

Several studies [39], [17], [10] have investigated the potential of using knowledge distillation between RGB and thermal modalities to improve performance in various recognition applications. Specifically, Heatnet [39] utilizes knowledge distillation from daytime prediction results to nighttime to improve the performance of the RGB-T semantic segmentation network. MS-UDA [17] and CEKD [10] distill the knowledge of RGB-T segmentation network to thermal image segmentation network. In contrast to these previous works, this paper specifically focuses on knowledge distillation between clean and masked images for RGB-T semantic segmentation tasks.

III. METHOD

A. RGB-T Mask2Former

1) *Preliminaries for Mask Classification:* Mask classification architecture [6], [5] is a universal image segmentation network capable of semantic, instance, and panoptic segmentation. The network groups input pixels into N segments by estimating N binary masks and N class labels. The network consists of three main modules: a *backbone* that extracts low-resolution features from an image, a *pixel decoder* that gradually upsamples these features to generate high-resolution per-pixel embeddings, and a *transformer decoder* that estimates object queries based on the image features. The class prediction is estimated via MLP layers with the object queries. The binary mask predictions are obtained by decoding the per-pixel embeddings with object queries. Please refer to these papers [6], [5] for details.

2) *Mask2Former for RGB-T images:* We adopted Mask2Former [5] for semantic segmentation as our baseline model and modified the model to take RGB and thermal images, as shown in Fig. 3. More specifically, we assigned a modality-wise backbone for each modality. After extracting

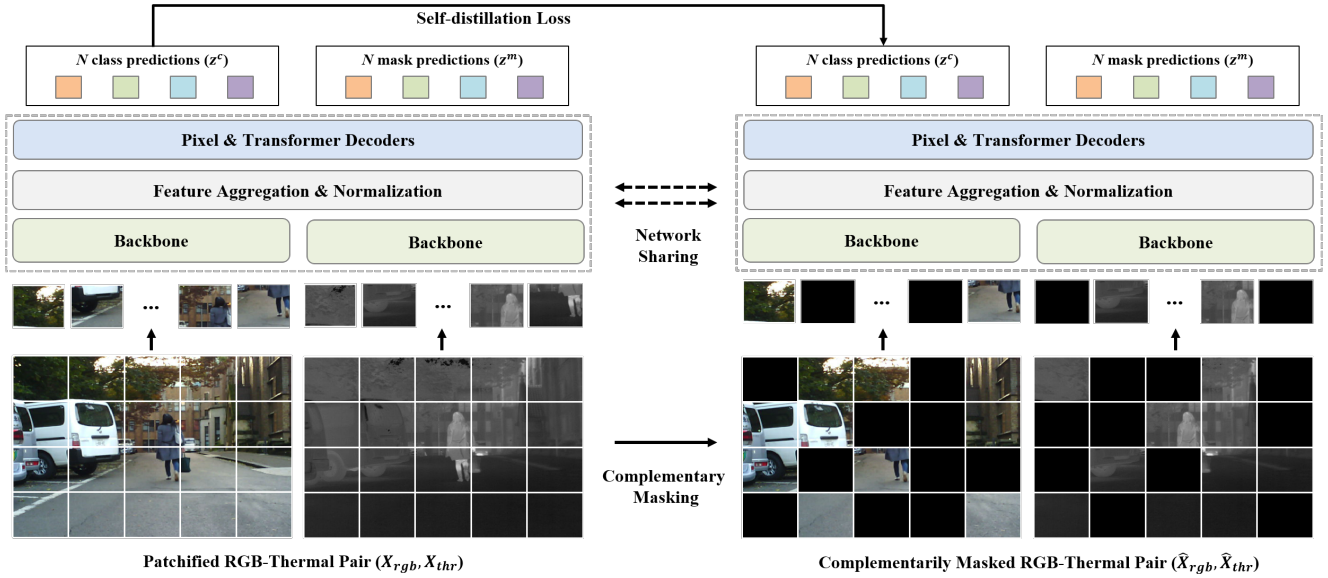


Fig. 3: **Overall pipeline of complementary masking and self-distillation for RGB-thermal semantic segmentation.** Our proposed training framework consists of complementary random masking and self-distillation loss. We randomly mask the patchified RGB-thermal pair in a complementary manner that guarantees at least one modality is valid. After that, the network estimates each prediction results from clean and masked RGB-thermal pairs. We enforce the network to predict the same class prediction results from the clean and masked RGB-thermal pairs. The proposed method resolves the over-reliant problem of RGB-T semantic segmentation networks and encourages the network to extract complementary and meaningful representations for robust and accurate semantic segmentation performance from RGB-T images.

