

MULTIMODAL TRANSFORMER USING CROSS-CHANNEL ATTENTION FOR OBJECT DETECTION IN REMOTE SENSING IMAGES

Bissmella Bahaduri, Zuheng Ming, Fangchen Feng, Anissa Mokraoui

L2TI Laboratory, University Sorbonne Paris Nord, Villetaneuse, France

ABSTRACT

Object detection in Remote Sensing Images (RSI) is a critical task for numerous applications in Earth Observation (EO). Differing from object detection in natural images, object detection in remote sensing images faces challenges of scarcity of annotated data and the presence of small objects represented by only a few pixels. Multi-modal fusion has been determined to enhance the accuracy by fusing data from multiple modalities such as RGB, infrared (IR), lidar, and synthetic aperture radar (SAR). To this end, the fusion of representations at the mid or late stage, produced by parallel subnetworks, is dominant, with the disadvantages of increasing computational complexity in the order of the number of modalities and the creation of additional engineering obstacles. Using the cross-attention mechanism, we propose a novel multi-modal fusion strategy for mapping relationships between different channels at the early stage, enabling the construction of a coherent input by aligning the different modalities. By addressing fusion in the early stage, as opposed to mid or late-stage methods, our method achieves competitive and even superior performance compared to existing techniques. Additionally, we enhance the SWIN transformer by integrating convolution layers into the feed-forward of non-shifting blocks. This augmentation strengthens the model's capacity to merge separated windows through local attention, thereby improving small object detection. Extensive experiments prove the effectiveness of the proposed multimodal fusion module and the architecture, demonstrating their applicability to object detection in multimodal aerial imagery. Our code is available at [here](#).

Index Terms— Multimodal fusion, cross-channel attention, convolutional shifting window, object detection, remote sensing imagery

1. INTRODUCTION

Object detection in Remote Sensing Images (RSI) including aerial images is a critical task enabling the identification and localization of objects within satellite or aerial imagery. It has numerous applications for Earth Observation (EO) such as environmental monitoring, climate change, urban planning, and military surveillance [1]. Given the availability of different

sensors onboard satellites and UAVs, multimodal fusion has been an active area of interest for research. Given the smaller object size and limitation of annotated data availability [2], multimodal data has been used to improve the performance of diverse computer vision tasks in RSI including object detection. Previous work has shown that using neural networks for fusing multiple modalities has remarkable benefits for tasks such as semantic segmentation [3], video description [4], and action recognition [5].

Multimodal fusion methods can be classified according to the stage at which the fusion is performed: early-stage fusion, mid-stage fusion, and late-stage fusion [6]. Early-stage fusion methods typically employ simple techniques like concatenation or averaging to integrate multimodal data at the early stage of the model. Although early-stage fusions are usually simple and easy to use, their performance lags behind more sophisticated mid- or late-stage methods. In mid-stage and late-stage fusion, modality features are extracted by independent feature extractors sub-networks, often referred to as streams in the literature, before being fused. The key distinction between mid-stage and late-stage fusion lies in their subsequent processing steps: in mid-stage fusion, the fused features are further processed by a main network, whereas in late-stage fusion, the feature extractor streams constitute the entirety of the network.

Depending on the fusion technique employed, multimodal fusion approaches can also be categorized into aggregation-based, alignment-based, and attention-based methods. Aggregation-based methods [7] rely on straightforward operations like concatenation and averaging to fuse modalities, a process that can occur at any stage—early, mid, or late. Alignment-based methods [5] leverage regularization loss to align embeddings from different modalities, typically executed at the late stage of fusion. The more recent attention-based methods [8, 4] utilize attention mechanisms for fusion, commonly occurring at mid or late stages. A detailed review of some related attention-based fusion methods will be presented in the next section.

