



YOLO-FireAD: Efficient Fire Detection via Attention-Guided Inverted Residual Learning and Dual-Pooling Feature Preservation

Weichao Pan¹, Bohan Xu², Xu Wang¹,
Chengze Lv¹, Shuoyang Wang¹, and Zhenke Duan³(✉)

¹ School of Computer Science and Technology, Shandong Jianzhu University

² International School of Information Science&Engineering, Dalian University of Technology

³ Wuhan SecureScape Technology Co., Ltd
duanzhenke@sscaphwh.com

Abstract. Fire detection in dynamic environments faces continuous challenges, including the interference of illumination changes, many false detections or missed detections, and it is difficult to achieve both efficiency and accuracy. To address the problem of feature extraction limitation and information loss in the existing YOLO-based models, this study propose **You Only Look Once for Fire Detection with Attention-guided Inverted Residual and \textbf{D}ual-pooling Downscale Fusion (YOLO-FireAD)** with two core innovations: (1) **Attention-guided Inverted Residual Block (AIR)** integrates hybrid channel-spatial attention with inverted residuals to adaptively enhance fire features and suppress environmental noise; (2) **Dual Pool Downscale Fusion Block (DPDF)** preserves multi-scale fire patterns through learnable fusion of max-average pooling outputs, mitigating small-fire detection failures. Extensive evaluation on two public datasets shows the efficient performance of our model. Our proposed model keeps the sum amount of parameters (1.45M, 51.8% lower than YOLOv8n) (4.6G, 43.2% lower than YOLOv8n), and mAP75 is higher than the mainstream real-time object detection models YOLOv8n, YOLOv9t, YOLOv10n, YOLO11n, YOLOv12n and other YOLOv8 variants 1.3-5.5%.

Keywords: Fire detection, Efficient models, YOLO, Attention mechanisms, Small object detection.

1 Introduction

Fire disasters pose significant threats to human safety, ecological systems, and critical infrastructure worldwide. With increasing urbanization and climate change impacts, developing rapid and reliable fire detection systems has become a crucial research frontier in computer vision and public safety 1.

Traditional sensor-based approaches often suffer from limited coverage and environmental sensitivity, while emerging vision-based methods leveraging deep learning have demonstrated superior adaptability in complex scenarios 2. However, existing

convolutional neural network (CNN) architectures for fire detection face three fundamental challenges in practical deployments: 1) Severe performance degradation under illumination variations and fire-like object interference, 2) The trade-off between effectiveness and efficiency, and 3) Information loss during feature downsampling particularly detrimental to small fire detection.

Current YOLO-based solutions [3], while achieving notable real-time performance, exhibit critical limitations. The standard residual blocks in lightweight networks inadequately capture long-range contextual dependencies crucial for distinguishing fire features from complex backgrounds. Furthermore, conventional pooling operations in feature pyramid networks tend to discard spatial details during downsampling, leading to suboptimal performance in detecting early-stage fires with limited visual signatures. Recent attempts incorporating attention mechanisms or multi-scale fusion either introduce prohibitive computational overhead or fail to maintain feature integrity across scales [4,5,6,7].

To address these challenges, this study proposes YOLO-FireAD, a novel fire detection framework that synergistically integrates two specialized modules. Specifically, the contributions of this study are as follows.

1. This study proposes the **Attention-guided Inverted Residual Block (AIR)**, which synergistically embeds spatial-channel hybrid attention into inverted residual operations. This design achieves adaptive fire feature amplification through gated position-aware convolutions and channel reweighting, reducing false alarms from fire-like interference while maintaining computational efficiency (39% parameter reduction vs. baseline).

2. This study proposes the **Dual Pool Downscale Fusion Block (DPDF)** addresses small-fire detection failures through parallel max-average pooling fusion with learnable coefficients. This strategy preserves both flame edge details and smoke continuity during downsampling, improving mAP50 by 1.7% compared to standard YOLOv8n while reducing GFLOPs by 15%.

3. This study uses common evaluation metrics combined with cross-dataset testing (fire_detection dataset and fire dataset) for validation. Extensive experiments show that the model proposed in this study, YOLO-FireAD model, has a good precision-efficiency balance, achieving 34.6% mAP50-95 (1.8% higher than YOLOv8) with 51.8% parameter reduction.

The remainder of this paper is organized as follows: Section 2 reviews related work in fire detection and lightweight CNN architectures. Section 3 details the technical implementation of AIR and DPDF modules. Section 4 presents experimental setup and comparative analysis. Section 5 discusses practical implications and limitations, followed by concluding remarks in Section 6.