modality-wise image features from RGB and thermal images, we aggregate the features with a simple winner-take-all strategy via *max* operation. This aggregation finds the most prominent feature from the RGB and thermal features across channel dimensions. We then normalized the aggregated feature map. At this stage, we can directly forward one modality feature to the decoder without aggregating the multi-modal features. After that, the aggregated multi-modal feature or single-modal feature is delivered to the pixel and transformer decoder to estimate N class and mask predictions. The final semantic segmentation mask can be obtained with a simple matrix multiplication of class and mask predictions.

B. Complementary Random Masking

Recently, masking strategy has been widely utilized in various language [9], [23], [2] and visual applications [1], [12], [41], [38], [43], [28] to learn meaningful representation. Especially, the image masking strategy is used to pre-train a large capacity backbone model to learn general representation for various downstream tasks, such as recognition [12], [41], video applications [38], 3D application [28]. Differing from these works focusing on learning general representation, we utilize a masking strategy to overcome the over-reliant problem of the RGB-T semantic segmentation task and to learn complementary and robust representation for each modality.

More specifically, as shown in Fig. 2 and Tab. I, common RGB-T segmentation networks easily rely on a single modality. Therefore, the network rarely extracts meaning-

ful representation from the other modality to segment and classify objects. This makes the network vulnerable to a wide range of fault cases, such as sensor disconnection, lens occlusion, and other input quality degeneration. Also, it loses the chance to learn complementary and useful representations for each modality to segment and classify objects. Therefore, we push the network in a situation where one modality is partially unavailable but allows the missing information can be complemented by the other modality. For this purpose, we propose a complementary random masking method for RGB-T semantic segmentation.

Complementary Patch Masking. We use Swin-Transformer [24] as our backbone model. Therefore, each modality image is patchified into a set of non-overlapped small patches, which will be fed into the transformer model for processing. Here, we randomly mask out the patches of one modality and replace the masked patches with learnable mask token vectors by following the convention of the token masking [9], [1], [43]. The other modality’s patches are masked out in the same manner by using the complementary mask. The complementary random masking process is defined as follows:

$$\begin{aligned}\hat{X}_{rgb} &= M * X_{rgb} + \hat{M} * L_{rgb}, \\ \hat{X}_{thr} &= \hat{M} * X_{thr} + M * L_{thr},\end{aligned}\quad (1)$$

where X_{input} is tokenized input image, M is random mask, \hat{M} is its complementary mask, defined as $1 - M$, and L_{input} is learnable token vector.

C. Self-distillation Loss

The proposed self-distillation loss consists of two losses L_{SDC} and L_{SDN} . The former loss L_{SDC} aims to make the network learn to extract complement representation when one modality information is partially unavailable. Therefore, we enforce the network to predict the same class prediction results from clean (*i.e.*, X_{rgb} and X_{thr}) and masked (*i.e.*, \hat{X}_{rgb} and \hat{X}_{thr}) RGB-thermal pairs. The proposed self-distillation loss for complementary representations is defined as follows:

$$L_{SDC} = L_1(z^c(X_{rgb}, X_{thr}), z^c(\hat{X}_{rgb}, \hat{X}_{thr})), \quad (2)$$

where $z^c(\cdot)$ is class prediction logit estimated from given tokenized inputs and $L_1(\cdot)$ is mean absolute error function.

The latter loss L_{SDN} aims to make the network extract robust representations from a partially masked single modality based on their non-local context rather than local features. For this purpose, we further enforce the class prediction consistency between the clean RGB-T pair and single masked modality (*i.e.*, \hat{X}_{rgb} or \hat{X}_{thr}). The proposed self-distillation loss for non-local representations is defined as follows:

$$L_{SDN} = L_1(z^c(X_{rgb}, X_{thr}), z^c(\hat{X}_{rgb})), \quad (3) \\ + L_1(z^c(X_{rgb}, X_{thr}), z^c(\hat{X}_{thr}))$$

D. Supervised Loss

We utilize the same supervised loss function L_{sup} , which consists of binary mask loss L_{mask} and classification loss L_{cls} , used in Mask2Former [5]. The supervised loss is defined as follows:

$$L_{sup} = L_{mask} + \lambda_{cls}L_{cls}, \quad (4)$$

where the mask loss L_{mask} is a combination of binary cross-entropy loss and dice loss [27], defined as $L_{mask} = \lambda_{ce}L_{ce} + \lambda_{dice}L_{dice}$.