Despite the achievements and good results of these methods, several challenges persist. While aggregation-based methods offer simplicity, they often overlook potential modal misalignments and intra-modality dependencies. On the other hand, alignment-based methods, while aimed at minimizing

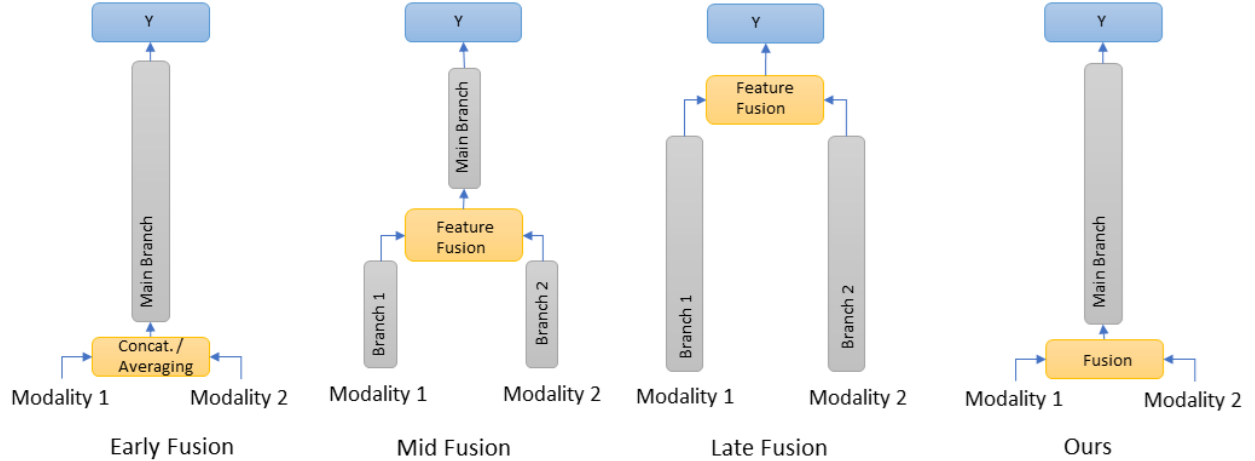


Fig. 1: An illustration of different fusion strategies. In contrast to mid-stage and late-stage fusions which involve fusion at the feature level extracted by distinct parallel subnetworks, our approach targets fusion at the early stage.

alignment loss, may not adequately address intramodal information exchanges. Moreover, regardless of the fusion approach, the use of subnetworks for feature extraction from different modalities remains widespread. These subnetworks are trained to extract modal features independently, with fusion occurring in subsequent steps, as shown in Fig 1. This family of methods increases computational complexity and adds additional technical hurdles, such as determining the optimal step and scale at which fusion is performed. This challenge is particularly pronounced when the independent parallel subnetworks are composed of recently developed Vision Transformers (ViT) [9], where computational costs increase quadratically. In addressing this issue, some works choose CNN for modality-specific streams and propose to use transformers for the multimodal fusion module [8]. This integration of CNNs and ViTs also occurs for object detection where existing fusion methods are mainly CNN-based [2, 8].

In this paper, we consider object detection in remote sensing images with the presence of two modalities, RGB and infrared (IR). We propose an early-stage multimodal cross-attention fusion method based on the cross-attention mechanism (Fig. 2). We show that different modalities can be fused at the early stage without requiring additional independent subnetworks for each modality. The key concept behind the proposed method involves processing the three RGB channels individually instead of considering them as a single modality. Notably, the absence of modality-specific subnetworks enables scalability to accommodate any number of modalities without a substantial increase in model size. Additionally, this strategy proves effective in eliminating the need for additional subnetworks dedicated to modality feature extraction. We demonstrate that the cross-channel attention fusion can outperform or match the performance of state-of-the-art mid-stage fusion based on the CNN backbone.

Additionally, we propose to use SWIN [10] as the main

backbone which uses the shifted window mechanism to enable the merging of neighboring patches separated by local window attention. Given that this shifting mechanism is only operative in half of the blocks, we propose augmenting the Feed-Forward network (FFN) with a convolutional layer while maintaining the same embedding dimension. This improvement aims to strengthen the network’s ability to capture local information and facilitate the integration of neighboring patches in different windows. Hereafter, we refer to these blocks as convolutional-shifting blocks. Furthermore, we show the necessity of this backbone improvement for the proposed fusion strategy and we believe that this is due to the rich information present in the fusion module output.

To summarize, the main contributions of this work are:

- 1) We introduce a new cross-channel attention module that allows for the early alignment of different modalities.
- 2) We propose a convolutional-shifting window that incorporates convolutional layers in FFN to assist in merging separated windows in local attention enhancing the detection performance.
- 3) The extensive experiments demonstrate the superiority of the proposed approach, highlighting its applicability for object detection in multimodal aerial imagery.