2 Methods

In this section, we offer a comprehensive description of each module within the network model, clarifying their specific functions. We also provide an overall explanation of the

model, detailing the involved components and structures, including the AIR block and DPDF block.

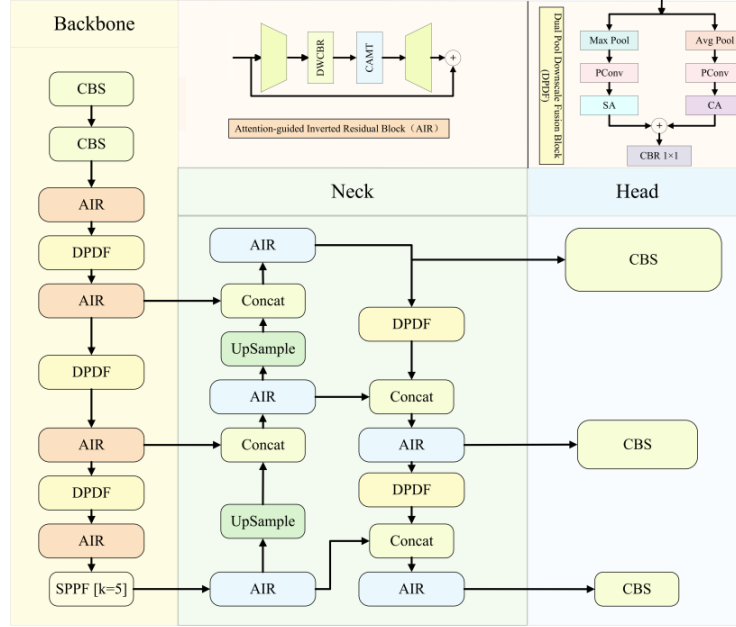


Fig. 1. Overview of the YOLO-FireAD model architecture.

2.1 Overview

As shown in Figure 1, YOLO-FireAD adopts a "Backbone-Neck-Head" three-tier architecture, which is optimised for early fire detection and high efficiency. Its core lies in combining the attention-guided feature enhancement of the AIR module and the multi-granularity feature retention mechanism of the DPDF module throughout the entire network, which improves the mAP50-95 of fire detection to 34.6% (1.8% higher than that of YOLOv8 8) while maintaining high efficiency (53% lower than the number of YOLOv8 parameters).

In the Backbone stage, based on the multi-stage feature extraction design, it consists of CBS base module, AIR attention module and DPDF downsampling module stacked alternately: the CBS module extracts the local texture features of the flame (such as edge flickering, smoke diffusion pattern) through 3×3 convolution; the AIR module embeds the attention mechanism, which enhances the response of the flame core region and suppresses the fire-like interference; and the DPDF module preserves the flame highlight details (maximum pooling) and smoke continuity (average pooling) simultaneously during resolution downscaling through a dual-pooling fusion strategy.

A bi-directional feature pyramid structure is used in the neck stage to fuse deep semantics with shallow details through UpSample and Concat operations. The DPDF module performs resolution alignment and noise filtering at the feature fusion nodes to

remove feature blurring due to smoke occlusion. The AIR module inserts cross-layer connectors to dynamically calibrate the attentional distributions of features at different scales.

The detection head, while maintaining multi-scale detection capability, achieves channel compression by stacking CBS modules and finally enters the loss function CIoU loss 9 as well as cross-entropy loss for localisation and classification.

2.2 Attention-guided Inverted Residual Block (AIR)

Aiming at the problem of confusion between fire-like interferences (such as light reflections) and real flames in flame detection, this study proposes the AIR module, which inherits and improves the classical inverted residual structure, and embeds the lightweight channel-spatial hybrid attention mechanism into the feature transformation process to achieve efficient feature discrimination through the three-stage design of feature decoupling-attention enhancement-lightweight compression. The module structure is shown in Figure 2.

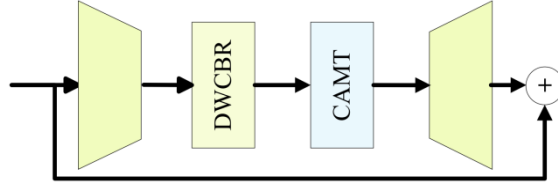


Fig. 2. Overview of the Attention-guided Inverted Residual Block structure.