Modality-wise Supervision. The current network architecture can estimate three types of prediction results according to the given input modalities (*i.e.*, RGB image, thermal image, and RGB-thermal pair). We also empirically found that the supervised loss for each prediction result of multiple input modalities can perform better than a single supervised loss for RGB-thermal pair. The modality-wise supervised loss is defined as follows:

$$L_{MWS} = L_{sup}(z_{gt}, z(X_{rgb}, X_{thr})), \quad (5) \\ + L_{sup}(z_{gt}, z(\hat{X}_{rgb})) + L_{sup}(z_{gt}, z(\hat{X}_{thr})),$$

where z_{gt} is ground truth class z^c and binary mask z^m , z is class and mask prediction from the given input modalities (RGB, thermal, or RGB-thermal). For the masked modalities, this loss only uses the first term (*i.e.*, RGB-T). The total loss is defined as follows:

$$L_{total} = L_{MWS} + L_{SDC} + L_{SDN} \quad (6)$$

IV. EXPERIMENTS

A. RGB-T Semantic Segmentation Datasets

In this study, we employ three publicly available RGB-T datasets to train and evaluate the proposed method.

MFNet dataset [11] consists of 820 daytime and 749 nighttime RGB-thermal images of urban driving scenes with a resolution of 640×480 . The dataset provides semantic labels for nine classes, including one unlabeled class and eight classes of common objects.

PST900 dataset [35] provides 894 RGB-thermal images with a resolution of 1280×720 , taken under the cave and subterranean environments for DARPA Subterranean Challenge. The dataset contains annotated segmentation labels for five classes, including one background class (*i.e.*, unlabeled) and four object classes.

KP dataset [14] is RGB-T paired urban driving scene dataset, providing 95K video frames (62.5K for daytime and 32.5K for nighttime) with a resolution of 640×512 . Originally, KAIST Multispectral Pedestrian Detection (KP) dataset provides detection bounding box labels only. But, thankfully, Kim *et al.* [17] provides annotated semantic segmentation labels for 503 daytime and 447 nighttime images. The labels include 19 object classes, which are the same classes as Cityscapes [7] dataset.

However, the dataset splits of Kim *et al.* [17] are undesirable to common RGB-T semantic segmentation network training. We divided 950 training images into 499 for training, 140 for validation, and 311 for testing. Daytime and nighttime images were appropriately distributed in each set. We provide our train/val/test splits that were used to train our network and other networks on KP dataset.

B. Implementation Details

Mask2Former. We employ the Swin transformr [24] (tiny, small, and base) as our backbone model. We use the multi-scale deformable attention Transformer (MSDeformAttn) [49] as a pixel decoder. We adopt the same Transformer decoder with DETR [3] for the Transformer decoder. The query N is set to 100 by default.

Training Settings. Our proposed method is implemented with PyTorch [29] and Detectron2 [42] library on a machine equipped with two NVIDIA RTX A6000 GPUs. The following training setting is used for all datasets. We use AdamW optimizer [25] and poly learning rate scheduler [4] with an initial learning rate of 10^{-4} . We train all segmentation networks with a batch size of 14 for 35K iterations. We utilize Swin transformer model [24] pre-trained on ImageNet-1K [30] (*i.e.*, tiny(T), small(S), and base(B)) as a backbone model. We apply random color jittering [22], random horizontal flipping, and random cropping to RGB and thermal images as data augmentation. For the coefficients of loss functions, we set λ_{cls} as 2.0 for predictions matched with a GT label and as 0.1 for the "no object" (*i.e.*, no match with any GT labels) by following [5]. Also, the coefficient λ_{ce} and λ_{dice} are set to 5.0

TABLE II: **Quantitative comparisons for semantic segmentation of RGB-T images on MF [11], PST900 [35], and KP [14] datasets.** We compared our proposed method with the previous state-of-the-art RGB-T semantic segmentation networks on MF, PST900, and KP benchmarks. Our proposed method demonstrates outperformed performance in all benchmark datasets. The best and second best performance in each block is highlighted in **bold** and underline, respectively.