2. RELATED WORK

Attention based multimodal fusion: Recently attention-based multimodal fusion has caught attention. In contrast to aggregation-based approaches, which rely on simple concatenation, attention-based fusion methods employ self and cross-attention mechanisms to fuse modality features at the mid or late-stage. Cross-modal attention is used to enrich the main modality with lidar information in [8], while [11] employs attention at various stages between two separate branches for multi-modal fusion. Additionally, attention is

utilized in [4] for the mid-stage fusion of image and audio modalities through the introduction of extra bottleneck tokens. Similarly, [12] introduces additional learnable fusion tokens and aggregates modalities for token fusion using Bi-LSTM.

Compared to aggregation-based methods, existing attention-based approaches replace concatenation with more sophisticated feature-level attention mechanisms. However, they still require feature extractors for each modality. The difference between our proposed cross-channel attention fusion and existing attention-based methods is that, in our proposed method, fusion occurs before any substantial processing and featurization.

Multimodal fusion for object detection: Over the past decade, deep learning methods have been widely explored for multimodal data fusion to improve object detection performance. [7] is an aggregation-based method and its performance is affected by the scale and stage at which aggregation occurs, which is a common problem for mid-stage fusion methods. [13] is an alignment-based method that aligns multimodal features based on similarity regulation. One of the limitations of this method is that it ignores modal-specific information.

In terms of architectural design, a notable commonality among these existing methods is their reliance on subnetwork branches for feature extraction. Indeed, [14] explored early, middle, and late-stage fusion using CNNs for object detection and claimed that mid-stage fusion yielded the optimal performance. The following works all followed this principle. [15] uses the transformer self-attention mechanism for fusion at multiple stages between two modality branches. Manish et al. [2] proposed a real-time framework for object detection in multi-modal remote sensing imagery by performing a mid-stage fusion of RGB and IR modalities. More recently, Zhang et al. [16] proposed the CNN-based SuperYOLO and studied various fusion strategies, including early-stage concatenation. It’s worth noting that their proposed multimodal fusion also relies on subnetworks comprising squeeze and excitation blocks [17] maintaining consistent input sizes. Different from these methods, we propose cross-channel attention fusion without the need for separate subnetworks. Furthermore, whereas the aforementioned methods consider RGB as a unified modality, we opt to treat each color channel independently. This approach enables us to leverage not only the cross-information between two modalities (RGB and IR) but also among all three RGB channels.

Window-based Vision Transformers/CNN. The landscape of computer vision has undergone a substantial transformation with the emergence of ViT [9], showcasing advancements across a broad spectrum of visual tasks. However, while ViTs exhibit commendable performance, they come with certain limitations, notably a substantial computational burden and lack of sense of locality. To address these limitations, [10] enhanced the vanilla ViTs by introducing hierarchical

architectures and localized windows. These enhancements have found practical applications in single-modal aerial image object detection [18]. Drawing from this body of work, more recent approaches endeavor to combine the strengths of CNNs with ViTs [19]. This combination capitalizes on the respective advantages of both CNNs and ViTs, thus offering promising prospects for computer vision applications such as classification [19], and Face Presentation Attack Detection (PAD) [20]. In this paper, we utilize SWIN architecture as the main backbone, with an enhancement on joining the separated windows by incorporating a convolutional layer in the feed-forward network of non-shifting blocks.

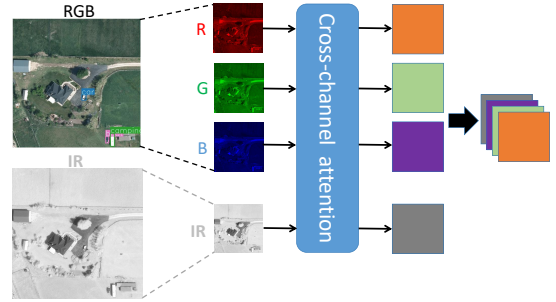


Fig. 2: Combining multimodal inputs using cross-channel attention instead of simple channel-wise concatenation.

3. PROPOSED METHOD

3.1. Overall architecture

Illustrated in Fig 3, the proposed architecture is composed of three components: the cross-channel attention fusion module, a feature extraction backbone featuring SWIN-like blocks enhanced with our proposed convolutional-shifting window mechanism, and a YOLO-based detection head as employed in [16]. The choice of this detection head is motivated by the simplicity inherent in the one-step detection method, widely adopted in the field, facilitating comparisons.