For the input features $X_{in} \in R^{C \times H \times W}$, firstly, the channel dimension reduction is achieved by reducing the convolution to reduce the amount of calculation:

$$X_{reduce} = \text{ReLU} \left(\text{BN}(\text{Conv}_{1 \times 1}(X_{in})) \right). \quad (1)$$

The reduction rate $r = 0.25$ here is experimentally optimised to make a balance between reducing information loss and computational complexity. The deep feature extraction stage uses a deep separable convolution 10 with more efficient parameterisation:

$$X_{dw} = \text{ReLU} \left(\text{BN} \left(\text{DepthwiseConv}_{3 \times 3}(X_{reduce}) \right) \right). \quad (2)$$

This operation extracts local features via a spatial convolution kernel $k \times k$ while maintaining channel independence, reducing the computational complexity from $O(k^2 C_{exp}^2)$ to $O(k^2 C_{exp})$.

Then X_{dw} undergo convolutional additive self-attention 11 through the following process: First, the input features are decomposed into query (Q), key (K), and value (V) triples by 1×1 convolution:

$$[Q, K, V] = \text{Split}(\text{Conv}_{1 \times 1}(X)). \quad (3)$$

Then the parallel processing strategy of space-channel decoupling is adopted:

$$\hat{Q} = \mathcal{F}_{channel}(\mathcal{F}_{spatial}(Q)), \quad (4)$$

$$\hat{K} = \mathcal{F}_{channel}(\mathcal{F}_{spatial}(K)), \quad (5)$$

where the spatial operation $\mathcal{F}_{spatial}$ uses gated convolution to enhance position perception:

$$\mathcal{F}_{spatial}(x) = x \odot \sigma(\text{Conv}_{3 \times 3}(\text{DepthwiseConv}(x))). \quad (6)$$

The channel operation $\mathcal{F}_{channel}$ establishes global dependencies through adaptive pooling:

$$\mathcal{F}_{channel}(x) = x \odot \sigma(\text{MLP}(\text{GAP}(x))). \quad (7)$$

Then feature enhancement is achieved by additive fusion and deep convolution:

$$X_{out} = \text{Dropout}(\text{DepthwiseConv}_{3 \times 3}((\hat{Q} + \hat{K}) \odot V)), \quad (8)$$

where \odot denotes element-by-element multiplication, the design reduces $O(C^2)$ parameter overhead while maintaining feature interaction capabilities.

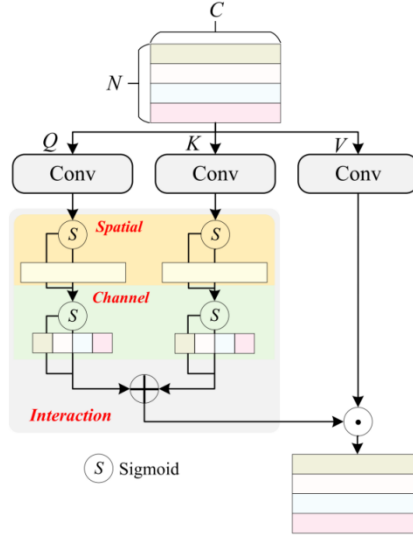


Fig. 3. Overview of the convolutional additive self-attention structure.

Finally, the number of channels is restored to facilitate subsequent feature extraction. The module structure is shown in Figure 3.

2.3 Dual Pool Downscale Fusion Block (DPDF)

For the problems of small target leakage and dynamic blurring in flame detection, the DPDF module achieves multi-granularity information fusion through a dual-path feature retention mechanism. As shown in Figure 4, the module synchronously preserves flame edge details (via maximum pooling) and smoke diffusion features (via average pooling) during resolution downsampling.

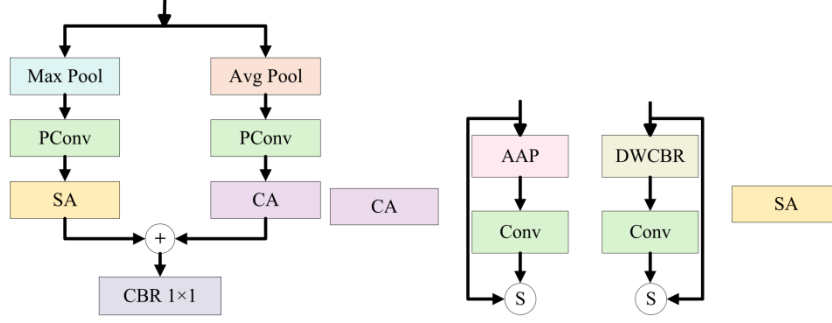


Fig. 4. Overview of the Dual Pool Downscale Fusion Block structure.

For dual path feature extraction, two pooling operations are first performed in parallel to obtain complementary features:

$$X_{max} = \text{MaxPool}_{2 \times 2}(X_{in}), \quad (9)$$

$$X_{avg} = \text{AvgPool}_{2 \times 2}(X_{in}). \quad (10)$$

The maximum pooling path reinforces the flame spike features (such as flickering highlight areas), and the average pooling path maintains the continuity of smoke diffusion.