(a) Quantitative comparison on MF day-night evaluation set [11] (9 classes)

Method	Unlabeled	Car	Person	Bike	Curve	Car Stop	Guardrail	Color Cone	Bump	mIoU
MFNet [11]	96.9	65.9	58.9	42.9	29.9	9.9	0.0	25.2	27.7	39.7
PSTNet [35]	97.0	76.8	52.6	55.3	29.6	25.1	<u>15.1</u>	39.4	45.0	48.4
RTFNet [36]	<u>98.5</u>	87.4	70.3	62.7	45.3	29.8	0.0	29.1	55.7	53.2
FuseSeg [37]	97.6	87.9	71.7	64.6	44.8	22.7	6.4	46.9	47.9	54.5
AFNet [44]	98.0	86.0	67.4	62.0	43.0	28.9	4.6	44.9	56.6	54.6
ABMDRNet [46]	98.6	84.8	69.6	60.3	45.1	33.1	5.1	47.4	50.0	54.8
FEANet [8]	98.3	87.8	71.1	61.1	46.5	22.1	6.6	<u>55.3</u>	48.9	55.3
GMNet [48]	97.5	86.5	73.1	61.7	44.0	42.3	14.5	<u>48.7</u>	47.4	57.3
CMX [21]	98.3	<u>90.1</u>	<u>75.2</u>	<u>64.5</u>	50.2	35.3	8.5	54.2	60.6	59.7
Ours (Swin-T)	98.2	90.0	73.1	63.7	47.9	40.7	9.9	54.4	54.2	59.1
Ours (Swin-S)	98.4	90.6	75.5	67.2	48.3	43.4	11.8	56.8	59.3	61.2
Ours (Swin-B)	98.2	90.0	75.1	<u>67.0</u>	45.2	49.7	18.4	54.2	54.4	61.4

(b) Quantitative comparison on PST900 evaluation set [35] (5 classes)

Method	Background	Fire-Extinguisher	Backpack	Hand-Drill	Survivor	mIoU
MFNet [11]	98.6	60.4	64.3	41.1	20.7	57.0
RTFNet-152 [36]	98.9	52.0	75.3	25.4	36.4	57.6
PSTNet [35]	98.9	70.1	69.2	53.6	50.0	68.4
ABMDRNet [46]	99.0	66.2	67.9	61.5	62.0	71.3
GMNet [48]	99.4	73.8	83.8	85.2	78.4	84.1
Ours (Swin-T)	99.5	<u>79.1</u>	86.0	86.2	78.7	85.9
Ours (Swin-S)	99.6	<u>76.2</u>	91.0	<u>87.2</u>	<u>80.8</u>	<u>86.9</u>
Ours (Swin-B)	99.6	79.5	<u>89.6</u>	89.0	82.2	88.0

(c) Quantitative comparison on KP day-night evaluation set [14], [17] (19 classes)

Method	Road	Sidewalk	Building	Wall	Fence	Pole	Traffic light	Traffic sign	Vegetation	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motorcycle	Bicycle	mIoU
MFNet [11]	93.5	23.6	75.1	0.0	0.1	9.1	0.0	0.0	69.3	0.2	90.4	24.0	0.0	69.6	0.3	0.3	0.0	0.0	0.6	24.0
RTFNet [36]	94.6	39.4	86.6	0.0	0.6	0.0	0.0	0.0	81.7	3.7	92.8	58.4	0.0	87.7	0.0	0.0	0.0	0.0	0.5	28.7
CMX [21]	97.7	53.8	90.2	0.0	47.1	46.2	10.9	45.1	87.2	34.3	93.5	74.5	0.0	91.6	0.0	59.7	0.0	46.1	0.2	46.2
Ours (Swin-T)	98.8	56.4	89.0	0.0	62.3	54.1	31.2	31.2	84.3	23.2	<u>94.4</u>	<u>83.6</u>	0.0	94.7	0.0	77.7	0.0	<u>51.4</u>	40.7	51.2
Ours (Swin-S)	98.8	<u>60.7</u>	92.1	0.0	60.4	55.1	31.1	53.2	<u>89.1</u>	27.7	95.0	81.4	17.7	<u>95.2</u>	1.1	83.3	0.0	49.9	<u>42.3</u>	<u>54.4</u>
Ours (Swin-B)	99.0	61.9	91.8	0.0	58.7	50.6	39.2	55.3	89.2	<u>23.2</u>	94.3	85.2	<u>2.9</u>	95.3	0.0	<u>80.5</u>	0.0	66.2	54.6	55.2

C. RGB-Thermal Semantic Segmentation

In this section, we compare our proposed method with the previous RGB-T semantic segmentation networks [11], [36], [35], [37], [48], [44], [46], [8], [21] on three benchmarks. We use mean Intersection-over-Union (mIoU) to evaluate the performance of semantic segmentation.