3.2. Cross-channel attention for multimodal fusion

We present our proposed cross-channel attention fusion module in Fig 4. Considering $X_{\text{RGB}} \in \mathbb{R}^{H \times W \times 3}$, and $X_{\text{IR}} \in \mathbb{R}^{H \times W}$ as RGB and IR modality image inputs respectively, following the standard practice in ViT methodologies, each channel (R, G, B, and IR) undergoes tokenization and patch splitting. The resulting patched channels can be denoted respectively as $X_r, X_g, X_b, X_{ir} \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s} \times d}$ where s represents the patch size and d denotes the embedding dimension. Subsequently, the obtained X_r, X_g, X_b, X_{ir} are fused based on the proposed cross-channel attention fusion module which relies on the cross-attention mechanism. Mathematically, the attention mechanism can be generally formulated as in equa-




Fig. 3: (a) The overall architecture based on Swin-like backbone for multi-modal object detection in RSI; (b) Convolutional-shifting window module.

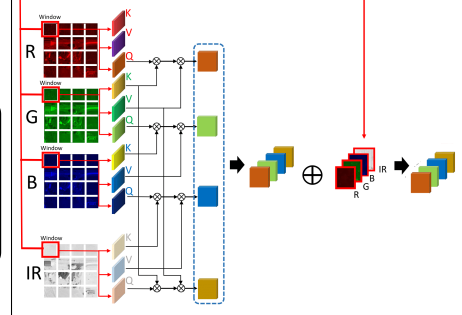


Fig. 4: Cross-channel attention module for RGB and IR images multimodal fusion.

tion (1):

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{Q \cdot K^T}{\sqrt{d}} \right) \cdot V \quad (1)$$

where Q, K, V are the query, key, and value matrices and d is the embedding dimension. Cross-attention and self-attention represent two applications of the attention mechanism. Cross-attention involves the asymmetric combination of two embedding sequences, where one sequence acts as a query input, while the other serves as key and value inputs, thereby capturing the interdependence between the two sequences. We demonstrate the effectiveness of cross-attention for fusion by contrasting it with a fusion technique based on self-attention. Detailed comparisons can be found in the experimental section.

In our fusion process, we propose to apply cross-attention individually to each channel (R, G, B, IR). Specifically, each channel acts as a query, and one of the remaining channels is chosen to fulfill both key and value roles in the cross-attention mechanism. We select the complementary channel to minimize its spectral proximity to the query channel. The formed four pairs are (R, G), (G, B), (B, IR), and (IR, G). The fused output comprises the concatenation of all these fused channels along with a residual branch. To streamline computational complexity, we omit the feed-forward network from the cross-channel attention, as its role in altering the channel dimension is redundant within this context.

We highlight that, in our cross-channel attention framework, we treat each channel as an independent modality, thus incorporating the relationship between channels alongside information exchange with the IR modality. This decomposition of RGB enables distinct fusion mechanisms for each RGB channel and the IR modality, recognizing potential variations in their relationships. To illustrate this, we also tested a cross-attention-based fusion technique treating RGB as a single modality. Detailed comparisons are provided in the experimental section. Moreover, we deliberately align the number of output channels of the fusion module with the input channels. This choice ensures direct or indirect linkage among the

four channels. Additionally, it substantially reduces computational complexity.

3.3. Convolutional-shifting window-based backbone

We base our backbone on SWIN which adopts a hierarchical approach with multiple stages. SWIN divides the input image into non-overlapping patches at multiple scales, enabling more efficient processing of both local and global information. The hierarchical transformer architecture allows SWIN to capture long-range dependencies while maintaining computational efficiency. In our SWIN-based backbone, three stages process the image at different scales starting from a higher resolution scale to low resolution in the third stage. Since the objects in RSI are often small and densely packed into a few pixels, we choose a higher number of blocks in the initial stage where the resolution remains high, while progressively reducing the number of blocks in later stages decreasing the resolution by a factor of 2. This backbone enables us to learn hierarchical multi-resolution features in a fine-to-coarse manner with a higher focus in the first stage for better detection of small objects (see Fig 3 (a)).