Partial convolution (PConv) 12 is used to reduce the computational overhead:

$$X'_{max} = \text{PConv}(X_{max}) = \text{DepthwiseConv}_{3 \times 3}(X_{max}^{[r:]}) \oplus X_{max}^{[r:]}, \quad (11)$$

where $r=4$ is the channel reduction rate, performing deep convolution for only 1/4 of the channels and passing through the rest. This design reduces FLOPs by 68% while maintaining 90% feature accuracy.

Feature calibration is implemented via the SA (Spatial Attention) and CA (Channel Attention) modules:

$$X''_{max} = \text{CA}(\text{SA}(X'_{max})), \quad (12)$$

$$X''_{avg} = \text{CA}(\text{SA}(X'_{avg})). \quad (13)$$

Among them, the SA module enhances the flame edge response using cavity convolution, and the CA module applies channel reweighting based on the flame chromaticity histogram.

Finally, adaptive weighted fusion is achieved through learnable coefficients:

$$X_{fusion} = \alpha \cdot X''_{max} + (1 - \alpha) \cdot X''_{avg}. \quad (14)$$

3 Experimental setup

In this section, a brief overview of the experimental setup and related resources will be given. In the following, this study introduces the experimental dataset, the experimental configuration, and the evaluation metric in turn.

3.1 Dataset

The two datasets used in this study, fire_detection dataset and fire dataset, are both from the Paddle community and contain a large number of multi-scale complex flame fire scenes and smoke.

3.2 Configuration

The experimental program was executed on ubuntu22.04 operating system with NVIDIA GeForce RTX 4090 graphics card driver. The deep learning framework was selected as pytorch with 2.1.0, jupyter notebook was used for the compiler, Python 3.10 was used as the specified programming language, and all the algorithms used in the comparative analysis were operationally consistent and ran in the same computational setup. The image size was normalized to 640×640×3, and the batch size was 128, the optimizer was AdamW, the learning rate was set to 0.002, the momentum to 0.9, and the number of training periods was 200.

3.3 Evaluation metric

Model performance was evaluated using detection accuracy metrics (precision, recall, f1 score), comprehensive mAP scores (mAP50, mAP75, mAP50-95), and computational efficiency parameters (Params, GFLOPs, Model Size) for holistic assessment [13].

4 Experimental details and analyses

In order to validate the superior performance of the YOLO-FireAD detection model proposed in this paper, a series of validations are conducted on the above dataset and evaluated and analyzed using several evaluation metrics mentioned above.

Firstly, the current mainstream object detection models and their variants are introduced, and comparison experiments with the model proposed in this paper to demonstrate the superiority of the proposed model. Then, the results of the model pro-

posed in this paper are evaluated, including comparative experimental analysis, generalization experimental analysis, and ablation study analysis. Finally, the limitations of the model proposed in this paper are analyzed.

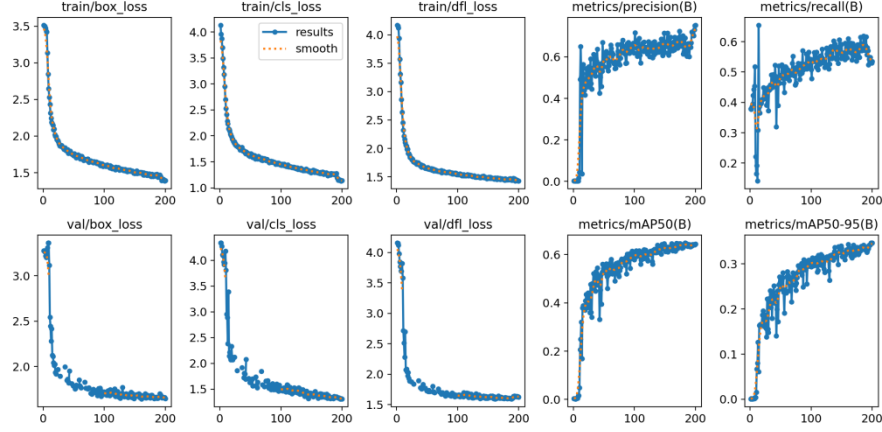


Fig. 5. YOLO-FireAD iterations on the fire_detection dataset.

Figure 5 records the changes of the bounding box loss, classification loss, focus loss, model precision, recall, and average precision for different thresholds in each iteration of the model on the training set and the test set. According to the charts, it can be seen that in the early stage of training, the model has high values for each loss and converges quickly, while the model metrics rise rapidly, indicating that the model is learning efficiently at this time. As the number of iterations increases, the rate of change of each curve gradually decreases, and after about 40 iterations, the curves begin to stabilise, and finally about 200 iterations, the curves basically remain stable.