1) Evaluation on MFNet Day-Night Dataset [11]:

The quantitative and qualitative comparison results are shown in Tab. II-(a) and Fig. 4. We trained the RGB-T Mask2Former [5], as described in Sec. III, along with our proposed method. Also, we provide a variant of the network with Swin-T, Swin-S, and Swin-B backbone models. Compared to the previous state-of-the-art method (*i.e.*, CMXNet [21]), our approach leads to a 1.7% performance gain in the mIoU metric. Furthermore, our methods (Swin-S and B) achieve the best or second-best performance in most IoU metrics over nine classes.

2) Evaluation on PST900 Dataset [35]: For the PST900 benchmark, our model (*i.e.*, Swin-B) achieves a high performance improvement by 3.9% against previous state-of-the-art result (*i.e.*, GMNet [48]), as shown in Tab. II-(b).

3) Evaluation on KP Day-Night Dataset [14], [17]:

KP dataset has a more diverse and numerous number of classes compared to the MF [11] and PST900 [35] datasets. Therefore, the increased complexity in the dataset makes it more difficult to accurately segment the objects in the RGB-T images, requiring more advanced techniques and consideration of multi-modality inputs. We trained publicly available RGB-T semantic segmentation networks (*i.e.*, MFNet [11], RTFNet [36], CMXNet [21]) on the KP dataset with their provided code bases.

As shown in Tab. II-(c), our method (*i.e.*, Swin-B) achieves 9.0% performance improvement against CMXNet [21]. Also, Fig. 4 shows that our method shows precise and accurate segmentation quality in partially occluded, noisy, and cluttered environments. This implies that the proposed complementary random masking and self-distillation loss make the network learn to extract non-local and complementary representations from each modality, even in challenging conditions. We think the results show that as the complexity of semantic segmentation is higher, our proposed method is more helpful in achieving accurate and robust semantic segmentation performance from RGB and thermal images.




Fig. 4: **Qualitative comparison for semantic segmentation of RGB-T images on MFNet [11], PST900 [35], and KP [14] datasets.** The first two rows are qualitative comparisons of MFNet dataset, the next two rows are PST 900 dataset results, and the remaining rows are KP dataset results. The proposed method shows reliable and accurate segmentation results across all datasets, including day-light, low-light, noisy images, and harsh cave conditions. Further results can be found in the supplementary video.

D. Ablation Study


In this ablation study, we investigate the components of the proposed method, as shown in Tab. III-(a). Baseline indicates an RGB-T Mask2Former model that is modified to take RGB and thermal image inputs, as described in Sec. III. Our empirical finding demonstrates that modality-wise supervision loss L_{MWS} , which provides supervision for each prediction result from multiple input modalities, yields +0.8% performance gain compared to a single supervised loss for RGB-thermal pairs (*i.e.*, Baseline).

Also, applying complementary random masking CRM brings +0.4% performance improvement by pushing the network to segment and classify objects even when partially occluded inputs are provided. The self-distillation losses for complementary and non-local representation (L_{SDC} and L_{SDN}) are getting +0.3% and +0.2% improvement, respectively. L_{SDC} aims to make the network learn to extract

complement representation when one modality information is missing. L_{SDN} aims to make each modality extract robust representation to segment objects based on their non-local context rather than local features. Lastly, when all components are combined together, we get +1.7% performance improvement compared to the Baseline model.

1) *Complementary Random Masking*: We study various types of complementary masking strategies, as shown in Fig. 5 and Tab. III-(b). Square masking randomly masks a square area with half the height and width of the image in a random position. Patch masking randomly masks half an image (*i.e.*, 0.5 ratios) with patches of different sizes (*e.g.*, 8, 16, 32, 64).