In addition, a well-recognized limitation of window-based Vision Transformers is their segregation of neighboring patches across different windows. To address this challenge, the SWIN Transformer introduces a shifting window mechanism, albeit restricted to only half of its blocks in a way that each nonshifting block is followed by a shifting window block. In our approach, we seek to enhance connectivity across all blocks and imbue the architecture with a heightened sense of locality. To achieve this, we introduce an extra convolutional layer with a kernel size of two by two positioned between two Fully Connected (FC) layers within the FFN while keeping the embedding dimension fixed (the orange block as shown in Fig 3 (b)). The use of a two-by-two kernel size and a fixed dimension is intended to prevent an increase in the computational complexity of the model. This augmentation not only promotes greater coherence but also enhances the network's perception of spatial proximity.

The enhanced blocks can be utilized across all three stages

of the process. Through empirical investigation, we have determined that the optimal performance is attained when the enhanced blocks are deployed in stages 1 and 2. A comparison of various combinations is provided in the experimental section for further insight.

4. EXPERIMENTS

4.1. Experimental setup

We apply our proposed cross-channel attention to the task of object detection in aerial images. We experiment with the VEDAI dataset [25]. The dataset contains scenes captured from the same altitude with a resolution of $12.5\text{cm} \times 12.5\text{cm}$ per pixel. Each image has 4 channels of RGB and IR. The dataset consists of 1246 images having diverse backgrounds containing highways, pastures, mountains, and urban areas. The images come in both 1024×1024 and 512×512 sizes and we have used the 512×512 images for all the experiments. The task is to identify and localize 11 types of vehicles such as car, pickup, truck, and camping. The dataset is divided into 10 folds for cross-validation evaluation. Following the same protocol as in SuperYOLO [16], we use the first folder for the ablation studies, and all 10 folders for overall evaluation compared with the state-of-the-art methods. The data is augmented with Hue Saturation Value (HSV), multi-scale, translation, left-right flip, and mosaic augmentation methods similar to SuperYOLO. We considered eight classes in the dataset and ignored classes that have under 50 instances in the dataset.

We have used the standard Stochastic Gradient Descent (SGD) [26] to optimize the network with a momentum of 0.937 and weight decay of 0.0005. The models were trained for 300 epochs using one Nvidia A100 GPU. We used the standard detection loss which combines localization, classification, and confidence losses to train the model, and the performances are evaluated using mAP_{50} , i.e., detection metric of mean Average Precision at IOU (Intersection over Union) = 0.5 for all categories.

4.2. Cross-channel attention fusion module

We evaluate the performance of the proposed Cross-Channel (CC) attention fusion module. To illustrate the effectiveness of the cross-channel attention fusion, we designed two fusion techniques based on vanilla self-attention and vanilla cross-attention. For the vanilla self-attention fusion, after the standard ViT patch encoding with separate projections, the obtained embeddings of each modality are concatenated before using a standard self-attention mechanism to perform the fusion. The difference between this configuration and the proposed cross-channel attention is that the self-attention is used instead of the cross-attention and the RGB channels are treated jointly. For the vanilla cross-attention, the cross-attention is used with two modalities (RGB and IR)

as separate inputs. Notably, this differs from the proposed cross-channel attention fusion, where attention is applied at the channel level within each modality. All the aforementioned attention mechanisms are employed in a multi-head way.

We tested the proposed fusion module on two types of backbone. We chose the CNN-based SuperYOLO’s backbone and our enhanced SWIN-like backbone which is ViT-based. For the CNN-based backbone, we compare the proposed fusion method with pixel-level fusion and feature-level fusion, both studied in SuperYOLO. For the ViT-based backbone, we compare the proposed fusion method with the early-stage concatenation method, vanilla self-attention fusion, and vanilla cross-attention fusion. Performance with a single modality is also presented as a reference. The results are shown in Table 2.

For the CNN-based backbone, CC-attention fusion outperforms the Pixel-level fusion and the Multimodal Feature-level (MF) fusion used in SuperYOLO. For the ViT-based backbone, the CC-attention fusion outperforms the RGB-IR concatenation by 3.3% and improves by respectively 15% and 8% compared to the scenario where only IR or RGB images are used. As expected, the proposed cross-channel attention fusion method outperforms the two tested attention-based fusion strategies. It’s also interesting to notice that the vanilla cross-attention method outperforms the vanilla self-attention method. These observations validate the advantages of cross-attention and confirm the benefits of treating RGB channels individually.