4.1 Comparative Experimental Analysis

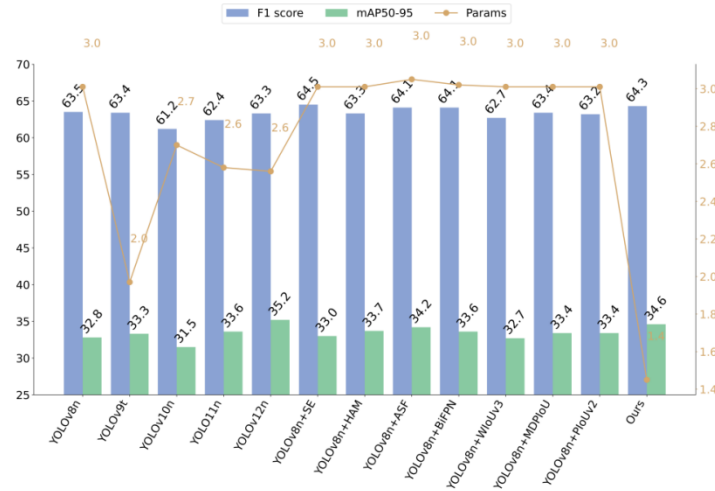


Fig. 6. Comparison of model metrics (F1 score, mAP50-95, Params) on fire_detection dataset.

Table 1. Test results of all models on fire_detection dataset.

Model	P	R	F1	mAP50	mAP75	mAP50-95	Params	GFLOPs	Size
YOLOv8n	64.6	62.4	63.5	64.2	31.9	32.8	3.01	8.1	6.1
YOLOv9t	66.8	60.3	63.4	64.4	31.8	33.3	1.97	7.6	4.6
YOLOv10n	68.2	55.5	61.2	60.0	30.0	31.5	2.70	8.2	5.6
YOLO11n	66.5	58.7	62.4	64.3	31.5	33.6	2.58	6.3	5.4
YOLOv12n	65.9	61.0	63.3	66.1	34.2	35.2	2.56	6.3	5.4
YOLOv8n+SE	70.4	59.5	64.5	63.5	31.4	33.0	3.01	8.1	6.1
YOLOv8n+HAM	3.3	63.3	63.3	64.8	32.7	33.7	3.01	8.1	6.1
YOLOv8n+ASF	69.1	59.8	64.1	65.1	32.8	34.2	3.05	8.6	6.2
YOLOv8n+BifPN	67.6	60.9	64.1	63.4	32.3	33.6	3.02	8.1	6.1
YOLOv8n+WIoUv3	63.9	61.5	62.7	63.4	31.7	32.7	3.01	8.1	6.1
YOLOv8n+MDPIoU	65.4	61.3	63.4	63.8	32.2	33.4	3.01	8.1	6.1
YOLOv8n+PIoUv2	66.1	60.3	63.2	63.9	33.8	33.4	3.01	8.1	6.1
Ours	75.3	53.3	64.3	64.9	35.5	34.6	1.45	4.6	3.3

In order to verify the performance of YOLO-FireAD, the comparison experiment compares YOLO-FireAD with YOLOv8n 8, YOLOv9t 14, YOLOv10n 15, YOLO11n 8, YOLOv12n, and the improved models of YOLOv8n 16,17,18,19,20,21,22 on fire_detection dataset, and analyse the model performance in terms of P, R, F1score, mAP50, mAP50-75,mAP50-95, Params, FLOPs, Model Size to analyse the model performance,

the experimental results are shown in Table 1, and Figure 6 is the visual graph of the experimental data of F1, mAP50, Params. It can be seen that the YOLO-FireAD model performs well in all performance indicators, and its P, R, F1, and mAP (mAP50, mAP75, and mAP50-90) are significantly better than the other models. The advantage of P is the most obvious, which is 10.7%, 9.4%, 12%, and 11.4% higher compared with YOLOv8n, YOLOv12n, YOLOv8n+HAM, and YOLOv8n+WIoUv3, respectively. In addition, YOLO-FireAD's mAP75 performs the best among all models, which indicates that YOLO-FireAD maintains high detection performance at different IoU thresholds, which gives it a strong advantage in complex bad detection tasks. While maintaining high accuracy, the complexity of YOLO-FireAD is also significantly lower than other models, with Params of 1.45M, which is only 48.2% of YOLOv8n, GFLOPs of 4.6G, and Model Size of 3.3MB, which are the lowest among all the compared models.