Generally, complementary random masking shows better performance than the Baseline model. But, the patch size or masking scheme seems to be important for network performance. For example, the complementary random square



(a) RGB/THR images (b) Random square (c) Random patch (8) (d) Random patch (16) (e) Random patch (32) (f) Random patch (64)

Fig. 5: **Illustration of various complementary masking strategies.** Square masking randomly masks a square area with half the height and width of the image in a random position. Patch masking randomly masks half an image (*i.e.*, 0.5 ratios) with patches of different sizes (*e.g.*, 8, 16, 32, 64).

TABLE III: **Ablation study on MFNet dataset [11].** We conduct an ablation study of the proposed method and various masking strategies. Swin-S backbone is used for ablation study.

Method	L_{MWS}	CRM	L_{SDC}	L_{SDN}	mIoU
Baseline					59.5
(1)	✓				60.3
(2)	✓	✓			60.7
(3)	✓	✓	✓		61.0
(4)	✓	✓		✓	60.9
Ours	✓	✓	✓	✓	61.2

Method	Square	Patch (8)	Patch (16)	Patch (32)	Patch (64)
mIoU	59.7	59.8	60.8	61.2	61.1

masking may hinder learning non-local representation from L_{SDN} loss by masking out a wide area. Similarly, too tiny patch size is also undesirable to learn complementary representations for each modality. Generally, random patches over a certain size show higher performance. Empirically, we found that patch size 32 shows the best performance.

V. CONCLUSION

In this paper, we have proposed a complementary random masking strategy and self-distillation loss for robust and accurate RGB-Thermal semantic segmentation. The proposed masking strategy prevents over-reliance on a single modality. It also improves the accuracy and robustness of the neural network by forcing the network to segment and classify objects even when one modality is partially available. Also, the proposed self-distillation loss encourages the network to extract complementary and meaningful representations by enforcing class prediction consistency between clean and masked RGB-thermal pairs. Based on the proposed method, we achieve state-of-the-art performance over three RGB-T semantic segmentation benchmarks.

ACKNOWLEDGMENT

This work was supported by the Ministry of Trade, Industry and Energy (MOTIE) and Korea Institute of Advancement of Technology (KIAT) through the International Cooperative R&D program (P0019782). Also, this research was supported by a grant (P0026022) from R&D Program funded by Ministry of Trade, Industry and Energy of Korean government.

REFERENCES

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part 1 16*, pages 213–229. Springer, 2020.
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 40(4):834–848, 2017.
- [5] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022.
- [6] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. 2021.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 3213–3223, 2016.
- [8] Fuqin Deng, Hua Feng, Mingjian Liang, Hongmin Wang, Yong Yang, Yuan Gao, Junfeng Chen, Junjie Hu, Xiyue Guo, and Tin Lun Lam. Feanet: Feature-enhanced attention network for rgb-thermal real-time semantic segmentation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4467–4473. IEEE, 2021.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [10] Zhen Feng, Yanning Guo, and Yuxiang Sun. Cekd: Cross-modal edge-privileged knowledge distillation for semantic scene understanding using only thermal images. *IEEE Robotics and Automation Letters*, 2023.