4.3. Convolutional-shifting window

We demonstrate the effectiveness of integrating the enhanced convolutional-shifting window across three stages of the backbone. We evaluate three scenarios to assess the impact. Table 3 shows the results. We can see that when adding a convolution layer inside the FFN in non-shifting blocks at stage 1, the model outperforms the FFN without convolution by 4.5% in terms of the mAP_{50} . Furthermore, introducing convolution at stages 1 and 2 gains 5.2% improvement.

4.4. Overall performance

The overall performance of our proposed approach is illustrated in Table 1. we can see that our proposed method achieves competitive results and outperforms the state-of-the-art CNN model by 0.3%. Additionally, our method outperforms SuperYOLO in detecting difficult classes with the least number of instances in the training set, namely the Boat, Van, and Other classes. We also illustrate in Fig 5 a visual comparison of our method and SuperYOLO for two different scenes where only our method has successfully detected and correctly classified the objects.

Method	Car	Pickup	Camping	Truck	Other	Tractor	Boat	Van	mAP ₅₀
YOLOv3 [21]	84.57	72.68	67.13	61.96	43.04	65.24	37.10	58.29	61.26
YOLOv4 [22]	85.46	72.84	72.38	62.82	48.94	68.99	34.28	54.66	62.55
YOLOv5 [23]	84.33	72.95	70.09	61.15	49.94	67.35	38.71	56.65	62.65
YOLOrs [2]	84.15	78.27	68.81	52.60	46.75	67.88	21.47	57.91	59.73
YOLO-Fine [24]	79.68	74.49	77.09	80.97	37.33	70.65	60.84	63.56	68.83
SuperYOLO [16]	89.30	81.48	79.22	67.27	54.29	78.88	55.95	71.41	72.22
Ours	89.13	82.70	76.38	61.57	56.32	77.94	60.36	75.84	72.53

Table 1: Class-wise mean Average Precision mAP_{50} for our proposed method comparing to the state-of-art on VEDAI Dataset.

Architecture	Fusion Method	mAP ₅₀
CNN-based (SuperYOLO)	Pixel fusion	76.90
	MF fusion	77.73
	CC attention	77.9
ViT-based (Ours)	IR	63.79
	RGB	70.55
	RGB-IR concatenation	74.23
	RGB-IR Vanilla self-attention	74.03
	RGB-IR Vanilla cross-attention	75.64
	RGB-IR CC attention	78.53

Table 2: The comparison of the proposed multi-channel attention fusion module with other fusion methods with CNN-based and ViT-based backbones on the VEDAI dataset (Fold-1).

Backbone	mAP ₅₀
Simple FFN	71.23
FFN with Conv2d in stage 1	75.75
FFN with Conv2d in stages 1,2	76.52

Table 3: Effect of the convolution in FFN at the 1st and the 2nd stages on the VEDAI dataset (Fold-1).

5. CONCLUSIONS

This paper introduces a new cross-channel attention multi-modal fusion module that allows for aligning different modalities by learning the relationship between different channels at the early stage. The achieved results demonstrate the module’s competitiveness compared to the mid-stage CNN-based fusion methods. Moreover, we propose the convolutional-shifting window, which incorporates convolutional layers in a feedforward network of non-shifting blocks in SWIN. This enhancement facilitates the joining of separated local windows within the SWIN framework, thus augmenting the model’s capability to localize and detect small objects. The experiments provide evidence of the competitive and better performance of the proposed approach compared to the mid-



Fig. 5: Visual results using our method and SuperYOLO.

stage fusion technique in the context of object detection using multimodal aerial imagery. The results also emphasize the utility of the convolutional-shifting window in enhancing the model’s spatial awareness and object detection capabilities.

For future endeavors, exploring the algorithm’s applicability across different modalities presents an intriguing avenue. Additionally, investigating pretraining strategies such as self-supervised learning could further enhance the robustness and performance of the method.

6. ACKNOWLEDGEMENTS

This work was supported in part by the Laboratoire Commun Intelligence, Reconnaissance et Surveillance Réactive (IRISER) LabCom IRISER Project.