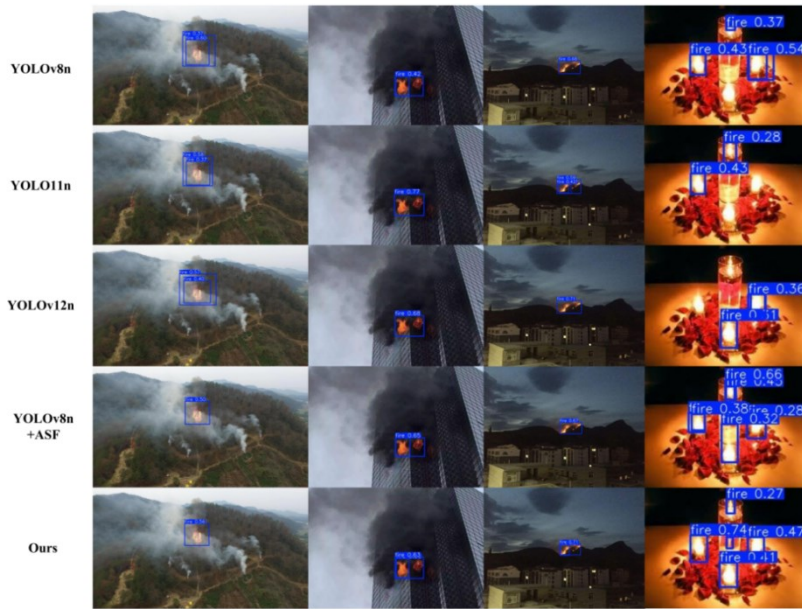


Fig. 7. Sample visualisation of several competitive model identification results.

Figure 7 shows several representative sets of detection results of several mainstream models and YOLO-FireAD in actual situations, in order to clearly see the differences in performance of different models in flame hazard detection. YOLOv8n is able to identify flames in most of the scenarios, but the recognition accuracy is slightly lower, and there is a certain leakage of detection in dense flame scenarios. Compared with YOLOv8n, YOLOv11n and YOLOv12n have higher leakage rates in dense flames, which indicates that there are some limitations in their flame recognition in complex scenarios. Compared to the above three models YOLOv8n+ASF has a significantly lower leakage rate when the flame is dense, but there is a certain leakage for small flames, which indicates that it has some limitations when dealing with small-sized objects. Compared with the above models, YOLO-FireAD performs significantly better

with the attention mechanism of its AIR module, which can dynamically adjust the weights of the feature maps to make the network more focused on the flame targets, especially in multi-dense flame scenes. The DPDF module synchronously retains the flame highlight details during resolution downscaling through the dual-pooling fusion strategy, avoiding the loss of the small flame features during downsampling, thus achieving a more accurate recognition of small flames.



Fig. 8. Heat map visualisation results, the first row is the original image, the second row is the YOLOv8n visualisation results and the third row is the YOLO-FireAD visualisation results.

Figure 8 shows the heat map visualisation results of YOLOv8n and YOLO-FireAD in the flame recognition task. From the figure, it can be seen that although YOLOv8n is able to identify most of the flames, there are some false detections and the identification of the flames is not obvious enough, and there are leakage detections. In contrast, YOLO-FireAD performs much better, accurately identifying the location of flames and more clearly identifying small flames.

Furthermore, we verify the generalization of the model on the Fire dataset, and the results are shown in Table 2. It can be seen that YOLO-FireAD also performs well. The mAP50-95 has the best performance among all the comparison models, reaching 8.12%, which is higher than YOLOv9t, YOLOv10n, YOLOv8n+HAM, YOLOv8n+WIoUv3, 1.39%, 1.13%, 1.37%, and 2.25%, respectively.

In order to further demonstrate the effectiveness of the two modules proposed in YOLO-FireAD, we conducted ablation experiments on the fire_detection dataset with the AIR, DPDF modules using YOLOv8n as the baseline algorithm, and the results of the experiments are shown in Table 3. Although the mAP50-95 of YOLO-FireAD decreased slightly after only introducing the AIR module, P increased from 64.6% to 68.0%, an increase of nearly 3.4%. At the same time, Params, FLOPs and Model Size were significantly reduced, especially Params was reduced from 3.01 to 1.84. That's a reduction of about 39 percent. This proves the role that AIR plays in lightweight.

When only the DPDF module is introduced into YOLO-FireAD, compared with YOLOv8n, the P, R, mAP50, mAP50-95 of the model are all improved, among which

the improvement of P is the most obvious, from 64.6% to 69.2%, and the reductions of Params, FLOPs, and Model Size are also achieved. This proves the usefulness of the DPDF module in improving the model recognition accuracy.