- [11] Qishen Ha, Kohei Watanabe, Takumi Karasawa, Yoshitaka Ushiku, and Tatsuya Harada. Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes. In *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS)*, pages 5108–5115, 2017.
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 2961–2969, 2017.
- [14] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 1037–1045, 2015.
- [15] Bin Kang, Dong Liang, Junxi Mei, Xiaoyang Tan, Quan Zhou, and Dengyin Zhang. Robust rgb-t tracking via graph attention-based bilinear pooling. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [16] Shehryar Khattak, Christos Papachristos, and Kostas Alexis. Keyframe-based thermal–inertial odometry. *Journal of Field Robotics*, 37(4):552–579, 2020.
- [17] Yeong-Hyeon Kim, Ukcheol Shin, Jinsun Park, and In So Kweon. Ms-uda: Multi-spectral unsupervised domain adaptation for thermal image semantic segmentation. *IEEE Robotics and Automation Letters*, 6(4):6497–6504, 2021.
- [18] Dong-Guw Lee, Myung-Hwan Jeon, Younggun Cho, and Ayoung Kim. Edge-guided multi-domain rgb-to-tir image translation for training vision tasks with challenging labels. *arXiv preprint arXiv:2301.12689*, 2023.
- [19] Dong-Guw Lee, Kyu-Seob Song, Young-Hoon Nho, Ayoung Kim, and Dong-Soo Kwon. Sequential thermal image-based adult and baby detection robust to thermal residual heat marks. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13120–13127. IEEE, 2022.
- [20] Chenglong Li, Xinyan Liang, Yijuan Lu, Nan Zhao, and Jin Tang. Rgb-t object tracking: Benchmark and baseline. *Pattern Recognition*, 96:106977, 2019.
- [21] Huayao Liu, Jiaming Zhang, Kailun Yang, Xinxin Hu, and Rainer Stiefelwagen. Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. *arXiv preprint arXiv:2203.04838*, 2022.
- [22] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016.
- [23] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [25] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- [26] Yawen Lu and Guoyu Lu. Superthermal: Matching thermal as visible through thermal feature exploration. *IEEE Robotics and Automation Letters*, 6(2):2690–2697, 2021.
- [27] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016.
- [28] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 604–621. Springer, 2022.
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*, pages 8026–8037, 2019.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sateesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [31] Ukcheol Shin, Kyunghyun Lee, Byeong-Uk Lee, and In So Kweon. Maximizing self-supervision from thermal image for effective self-supervised learning of depth and ego-motion. *IEEE Robotics and Automation Letters*, 7(3):7771–7778, 2022.
- [32] Ukcheol Shin, Kyunghyun Lee, Seokju Lee, and In So Kweon. Self-supervised depth and ego-motion estimation for monocular thermal video using multi-spectral consistency loss. *IEEE Robotics and Automation Letters*, 7(2):1103–1110, 2021.
- [33] Ukcheol Shin, Kwanyong Park, Byeong-Uk Lee, Kyunghyun Lee, and In So Kweon. Self-supervised monocular depth estimation from thermal images via adversarial multi-spectral adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5798–5807, 2023.
- [34] Young-Sik Shin and Ayoung Kim. Sparse depth enhanced direct thermal-infrared slam beyond the visible spectrum. *IEEE Robotics and Automation Letters*, 4(3):2918–2925, 2019.
- [35] Shreyas S Shivakumar, Neil Rodrigues, Alex Zhou, Ian D Miller, Vijay Kumar, and Camillo J Taylor. Pst900: Rgb-thermal calibration, dataset and segmentation network. In *IEEE Int'l Conf. on Robotics and Automation (ICRA)*, pages 9441–9447, 2020.
- [36] Yuxiang Sun, Weixun Zuo, and Ming Liu. Rtnet: Rgb-thermal fusion network for semantic segmentation of urban scenes. *IEEE Robotics and Automation Letters (RAL)*, 4(3):2576–2583, 2019.
- [37] Yuxiang Sun, Weixun Zuo, Peng Yun, Hengli Wang, and Ming Liu. Fuseseg: Semantic segmentation of urban scenes based on rgb and thermal data fusion. *IEEE Trans. on Automation Science and Engineering (TASE)*, 2020.
- [38] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022.
- [39] Johan Vertens, Jannik Zürn, and Wolfram Burgard. Heatnet: Bridging the day-night domain gap in semantic segmentation with thermal images. In *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [40] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Minghui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 2020.
- [41] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022.
- [42] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [43] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022.
- [44] Jiangtao Xu, Kaige Lu, and Han Wang. Attention fusion network for multi-spectral semantic segmentation. *Pattern Recognition Letters*, 146:179–184, 2021.
- [45] Xiaoxiao Yang, Yejiang Qian, Huijie Zhu, Chunxiang Wang, and Ming Yang. Baanet: Learning bi-directional adaptive attention gates for multispectral pedestrian detection. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2920–2926. IEEE, 2022.
- [46] Qiang Zhang, Shenlu Zhao, Yongjiang Luo, Dingwen Zhang, Nianchang Huang, and Jungong Han. Abmdnet: Adaptive-weighted bi-directional modality difference reduction network for rgb-t semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2633–2642, 2021.
- [47] Shenlu Zhao, Yichen Liu, Qiang Jiao, Qiang Zhang, and Jungong Han. Mitigating modality discrepancies for rgb-t semantic segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [48] Wujie Zhou, Jinfu Liu, Jingsheng Lei, Lu Yu, and Jenq-Neng Hwang. Gmnet: graded-feature multilabel-learning network for rgb-thermal urban scene semantic segmentation. *IEEE Transactions on Image Processing*, 30:7790–7802, 2021.
- [49] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.