7. REFERENCES

- [1] Muhammad Ahmad et al., “Hyperspectral image classification—traditional to deep models: A survey for future prospects,” *IEEE journal of selected topics in applied earth observations and remote sensing*, vol. 15, pp. 968–999, 2021.
- [2] Manish Sharma et al., “Yolors: Object detection in multimodal remote sensing imagery,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 1497–1508, 2020.
- [3] Seong-Jin Park, Ki-Sang Hong, et al., “Rdfnet: Rgb-d multi-level residual feature fusion for indoor semantic segmentation,” in *Proceedings of the IEEE ICCV*, 2017, pp. 4980–4989.
- [4] Arsha Nagrani, Shan Yang, et al., “Attention bottlenecks for multimodal fusion,” *NeurIPS*, vol. 34, pp. 14200–14213, 2021.
- [5] Sijie Song, Jiaying Liu, et al., “Modality compensation network: Cross-modal adaptation for action recognition,” *IEEE Transactions on Image Processing*, vol. 29, pp. 3957–3969, 2020.
- [6] Di Feng, Christian Haase-Schütz, et al., “Deep multimodal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 3, pp. 1341–1360, 2020.
- [7] Caner Hazirbas, Lingni Ma, et al., “Fusenet: Incorporating depth into semantic segmentation via fusion-based cnn architecture,” in *Computer Vision—ACCV 2016*. Springer, 2017, pp. 213–228.
- [8] Satyam Mohla, Shivam Pande, et al., “Fusatnet: Dual attention based spectrospatial multimodal fusion network for hyperspectral and lidar classification,” in *Proceedings of the IEEE/CVF CVPR Workshops*, 2020, pp. 92–93.
- [9] Alexey Dosovitskiy et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [10] Ze Liu et al., “Swin transformer: Hierarchical vision transformer using shifted windows,” in *Proceedings of the IEEE/CVF ICCV*, 2021, pp. 10012–10022.
- [11] Peng Sun, Wenhui Zhang, et al., “Deep rgb-d saliency detection with depth-sensitive attention and automatic multi-modal fusion,” in *Proceedings of the IEEE/CVF CVPR*, 2021, pp. 1407–1417.
- [12] Yuxing Chen, Maofan Zhao, et al., “Incomplete multimodal learning for remote sensing data fusion,” *arXiv preprint arXiv:2304.11381*, 2023.
- [13] Konstantinos Bousmalis, George Trigeorgis, et al., “Domain separation networks,” *NeurIPS*, vol. 29, 2016.
- [14] J Liu, S Zhang, , et al., “Multispectral deep neural networks for pedestrian detection. arxiv 2016,” *arXiv preprint arXiv:1611.02644*.
- [15] Qingyun Fang et al., “Cross-modality fusion transformer for multispectral object detection,” *Available at SSRN 4227745*, 2022.
- [16] Jiaqing Zhang et al., “Superyolo: Super resolution assisted object detection in multimodal remote sensing imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–15, 2023.
- [17] Jie Hu et al., “Squeeze-and-excitation networks,” in *Proceedings of the IEEE CVPR*, 2018, pp. 7132–7141.
- [18] Xuan Cao et al., “Swin-transformer-based yolov5 for small-object detection in remote sensing images,” *Sensors*, vol. 23, no. 7, pp. 3634, 2023.
- [19] Haiping Wu et al., “Cvt: Introducing convolutions to vision transformers,” in *Proceedings of the IEEE/CVF ICCV*, 2021, pp. 22–31.
- [20] Zuheng Ming et al., “Vitrapad: video transformer using convolution and self-attention for face presentation attack detection,” in *2022 IEEE ICIP*. IEEE, 2022, pp. 4248–4252.
- [21] Joseph Redmon et al., “Yolov3: An incremental improvement,” *arXiv preprint arXiv:1804.02767*, 2018.
- [22] Alexey Bochkovskiy, Chien-Yao Wang, et al., “Yolov4: Optimal speed and accuracy of object detection,” *arXiv preprint arXiv:2004.10934*, 2020.
- [23] G. J. et al, “ultralytics/yolov5: v5.0,” 2021.
- [24] Minh-Tan Pham, Luc Courtrai, et al., “Yolo-fine: One-stage detector of small objects under various backgrounds in remote sensing images,” *Remote Sensing*, vol. 12, no. 15, pp. 2501, 2020.
- [25] Sebastien Razakarivony et al., “Vehicle detection in aerial imagery: A small target detection benchmark,” *Journal of Visual Communication and Image Representation*, vol. 34, pp. 187–203, 2016.
- [26] Herbert Robbins and Sutton Monro, “A stochastic approximation method,” *The annals of mathematical statistics*, pp. 400–407, 1951.