Table 2. Test results of generalisability of all models verified at Fire dataset

Model	P	R	mAP50	mAP50-95
YOLOv8n	37.0	28.8	25.6	7.83
YOLOv9t	29.8	28.5	22.4	6.73
YOLOv10n	34.4	29.4	22.7	6.99
YOLO11n	40.7	31.6	27.2	7.85
YOLOv12n	39.7	27.1	24.5	8.00
YOLOv8n+SE	43.4	22.5	22.8	7.24
YOLOv8n+HAM	35.5	28.6	22.5	6.75
YOLOv8n+ASF	39.0	28.3	25.9	7.48
YOLOv8n+BifPN	38.7	31.5	25.8	7.80
YOLOv8n+WIoUv3	33.1	28.0	21.4	5.87
YOLOv8n+MDPIoU	32.1	28.9	21.6	7.23
YOLOv8n+PIoUv2	35.0	29.0	23.6	7.22
Ours	35.9	25.7	25.0	8.12

Table 3. Test results of YOLO-FireAD ablation experiment on fire_detection dataset.

AIR	DPDF	P	R	mAP50	mAP50-95	Params	GFLOPs	Size
		64.6	62.4	64.2	32.8	3.01	8.1	6.1
✓		68.0	56.3	63.0	32.6	1.84	5.4	3.9
	✓	69.2	63.8	65.9	34.5	2.52	6.9	5.2
✓	✓	75.3	53.3	64.9	34.6	1.45	4.6	3.3

When the two are combined, the overall model performance is optimal, with P further improved to 75.3%, which is 10.7% higher compared to YOLOv8n, while mAP50-95 is improved from 32.8% to 34.6%. Params, FLOPs, and Model Size are all reduced to 1.45M, 4.6G, and 6.1MB respectively, which can be inferred that both AIR and DPDF play an indispensable role in YOLO-FireAD. The combination of the two makes YOLO-FireAD able to significantly increase the detection rate while guaranteeing the accuracy of the model through the enhancement of key features and the optimal fusion of multi-granularity features.

4.2 Limitation Analysis

Although the YOLO-FireAD model performs well in flame detection tasks, it still has some limitations. For example, although the AIR module can effectively enhance the response in the flame core region, it may misdetect fire-like interferences in the background (such as light reflections, glare areas, etc.) under complex backgrounds or extreme lighting conditions, leading to a decrease in detection accuracy. In addition, although the DPDF module performs well in multi-granularity feature fusion, when dealing with dynamic scenes (e.g., fast moving flames or fast spreading smoke), inaccurate feature fusion may occur, which affects the detection performance of dynamic flames. Meanwhile, due to the relatively small number of samples in the training data for some specific scenes (such as small target flames or smoke occlusion scenes), the model's detection capability in these scenes may be somewhat limited.

5 Conclusion

This study presents YOLO-FireAD, an efficient fire detection framework that addresses critical challenges in real-world deployment through two novel modules: the Attention-guided Inverted Residual Block (AIR) and Dual Pool Downscale Fusion Block (DPDF). AIR integrates a lightweight channel-spatial hybrid attention mechanism into inverted residual learning, enabling adaptive feature recalibration while suppressing fire-like interference. DPDF mitigates feature degradation during downsampling by fusing max-pooling's edge-preserving capability and average-pooling's contextual consistency. Extensive experiments across two benchmarks demonstrate the framework's superiority, outperforming YOLOv8 by 1.8% mAP with 51.8% fewer parameters.

However, limitations persist in extreme lighting conditions and dynamic fire-spread scenarios due to residual feature ambiguity. Future work will focus on two directions: Integrating infrared-visible multimodal data to enhance detection reliability in smoke-occluded scenarios, and optimizing cross-platform deployment for heterogeneous edge devices.

References

1. T. Wang, J. Wang, C. Wang, et al., Improving YOLOX network for multi-scale fire detection, *Vis. Comput.*, vol. 40, pp. 6493–6505, 2024. [doi: 10.1007/s00371-023-03178-1]
2. L. Cao, Z. Shen, and S. Xu, Efficient forest fire detection based on an improved YOLO model, *Vis. Intell.*, vol. 2, p. 20, 2024. [doi: 10.1007/s44267-024-00053-y]
3. W. Pan, X. Wang, and W. Huan, Real-time dynamic scale-aware fusion detection network: take road damage detection as an example, *J. Real-Time Image Proc.*, vol. 22, p. 55, 2025. [doi: 10.1007/s11554-025-01634-w]
4. W. Liu, Z. Shen, and S. Xu, CF-YOLO: a capable forest fire identification algorithm founded on YOLOv7 improvement, *SIViP*, vol. 18, pp. 6007–6017, 2024. [doi: 10.1007/s11760-024-03288-w]

5. X. Cao, J. Wu, J. Chen, et al., Complex Scenes Fire Object Detection Based on Feature Fusion and Channel Attention, *Arab J. Sci. Eng.*, 2024. [doi: 10.1007/s13369-024-09471-y]
6. J. Yao, J. Lei, J. Zhou, et al., FG-YOLO: an improved YOLOv8 algorithm for real-time fire and smoke detection, *SIViP*, vol. 19, p. 346, 2025. [doi: 10.1007/s11760-025-03894-2]
7. W. Pan, X. Wang, and W. Huan, EFA-YOLO: An Efficient Feature Attention Model for Fire and Flame Detection, *arXiv*, 2024. [eprint: 2409.12635].
8. G. Jocher, J. Qiu, and A. Chaurasia, Ultralytics YOLO (Version 8.0.0) [Computer software], 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
9. Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, and D. Ren, Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression, *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 12993–13000, 2020. [doi: 10.1609/aaai.v34i07.6999].
10. F. Chollet, Xception: Deep Learning with Depthwise Separable Convolutions, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1800–1807. [doi: 10.1109/CVPR.2017.195]
11. T. Zhang, L. Li, Y. Zhou, W. Liu, J.-N. Hwang, and X. Ji, CAS-ViT: Convolutional Additive Self-attention Vision Transformers for Efficient Mobile Applications, *arXiv*, 2024. [eprint: 2408.03703].
12. J. Chen, S.-h. Kao, H. He, W. Zhuo, S. Wen, C.-H. Lee, and S.-H. G. Chan, Run, Don't Walk: Chasing Higher FLOPS for Faster Neural Networks, in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 12021–12031. [doi: 10.1109/CVPR52729.2023.01157]
13. W. Chen, J. Luo, F. Zhang, et al., A review of object detection: Datasets, performance evaluation, architecture, applications and current trends, *Multimed. Tools Appl.*, vol. 83, pp. 65603–65661, 2024. [doi: 10.1007/s11042-023-17949-4]
14. C.Y. Wang, I.H. Yeh, and M.H.Y. Liao, YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information, in *Computer Vision – ECCV 2024*, A. Leonardis, E. Ricci, S. Roth, O. Russakovsky, T. Sattler, G. Varol, Eds., *Lecture Notes in Computer Science*, vol. 15089, Springer, Cham, 2025. [doi: 10.1007/978-3-031-72751-1_1]
15. A. Wang, H. Chen, L. Liu, K. Chen, Z. Lin, J. Han, and G. Ding, YOLOv10: Real-Time End-to-End Object Detection, in *Advances in Neural Information Processing Systems*, vol. 37, 2024, pp. 107984–108011.
16. J. Hu, L. Shen, and G. Sun, Squeeze-and-Excitation Networks, in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141. [doi: 10.1109/CVPR.2018.00745]
17. G. Li, Q. Fang, L. Zha, X. Gao, and N. Zheng, HAM: Hybrid attention module in deep convolutional neural networks for image classification, *Pattern Recognition*, vol. 129, p. 108785, 2022. [doi: 10.1016/j.patcog.2022.108785].
18. M. Kang, C.-M. Ting, F.F. Ting, and R.C.-W. Phan, ASF-YOLO: A novel YOLO model with attentional scale sequence fusion for cell instance segmentation, *Image and Vision Computing*, vol. 147, p. 105057, 2024. [doi: 10.1016/j.imavis.2024.105057].
19. Z. Tong, Y. Chen, Z. Xu, and R. Yu, Wise-IoU: Bounding Box Regression Loss with Dynamic Focusing Mechanism, *arXiv*, 2023. [eprint: 2301.10051].
20. S. Ma and Y. Xu, MPDIoU: A Loss for Efficient and Accurate Bounding Box Regression, *arXiv*, 2023. [eprint: 2307.07662].
21. C. Liu, K. Wang, Q. Li, F. Zhao, K. Zhao, and H. Ma, Powerful-IoU: More straightforward and faster bounding box regression loss with a nonmonotonic focusing mechanism, *Neural Networks*, vol. 170, pp. 276–284, 2024. [doi: 10.1016/j.neunet.2023.11.041].



2025 International Conference on Intelligent Computing

July 26-29, Ningbo, China

<https://www.ic-icc.cn/2025/index.php>

22. J. Doherty, B. Gardiner, E. Kerr, and N. Siddique, BiFPN-YOLO: One-stage object detection integrating Bi-Directional Feature Pyramid Networks, *Pattern Recognition*, vol. 160, p. 111209, 2025. [doi: 10.1016/j.patcog.2024.111209].